ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

DEPARTMENT OF MANAGEMENT SCIENCE AND TECHNOLOGY

CORRELATION MODELLING WITH APPLICATION TO RISK MANAGEMENT

Thesis by

Chrysi E. Markopoulou

A dissertation submitted for the degree of

Doctor of Philosophy

Supervisory Committee:
Professor Apostolos P. Refenes
Professor Christos Giannikos
Professor Diomidis Spinellis

Athens, July 2014

Acknowledgements

I would like to express my sincere appreciation and gratitude to my supervisor and mentor, Professor Apostolos Refenes, for his continuous support, guidance and inspiration throughout the years, from my MBA degree to the completion my PhD dissertation. I am thankful to the members of my supervisory committee, Professor Diomidis Spinellis and Professor Christos Giannikos for their insightful comments. I would also like to thank Professor Dionysis Chionis, Professor Achilleas Zapranis, Assistant Professor Dimitrios Psychoyios and Assistant Professor Vasiliki Skintzi who honored me by joining my evaluation committee.

I feel particularly indebted to Vasiliki Skintzi for the everlasting fruitful discussions, continuous guidance and unwavering encouragement. Furthermore, I would like to thank Dr. Spyros Xanthopoulos for his valuable advices during the first years of my studies.

Finally, I would like to express my deep gratitude to my parents and Panos. My parents (Eythimios and Konstanina) and my brothers (Vassilis and Antonis) have taught me the preciousness of working hard and that nothing worth having comes easy. Finally, I feel grateful to Panos who has shown endless support, understanding and encouragement.

Financial Support

This research has been co-financed by the European Union (European Social Fund – ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: Heracleitus II. Investing in knowledge society through the European Social Fund.

CORRELATION MODELLING WITH APPLICATION TO RISK MANAGEMENT

Abstract

Accurate estimation and prediction of correlation is of paramount importance in asset allocation, risk management and hedging applications, particularly in light of recent studies that provide evidence of increased correlation during periods of high volatility, leading to diminishing diversification benefits in states of nature that are most needed. The time-variability of the correlation process has fuelled extensive literature on dynamic correlation modelling. In an attempt to depart from correlation estimation based on projections from historical data, two alternative measures of correlation, namely the implied and the realized correlation, have been proposed in the recent literature. Remarkably, in contrast to volatility estimates, existing studies on the informational efficiency and forecasting performance of respective correlation measures are rather limited.

This thesis focuses on exploring the dynamics that govern the evolution of correlation risk premium and its components, namely implied and realized correlation, and assessing the impact of predictability to portfolio allocation, hedging and trading decisions. First, the time-variation and certain distribution characteristics of the correlation risk premium, defined as the difference of realized and implied correlation, are examined. The information content of market –specific and macroeconomic variables, which have been previously reported as proxies of the business cycles, in predicting future premium is also evaluated. Secondly, a model-free measure of implied correlation is proposed and the question of predictable dynamics in the evolution of the series is investigated both in statistical and economic terms. A trading strategy designed to exploit daily changes of the series sets the foundation for addressing the efficient market hypothesis. Finally, based on the distributional properties of realized volatility,

correlation and hedge ratio, an alternative forecasting methodology is applied to predict the realized hedge ratio and to explore the additive value of intraday data in a dynamic hedging context while the hedging performance is compared in terms of portfolio optimization and risk management.

The thesis has reached a number of conclusions. First, correlation and correlation risk premium vary substantially over time and increase sharply during turbulent periods, while culminated during the Asian and Russian financial crisis in 1997-1998 and the subprime mortgage crisis of 2007-2009. The previously documented correlation risk premium is no longer significant during the recent 2007 – 2009 crisis, suggesting the disappearance of arbitrage opportunities. Secondly, the predictability of model-free implied correlation series suggested by statistical measures cannot be exploited in terms of economic gains, suggesting that the S&P 100 options market is efficient. Finally, forecasting the dynamics of the realized hedge ratio directly reveals predictable patterns in the evolution of the hedge ratio, resulting in improved hedging performance, in terms of both economics gains and risk measures.

Papers under review

1. Markopoulou C., Skintzi V., Refenes A. (2014). Realized Hedge Ratio: Predictability and Hedging Performance.
2. Markopoulou C., Skintzi V., Refenes A. (2012). On the Predictability of Model-free Implied Correlation.

Table of contents

List of Tables

List of Figures

List of abbreviations

Abbreviations	
ADF	Augmented Dickey-Fuller
CRP	Correlation Risk Premium
DCC	Dynamic Conditional Correlation
ECM	Error Correction Model
EEP	Early Exercise Premium
EMH	Efficient Market Hypothesis
ES	Expected Shortfall
HAR	Heterogeneous Autoregressive
ICI	S&P 500 Implied Correlation Index
KPSS	Kwiatkowski-Phillips-Schmidt-Shin test
MAE	Mean Absolute Error
MCP	Mean Correct Prediction
MDM	Modified Diebold-Mariano test
MFIC	Model-Free Implied Correlation
MFIV	Model-Free Implied Volatility
PP	Phillips-Perron test
RCov	Realized Covariance
RC	Realized Correlation
RMSE	Root Mean Squared Error
RMVHR	Realized Minimum Variance Hedge Ratio
RS	Regime Switching
RV	Realized Volatility
VaR	Value-at-Risk
VRP	Volatility Risk Premium

Chapter 1

Introduction

1.1 Motivation of the Thesis

Correlation is central to most financial applications, including asset and derivatives pricing, asset allocation, risk management and hedging. It is thus not surprising that significant research effort has been devoted in the modelling and the forecasting of the correlation distribution. The recent financial crisis, which inaugurated in the U.S.A in 2007 and propagated globally the years that followed, has highlighted the vital importance of accurate estimation and forecasting of volatility and correlation. Although volatility has stirred the research interest, respective research in the correlation context is notably less intense. This thesis focuses on the modelling and predictability of alternative measures of correlation.

While early studies considered the second moments of distribution to be constant over time, later advances recognized the time-variability of correlation. A number of studies have discussed the asymmetry in correlation distribution and provided evidence of increased correlation during periods of higher volatility resulting to an increase of the aggregate risk borne by investors and diminished diversification benefits. (See Erb et al. 1994; Ang & Chen, 2002; Hong et al. 2007; Campbell, Koedijk and Kofman, 2002; Longin & Solnik, 2001, amongst others)

It is therefore natural to consider that investors would desire to protect themselves from unexpected escalation of correlation and are willing to pay an additional premium for the assets that perform well during high correlation and low diversification benefits states of nature. Correlation risk has been the subject of several recent studies that have unanimously provided evidence of a negative price, thus affecting the risk-return tradeoff and the market price of equity returns. Interestingly, market participants have recognized the possibility of trading correlation risk. The dispersion strategy and the correlation swap are two novel trading products, designed to exploit the differential pricing of index and individual options and the correlation risk premium. In response to the increased interest, the CBOE launched in July 2009 the S&P 500 Implied Correlation Index (ICI). The ICI is a measure of the average correlation of the S&P 500 index components, as implied by the option prices on the S&P 500 index options and the prices of the 50 largest individual stock options based on their capitalization.

Since correlation is an unobserved component of the asset return distribution, an important issue that arises is the modelling of correlation. In the univariate context, the seminal work of Engle (1982) on the autoregressive conditional heteroskedasticity (ARCH) model has laid the groundwork for the extensive volatility modelling literature that followed. Based on the ARCH specification, several alternative models that attempt to capture well-documented traits of the volatility process, for instance long-memory and asymmetry, have been proposed in the literature. More recently, researchers have focused on the extension of these models from the univariate to the multivariate dimension (see Bauwens et al., 2006 for a detailed review). However, the critical assumption of positive definiteness of the variance-covariance matrix along with imposed restrictions of the optimization of the log-likelihood function complicates the estimation procedure of these models. Moreover, stochastic and multivariate GARCH models rely on the use of historical information as the relevant information set to forecast future correlation.

In an attempt to depart from model complexity and restrictions, the research interest has shifted to alternative approaches of modeling the second moment of asset distributions that do not rely on the historical dataset. Among the prevailing methodologies is the measure of implied correlation, deduced from currently traded option prices, and the measure of realized correlation, computed from high-frequency returns.

Inherently, option prices reflect the current market view of future price movements of the underlying stock thus containing valuable information regarding the market forecast of future correlation. In contrast to implied volatility, which has attracted vast research interest, implied correlation has received considerably less attention over the past years. Only a limited number of studies have examined the notion of implied correlation in the context of various asset classes[1]. Longstaff et al. (2001), De Jong et al. (2004) and Han (2007) derive interest rate correlations implied from options in fixed income markets. Skintzi & Refenes (2005) use option price data on the Dow Jones Industrial Average Index and component stocks to derive an ‘implied correlation index’ that measures the average portfolio diversification.

In another stream of literature, the increased availability of multi-dimensional high frequency data has given rise to a new area of research in modeling and forecasting the realized volatility and covariance. Andersen and Bollerslev (1998) introduced the notions of realized volatility and covariance as model-free estimators of the true latent process. In essence realized volatility is defined as the sum of squared intraday asset returns and realized covariance as the cross product of returns. Realized correlation is directly deducted from realized covariance and volatility. Andersen, Bollerslev, Diebold and Labys (2001) (ABDL, hereafter), Andersen, Bollerslev, Diebold and Ebens (2001) (ABDE, hereafter) and Barndorff-Nielsen & Shephard (2002) showed that, according to the theory of quadratic variation, as the sampling frequency tends to infinity, the realized measures are unbiased and efficient estimators of the integrated processes, which essentially become observable, thus enabling direct estimation. The notable advantage of the realized measures is that they are essentially “model-free”, without relying on any parametric model that induces econometric or mathematical misspecification. Thereafter, the use of non-parametric measures of realized volatility and correlation in risk management, asset allocation and derivatives hedging applications has propagated.

The information content of implied and realized correlation in providing accurate forecasts of future ex-post realized correlation has been at the core of several studies.

[1] See Christoffersen et al. (2012) for a review of studies using option-implied information to derive moments, correlation and density forecasts.

Siegel (1997), Campa & Chang (1998) and Lopez & Walter (2000) study the forecasting performance of implied correlations derived from foreign exchange options and find evidence of superior forecasting performance when compared to historical correlation. In a similar context, Pong et al. (2004) find that the forecasting performance of realized volatility is closely followed by the implied counterpart.

An interesting question that arises, yet distinct from the previously mentioned, is whether the evolution of implied correlation per se is predictable. Forecasting the dynamics of correlation can be a useful tool in asset pricing and portfolio allocation. The presence of exploitable patterns in the dynamics of the series casts calls into question the efficient market hypothesis. In particular, the options market efficiency is addressed through either the assessment of the no-arbitrage principle or the exploitability of trading strategies in terms of economic gains (Jensen, 1978). The second approach has been widely used in the implied volatility context while the empirical results remain mixed. In particular, several studies (see Harvey & Whaley, 1992, Guo 2000, Brooks & Oozeer, 2002 and Konstantinidi, Skiadopoulos and Tsagkaraki, 2008, amongst others) provide evidence of significant predictability in implied volatility using time-series models, which, however, cannot be exploited to achieve significant economic profits. In contrast, Chiras and Manaster (1978) and Goyal & Saretto (2009) provide evidence of statistically and economically significant predictability in the dynamics of volatility implied in at-the-money stock option prices. Remarkably, despite the plethora of studies in the predictability of volatility, related literature in forecasting correlation dynamics is still in its infancy.

From a risk management perspective, accurate correlation modelling and forecasting is of paramount importance for hedging decisions. While earlier studies considered the hedge ratio to be constant over time, following the recognition of time variability in the futures and spot returns distributions, a variety of models, from the generalized autoregressive conditionally heteroskedastic (GARCH) family as well as models that capture the long-run dependencies of the volatility process, have been employed to model the hedge ratio. Despite the well-documented superior forecasting ability of realized measures as opposed to estimators from daily or lower frequency data, not much attention has been paid to the informational efficiency of high frequency dataset to forecast the hedge ratio. Among the limited amount of studies, Lai and Sheu (2008 and 2010) and Shen and Stoja (2010) have provided evidence of increased hedging effectiveness when intraday data are employed to model the hedge ratio.

1.2 Contribution of the Thesis

This thesis is primarily focused on the pricing of correlation risk and the information content of alternative correlation measures in optimizing portfolio allocation, risk management and hedging decisions. Firstly, the distributional characteristics of the correlation risk premium and the informational content of volatility risk premium and macroeconomic variables are assessed. Secondly, an improved, non-parametric measure of implied correlation measure inferred from options' daily closing prices is proposed and the question of market efficiency is addressed. Thirdly, the superior informational content of high frequency data in optimizing hedging and risk management decisions is examined.

A number of studies have provided evidence of increased correlation during periods of higher volatility and low returns. An increase of market-wide correlation affects the investors' risk-return tradeoff, resulting in diminished diversification benefits in states of nature that are most needed. It is thus natural to consider that, depending on the relevant degree of risk aversion, an investor will require an additional compensation for undertaking the risk of correlation fluctuations. While a significant number of studies have focused on modelling the risk premium of volatility and examining the significance of several factors in forecasting it, understanding the dynamics of the risk premium of correlation is relatively poor, despite the fact that it is equally important.

This thesis provides a thorough assessment of the distributional properties of the risk premium of correlation and the informational content of several macroeconomic and market-specific variables in forecasting the dynamic evolution of the series. In addition, the evolution of correlation risk premium has been examined for the presence of several stylized facts, previously documented to affect the correlation process. The correlation risk premium is defined as the difference of realized and implied correlation. The former has been estimated from options' daily closing prices. For the latter, I propose a model-free estimate of market-wide correlation and diversification benefit. The Model-Free Implied Correlation (MFIC) is a refined estimate of implied correlation that, in contrast to previous studies and the official CBOE S&P 500 ICI, is firstly, derived from currently observed option prices without relying on any specific option valuation model and secondly, corrected for the early exercise premium. For comparison purposes, I also examine the evolution of the volatility risk premium for the index and the individual options and study the relationship with the correlation risk premium. An extensive dataset allows the evaluation of the time-varying characteristics of the correlation process during several periods of financial turmoil, and most importantly during the subprime mortgage crisis of 2007-2009 that was later transmitted globally.

The second contribution of the thesis lies in the assessment of predictable patterns in the evolution of the model-free implied correlation series. Despite the significant amount of studies on evaluating the performance of option-implied information on various financial applications, only a limited number of studies have dealt with the predictability of option-implied measures, while the vast majority of those focus on volatility. From a practical perspective, the presence of predictability in the evolution of the series enables market participants to form profitable trading strategies, calling into question the market efficiency hypothesis.

This thesis performs a comprehensive study of the dynamics and predictable patterns in the evolution of the implied correlation series per se. For this purpose, I propose the modelling of the correlation process evolution through alternative time-series specifications that capture different aspects of the distributional characteristics, for instance ARFIMA and HAR models are employed to model the long-memory properties while a regime switching attempt to capture the asymmetry in correlation. In addition, the study assesses the performance of alternative methods of combining forecasts, which essentially accumulate forecasts derived from the individual candidate models that have access to different information sets. I find that under several statistical measures, the evolution of the model-free implied correlation series is predictable, thus revealing exploitable inefficiency of the S&P 100 market. Finally, I develop a trading strategy that exploits daily changes in the implied correlation series and I access the economic significance of obtained forecasts. I show that, in the absence of transaction costs, the existence of predictable patterns in the dynamics of the model-free implied correlation series, suggested by statistical measures, can be further exploited to generate abnormal profits. However, when market frictions are introduced, the profitability of the proposed trading strategy is eliminated, suggesting that the efficiency hypothesis fails be rejected for the S&P 100 options market.

The third contribution of the thesis lies in the modelling and forecasting of realized hedge ratio from high frequency data and the assessment of comparative improvement over conventional methods that utilize daily data. The hedge ratio denotes the number of futures contracts that the hedger should hold in order to hedge one unit risk of the underlying spot market. The majority of the studies focus on the modelling of the variance-covariance matrix of spot and futures returns and ultimately, the estimation of the optimal hedge ratio. Andersen et al. (2006) defined realized beta as the ratio of realized covariance and variance and find that the well-documented long memory traits of the individual processes are eliminated when forming the beta ratio.

The contribution of this thesis in the hedging context is threefold. First, motivated by the methodology and suggested findings on realized beta, I assess and analyze thoroughly the differential distributional properties of realized variance, covariance and hedge ratio. Second, intrigued by the findings of the distributional analysis, I propose the employment of alternative time-series specifications to model and forecast the evolution of the realized hedge ratio per se. The forecasting performance of employed econometric models is compared to the forecasts obtained from random walk. I find that the hypothesis of efficient spot and futures markets of the EUR/USD and the GBP/USD exchange rates as well as the S&P 500 and the FTSE 100 indices is strongly rejected. Finally, I compare the proposed modelling methodology of the realized hedge ratio to widely used methodologies of deriving hedge ratios from daily returns and I find superior forecasting performance of the proposed modelling approach in risk management and portfolio optimization decisions. Obtained results hold across different asset classes but are more conspicuous in the case of stock indices.

1.3 Overview of the Thesis

A brief overview of the thesis is outlined below. Chapter 2 presents a thorough literature review on the modelling of the correlation process and its application on several aspects of financial theory, from the minimum-variance portfolio optimization theory and the Capital Asset Pricing Model (CAPM) to risk management, asset allocation and hedging decisions. To this end, the Chapter includes a detailed survey of econometric specifications that attempt to capture several aspects of the dynamics that govern the evolution of the process. In addition to traditional estimators, two alternative measures of correlation are presented, namely implied and realized correlation. A large body of literature supports the informational efficiency of the measures in option pricing, asset pricing and hedging applications. Secondly, the Chapter presents the fundamental concept of market efficiency and provides a detailed literature on relevant studies that have examined the efficiency of the options and the futures markets from alternative perspectives. Accurate correlation modelling and forecasting is of utmost importance in the hedging decisions, with the futures hedging being one of the most simple and widespread hedging strategies. A review of alternative hedge ratios and estimation methods is presented at the last Section of the Chapter.

Chapter 3 discusses the distributional properties of the correlation risk premium and several stylized facts of correlation distribution. First, the methodology for the construction of the model-free implied correlation series and the estimation of the early exercise premium is thoroughly described. In addition, the presence of long memory and asymmetric response to equity returns in the correlation risk premium is examined. Moreover, a Granger causality test is applied to test whether correlation risk premium drives volatility risk premium of the index, or vice versa. Finally, the informational content of an economic determinants model in forecasting correlation risk premium is assessed.

Chapter 4 examines the predictability of the proposed model-free implied correlation measure. A thorough evaluation of the time-series evolution of MFIC is performed and the existence of predictable patterns in the dynamics of the series is examined. The out-of-sample forecasts of the series are firstly assessed with statistical measures that gauge the prediction error in terms of both magnitude and directional efficiency. The fundamental theoretical background of the proposed trading strategy and the implementation steps are described thoroughly in this Chapter. The statistical measures provide strong evidence in favour of existing predictable pattern in the S&P 100 option market. However, results from the trading strategy suggest that statistically significant profits can be generated only in the absence of transaction costs.

In Chapter 5, the dynamics of the realized hedge ratio and the hedging performance is examined. Initially, the distributional properties of realized variance, realized covariance and realized hedge ratio are examined in detail. Following the evidence provided from the statistical analysis, a number of econometric models employed to forecast directly the realized hedge ratio series are thoroughly discussed. Statistical measures are employed to assess the dynamic properties and predictability of the realized minimum variance hedge ratio series. The out-of-sample hedging effectiveness of the realized hedge ratio is compared to the hedge ratio estimated from daily returns, namely constant and rolling OLS, Error-Correction Model as well from a model where the returns follow a DCC-GARCH model. Finally, the impact of transaction costs in dynamic hedging strategy as well as of different sampling frequencies in the derivation of the realized hedge ratio is assessed. The main findings suggest that forecasting directly the realized hedge ratio leads to marginal increase of the percentage of risk reduction while the hedger’s benefit is substantial when both the average return and the variance of the hedge portfolio is taken into account (as measured by the Sharpe ratio and expected utility). The results hold across the different asset classes, although, as expected, the benefits are lower in the case of exchange rates.

Section 6 summarizes the issues addressed along with the main conclusions of this thesis and discusses possible topics for future research.

Chapter 2

Literature Review

Understanding the dynamics that govern the evolution of correlation as the measure of dependence between financial instruments is of paramount importance in a number of financial applications including asset pricing, asset allocation and risk management. Accurate estimation and forecasting of correlation process can provide significant improvement in derivatives pricing and efficient asset allocation decisions within a trading and hedging framework. The first three Sections of this Chapter present the importance of correlation in the portfolio theory, studies that provided evidence of time-variability in the correlation process and alternative estimators that have attracted the vast majority of existing studies on correlation modelling. In Section 4, a detailed literature of implied and realized correlation estimates in presented. The notion and the emergence of correlation risk as a priced risk factor is discussed in Section 5. The next Section reviews the efficient market hypothesis with a special view to the options and futures markets. The last Section discusses the alternative measures of hedge ratio and its properties.

2.1 Portfolio Theory and Diversification

The fundamental principle underlying the mean – variance portfolio model (Markowitz, 1952) is the optimization of the risk-return tradeoff associated with investment portfolio decisions. In essence, an investor will seek to maximize the expected portfolio return while minimizing the variance of the asset returns. The basic inputs for the model are the expected return, the standard deviation of each asset as well as the correlation matrix between the assets. The expected return $E(R_{P,t})$ and variance, $\sigma^2_{P,t}$, of the portfolio are defined as follows:

$$E(R_{P,t}) = \sum_{i=1}^{N} w_{i,t} E(R_{i,t}) \tag{2.1}$$

$$\sigma^2_{P,t} = \sum_{i=1}^{N} w^2_{i,t}\sigma^2_{i,t} + 2\sum_{i=1}^{N}\sum_{j>i} w_i w_j \sigma_i \sigma_j \rho_{ij} \tag{2.2}$$

where, w_i is the relative weight of asset i to the portfolio, $E(R_i)$ and $\sigma_{i,t}$ are the return and the variance of the asset i at time t, respectively, and ρ_{ij} is the correlation of asset i and j.

Correlation is a direct measure of the diversification benefit that derives from investing in two or more assets. A value of correlation coefficient equal to one signifies the perfect comovement of asset returns and elimination of diversification benefit. In this case, the portfolio variance is the weighted sum of individual assets' variance, i.e.

$$\sigma^2_{P,t} = \sum_{i=1}^{N} w^2_{i,t}\sigma^2_{i,t} + 2\sum_{i=1}^{N}\sum_{j>i} w_i w_j \sigma_i \sigma_j \tag{2.3}$$

In terms of portfolio optimization, an investor will always seek assets that are negatively correlated so as the loss from one asset is offset by the gain of the other asset. The maximum diversification benefit is obtained for a perfectly negative correlation value of -1. In that case, the portfolio variance eliminates to

$$\sigma^2_{P,t} = \sum_{i=1}^{N} w^2_{i,t}\sigma^2_{i,t} \tag{2.4}$$

The Capital Asset Pricing Model (CAPM) laid the foundation to asset pricing theory. In essence, the model suggests that the equilibrium price of any asset is a function of the time-value of money, represented by the risk – free rate, and the reward that the investor requires to uptake additional risk. The equilibrium relationship is as follows:

$$E(R_i) = R_f + \beta_i (E(R_m) - R_f) \tag{2.5}$$

where R_f is the risk free rate, R_m is the return of the market and β_i is the beta coefficient that represents the sensitivity of the asset to the market or systemic risk, defined as:

$$\beta_i = \frac{\sigma_{i,M}}{\sigma_M^2} = \rho_{i,M} \frac{\sigma_i}{\sigma_M} \tag{2.6}$$

where $\rho_{i,M}$ is the correlation coefficient between asset *i* and the market portfolio returns, and σ_i and σ_M are the standard deviations of asset and market returns, respectively. Evidently, accurate correlation estimation is of utmost importance to the derivation of equilibrium market prices.

2.2 Time variability and Asymmetry of Correlation

One of the critical assumptions of the CAPM model is that the variance and covariance of the assets is time-invariant (Jobson and Korkie, 1981). Earlier studies assessed the time-varying assumption of correlation structure over different sample periods and find evidence of stable correlation structure. Kaplanis (1988) examined the stability of correlation between returns of ten stock markets for the period of 1976-1982. Applying the test procedure suggested by Jenrich (1970) over adjacent sub-periods of 46 months, the null hypothesis of constant correlations fails to be rejected. Ratner (1992) utilizes a sample of international stock market indices during 1973-1989 and also rejects the null of unstable correlation coefficients. Finally, Tang (1995) uses data of 17 stocks traded in the Hong Kong stock market during 1981-1992 and finds that the covariance and correlation structures are stable over time with the latter being significantly more stable.

In contrast, a plethora of studies provides evidence supporting the time-variation in the correlation structure. Early studies from Makridakis and Wheelwright (1974) and Bennett and Kelleher (1988) documented that co-movement of international stock markets are unstable over time. Koch and Koch (1991) provided evidence of growing market interdependence within the same region while the result is more pronounced in

1980 and 1987 as compared to 1972. Another approach adopted in the study of time-varying correlation is based on the autoregressive conditional heteroskedasticity (ARCH) model developed by Engle (1982) and its variations. Longin and Solnik (1995) employ a bivariate GARCH model to the asset returns and use the Jenrich test for the null assumption of constant conditional correlation. Results for monthly excess returns of seven major countries over the period of 1960 – 1990 suggest that unconditional correlation matrix is unstable over time. Darbar and Deb (1997) employ a GARCH model and find that correlation of international equity markets changes over time.

Additionally, there is a large body of previous research that has provided evidence of increased correlation during turbulent periods, associated with high volatility and low returns (e.g. Karolyi and Stulz 1996; Ang and Bekaert 2002; Longin and Solnik 2001; Ang and Chen 2002; Bae, Karolyi and Stulz 2003). As the correlation between assets that span the risk-return spectrum increases, the aggregate risk borne by investors will increase and the diversification benefit will eliminate. In particular, Longin and Solnik (2001) study the conditional correlation of international equity returns. To this end, they employ the notion of extreme value theory to model the distribution tails and find that correlation is mainly affected by the market trend, while high volatility per se does not necessarily lead to an increased correlation. More importantly, based on an empirical distinction of bull and bear markets, they find that correlation increases sharply only during bear markets. Campbell, Koedijk and Kofman (2002) use VaR-based conditional correlation estimators for international stock index returns during 1990-1999 and reach the sample conclusion of increased correlation in financial turmoil periods, suggesting that downward movements induce financial contagion across markets. Garcia and Tsafack (2011) study the co-movement in international equity and bond markets and find that strong dependence between the same asset classes at an international level, especially after accounting for asymmetry. Finally, Amira, Taamouti and Tsafack (2011) find evidence of asymmetric effect of volatility on correlation, which however fades out when market returns are included as an explanatory factor, thus proposing that the correlation increase is mainly driven by past returns and market direction.

Turning to the economic significance of asymmetry considerations, Ang and Chen (2002) employ several statistical tests and find that a two-regime switching model is able to capture and explain the correlation asymmetry. Additionally, they provide evidence of increased correlation between stock portfolios and US market returns during periods of negative market return while the results are more pronounced in the case of small or value firms. Hong et al. (2007) suggest a model-free test of asymmetric correlation and, similarly to Ang and Chen (2002), find that asymmetry holds across alternative Fama and French portfolios. In addition, they augment the optimal portfolio framework to account for asymmetric dynamics and find that the net utility benefit for a mean-variance investor reaches 2% on an annual base. The economic gains arising from the inclusion of asymmetries in a portfolio allocation decisions is further supported by Das and Uppal (2003) and Patton (2004).

Furthermore, the recent financial turmoil shed light on the notion of financial contagion, or else the transmission of financial distress across assets and countries interpreted as increased correlation. Although, given the recency of the crisis, little empirical evidence is available, obtained results consent on the increased importance of financial contagion. Longstaff (2010) finds that financial contagion was mainly disseminated through correlated liquidity and risk premium channels rather than information channel. Dungey et al. (2010) consider a dataset of stock and bond markets of six countries and study the effect of financial contagion during five major worldwide crises; namely, the Russian and Brazilian crisis in 1998 and 1999, respectively, the internet bubble burst in 2000, the Argentinian crisis in 2002-2005 and finally, the subprime mortgage crisis of 2007-2009. They find that financial crises are identical while the contagion effect is more pronounced in the case of Russian and subprime crises. In addition, Syllignakis and Kouretas (2011) apply the Dynamic Conditional Correlation (DCC) multivariate GARCH model to study the contagion effects of US, German and Russian markets to seven emerging markets in Europe and find that correlation increased drastically. The increase can be explained by several macroeconomic determinants, domestic and foreign monetary variables, as well as exchange rate movements. Cheung, Fung and Tsai (2010) suggest that the TED spread, defined as the difference of three-months futures contracts on US T-bills and Eurodollars has explanatory power for the contagion effect across major financial markets. Finally, Tacchella et al. (2012) examine the time evolution of correlation coefficients for 494 stocks listed in the S&P 500 index and observe an increase in correlations, which however, varies across economic sectors. In a similar context, Bekaert et al. (2012) suggest that contagion between the US equity market and equity market in 55 other countries are significant, yet small. In contrast, they provide evidence of increased correlation between domestic equity market and domestic individual equity portfolios.

2.3 Correlation Modelling

Expected return is directly observable whereas the second moments of the distribution, i.e. volatility and correlation, need to be estimated within the context of a particular model. Correlation modelling has stimulated the research interest over the past years, although admittedly, the relevant studies in the volatility context are more plentiful. While earlier estimators treated correlation as a constant variable, later studies on correlation modelling were intrigued by the empirical findings on time-variation of correlation structure. To this end, several models, from naïve to moving average and GARCH-type models have been proposed in the literature.

2.3.1 Simple Historical and Moving Average models

The simplest correlation estimator is the historical correlation defined as the ratio of covariance and sum product of assets' standard deviations. The implicit assumption is that correlation matrix remains unchanged over time and past distributional properties will be similar in the future. In essence, historical correlation is defined as:

$$\rho_{ij} = \frac{\sum_{i=1}^{N-1}(R_i - \bar{R}_i)(R_j - \bar{R}_j)}{\sqrt{\sum_{i=1}^{N-1}(R_i - \bar{R}_i)^2}\sqrt{\sum_{j=1}^{N-1}(R_j - \bar{R}_j)^2}} = \frac{Cov(R_i, R_j)}{\sigma_i \sigma_j} \qquad (2.7)$$

where R_i, R_j are the asset returns with mean value of $\bar{R}_i$ and $\bar{R}_j$ and standard deviation (or volatility) equal to σ_i and σ_j, respectively.

The simple historical model assigns equal weight to past-observed values. A widely used alternative method was proposed by JP Morgan RiskMetrics™. The model introduces a smoothing parameter, called lambda, which assigns more weight to recent observations. The Exponentially Weighted Moving Average (EWMA) model is defined as:

$$\rho_{ij,t} = \frac{\sum_{t=T-n-1}^{T-1} \lambda^{T-t-1}(R_{it} - \bar{R}_i)(R_{jt} - \bar{R}_j)}{\sqrt{\sum_{t=T-n-1}^{T-1} \lambda^{T-t-1}(R_i - \bar{R}_i)^2}\sqrt{\sum_{t=T-n-1}^{T-1} \lambda^{T-t-1}(R_j - \bar{R}_j)^2}} \tag{2.8}$$

where λ determines the weighting of returns and takes a value between 0 and 1, with a value closer to 1 denoting that the change in weights between time periods is small. The Risk Metrics database proposes λ to be equal to 0.94.

2.3.2 GARCH-type models

Following the recognition of time-varying traits of the volatility, the seminal work of Engle (1982) introduced the Autoregressive Conditional Heteroskedastic (ARCH) model, that was later generalized by Bollerslev (1986) to the GARCH model to model the evolution of volatility process. The main advantage of the specification is that the conditional variance is an autoregressive function of past information contained in the information set I_t. A GARCH(p,q) model is defined as follows:

$$\begin{aligned} & r_t \mid I_t \sim N(0, h_t^2) \\ & h_t^2 = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \ldots + \alpha_p\varepsilon_{t-p}^2 + \beta_0 + \beta_1 h_{t-1}^2 + \ldots + \beta_q h_{t-q}^2 \end{aligned} \tag{2.9}$$

where *p* is the number of lags of the error terms, ε, and *q* is the number of lags of conditional variance.

Volatility modelling has attracted the vast majority of previous studies. Apart from the simple GARCH models, several extensions have also been developed so as to account for asymmetry (Exponential GARCH - EGARCH, Glosten-Jagannathan-Runkle GARCH - GJR-GARCH, and Asymmetric Power ARCH - APARCH) and degree of integration (Integrated GARCH – IGARCH, Fractionally Integrated GARCH – FIGARCH among others).

The extension of above mentioned univariate models in the multivariate context allows the investigation of comovement of financial returns. Multivariate GARCH (MGARCH) models have been employed in the relevant literature of volatility spillover and correlation transmission, i.e. contagion. The main issue to be addressed during the estimation of MGARCH models is the high dimensionality, while maintaining the positive definiteness of the variance-covariance matrix.

The VEC-GARCH model, developed by Bollerslev et al. (1988), was the first multivariate extension of the simple GARCH model. The model parameterizes every variance and covariance as a function of lagged returns and lagged cross-product of returns. The model may be written as follows:

$$vech(H_t) = c + \sum_{j=1}^{q} A_j vech(\varepsilon_{t-j}\varepsilon_{t-j}^{'}) + \sum_{j=1}^{p} B_j vech(H_{t-j}) \tag{2.10}$$

where H_t is variance-covariance matrix of returns, c is a N(N+1)/2 + 1 vector of parameters, $vech(\cdot)$ is a column stacking operator of the lower triangular part of the argument square matrix and A_j and B_j are $N(N+1)/2 \times N(N+1)/2$ parameter matrices. Despite the flexibility in use, the VEC model suffers from the dimensionality problem as the total number of parameters to be estimated equal to $(p+q)\times N(N+1)/2)^2 + N(N+1)/2$. In addition, the conditions imposed to ensure the positive definiteness of the VCV matrix are rather restrictive.

The BEKK model, introduced in Engle and Kroner (1995), is a restricted version of the VEC model, in which the conditional covariance matrices are positively defined by construction. The model is given by:

$$H_t = CC' + \sum_{j=1}^{q}\sum_{k=1}^{K} A_{kj} A_{kj}' \varepsilon_{t-j}\varepsilon_{t-j}' + \sum_{j=1}^{p}\sum_{k=1}^{K} B_{kj} B_{kj}' H_{t-j} \tag{2.11}$$

where A_{kj}, B_{kj} and C are $N\times N$ parameter matrices and C is lower triangular.

The conditional correlation at time t is derived as the ratio of estimated covariance, $h_{ij,t}$, and product of estimated standard deviations, h_i and h_j, where each distribution moment is estimated independently from any the above-mentioned GARCH-type models:

$$\rho_{ij,t} = \frac{h_{ij,t}}{h_{i,t} h_{j,t}} \tag{2.12}$$

In a first attempt to model directly the correlation matrix, Bollerslev (1990) introduced the Constant Correlation model (CCC-GARCH) which is essentially a nonlinear combination of univariate GARCH processes. The correlation is considered constant over time while covariances are time varying. They basic concept is that the conditional variance-covariance matrix (H_t) is decomposed into conditional volatilities (D_t) and correlations (R) matrices as follows:

$$H_t = D_t R D_t \tag{2.13}$$

A direct interpretation of the above equation suggests that each correlation estimator of the conditional correlation matrix equals to $\rho_{ij}=h_{ij,t}/(h_{ii,t}\times h_{jj,t})$, where $h_{ij,t}$ can be estimated from univariate GARCH processes and ρ_{ij} is time-invariant.

Engle (2002) extended the CCC-GARCH model to the Dynamic Conditional Correlation (DCC) model, by allowing the correlation matrix of returns to be time varying, i.e.

$$H_t = D_t R_t D_t \tag{2.14}$$

The estimation procedure is as follows. Firstly, a univariate GARCH process is fitted to model the volatilities, $h_{i,t}$, of the k assets. The standardized residuals, z_t, later used to model the correlation of returns, are obtained by dividing the error terms, ε_t, with h_t. The DCC-GARCH model is then given by:

$$\begin{aligned} Q_t &= (1-a-b)\bar{Q} + a z_{t-1} z'_{t-1} + b Q_{t-1} \\ R_t &= D_{Q_t}^{-1/2} Q_t D_{Q_t}^{-1/2} \end{aligned} \tag{2.15}$$

where $\bar{Q}$ is the unconditional correlation matrix and a, b are scalars.

2.4 Alternative measures of correlation

In the previous Section, the dominant and most widespread methodologies of modelling volatility and correlation have been presented. Elton and Gruber (1973) were the first to emphasize and address the complexities involved in the calculation of correlation. The main difficulty arises from the dimensionality problem. Essentially, the number of pairwise correlation coefficients, for a portfolio of N assets, escalates up to $N(N-1)/2$. As an example, for a portfolio consisting of 100 assets, 4.950 would have to be produced.

In an attempt to depart from the complexities arising from the modelling of the correlation process, over the last decades, alternative measures of the second moment of distribution have been presented in the literature and attracted the interest of both academics and practitioners. The first one, implied volatility and implied correlation, is derived from option prices while the second, the realized counterpart, takes advantage of the availability of high-frequency data and provides a straightforward measurement of volatility and correlation.

2.4.1 Implied Correlation

The concept of extracting moments of asset return distribution from option prices has been central to an enormous number of studies over the last years. By definition, moments of distribution implied from observed option prices reflect the market expectation of future asset price changes and the subsequent indisputable valuation of said uncertainty that is driven by the market forces without relying on any specific econometric modelling specification. The estimation of distribution moments of stock price from option prices, as forward-looking measures, is widely considered to outperform historical correlation measures in terms of informational efficiency. Thereafter, conditioned on the market efficiency and the precision of the option-pricing model, measures deduced from option prices subsuming the information set available at any time t, are expected to be efficient and unbiased estimators of the true realized processes. Bates (1991) suggests that option prices reflect the market participants' expectations by giving a direct indication of the aggregate subjective distributions of investors.

Despite the numerous studies on implied volatility, implied correlation has received notably less attention in the literature. Earlier studies from Campa and Chang (1998) and Lopez and Walter (2000) show that correlation implied from options on foreign exchange rates demonstrate superior forecasting ability of future realized correlation. For comparison purpose, they also derive the historical correlation, EWMA as well as correlations based on GARCH models and find that implied correlation forecasts dominates the more traditional approaches. Castrén and Mazzotta (2005) conduct a comparative study of the forecasting ability of FX options implied correlation and return-based correlation measures and find mixed results across currencies. However, when they combine the information set of alternative estimators, they obtain the highest R-squared.

Skintzi and Refenes (2005) were the first to derive option-implied measure of correlation from equity data. Using data from the Dow Jones Industrial Average Index, the Implied Correlation Index (ICX) is defined as the average correlation coefficient that captures the difference between implied volatility of an index option and the weighted average of implied volatility of the constituent stocks. They assess the statistical properties and the forecasting performance of the Index and find that ICX provides accurate forecasts of realized correlation. Driessen, Maenhout and Vilkov (2009) obtain the implied correlation from option prices on the S&P 100 index and constituent stocks and provide evidence of high predictive power of future realized correlation. In a more recent paper, Driessen, Maenhout and Vilkov (2012) find that the predictive power is higher for 6-month and 1-year forecast horizon.

The usefulness of option-implied measures to portfolio allocation applications has been the subject of limited recent studies. Kostakis, Panigirtzoglou, and Skiadopoulos (2011) use data on the S&P 500 index and construct a portfolio consisting of a risky and a risk-free asset. Results from the evaluation of out-of-sample portfolio performance under several statistical and economic functions suggest that use of option implied information provide significant improvement of investors' risk-adjusted returns. DeMiguel et al. (2012) apply option-implied information in a mean-variance portfolio selection context and examine the improvement in terms of Sharpe ratio, portfolio volatility, certainty-equivalent return, and turnover. They find that implied volatility reduces the volatility but does not contribute to the improvement of the Sharpe ratio. In contrast, the inclusion of option implied skewness information increases both the Sharpe ratio and the expected returns.

The usefulness of implied measures has also been examined within a dynamic hedging context. Christoffersen, Jacobs, and Vainberg (2008) expanded existing literature on implied volatility and correlation and proposed a methodology to derive a forward-looking measure of beta from option prices. They find evidence of forecasting superiority of implied beta when compared to historical or return-based estimates. Buss and Vilkov (2012) propose a parametric solution to the calculation of implied correlation, compute option-implied betas and test their accuracy as realized beta predictors. Results suggest that, in line with Christoffersen, Jacobs, and Vainberg (2008), implied betas are accurate predictors of realized beta.

Interestingly, the Chicago Board Options Exchange (CBOE), having recognized the increased interest on correlation as a driver of diversification benefit, launched the CBOE S&P 500 Implied Correlation Index in July 2009. Essentially, the index is a measure of the average correlation between the S&P 500 index options and the basket of options, consisting of the largest 50 stocks. However, the construction of the index has several drawbacks. First, the options used for the derivation of the index are not filtered to eliminate microstructure and liquidity issues. Second, the implied volatility of the options is model-based; the Black-Scholes and the Barone-Adesi Whaley option valuation models are used for the index and stock options, respectively. To address these drawbacks, I develop a model-free alternative of the implied correlation index in Chapter 3 of this thesis.

2.4.2 Realized Correlation

The forecasting performance of correlation estimators, analyzed extensively in Section 2.3, rely heavily on the implicit assumption on asset returns distribution. Although several steps are taken forward the reduction of the dimensionality curse, the practical application of multivariate models remains feasible for very low dimensions. The increased availability of high-frequency data has set the foundations for a new area of research that involves estimating, modeling and forecasting conditional volatility and correlation from intraday data.

Realized Volatility (RV) is obtained by summing the squared intraday returns whereas Realized Covariance (RCov) is the cross product of intraday returns. It follows that Realized Correlation (RC) is obtained by dividing realized covariance by realized volatility. Based on the theory of quadratic variation, Andersen, Bollerslev, Diebold, & Labys, henceforth ABDL, (2001, 2003) and Andersen, Bollerslev, Diebold, & Ebens, henceforth ABDE (2001), showed that, as the sampling frequency tends to infinity, the realized measures are consistent estimators of the true latent processes. Realized measures are essentially "model-free" thus exhibiting notably advantageous properties over previously reported estimators that rely on parametric models that induce econometric or mathematical misspecification.

Similar to the measure of implied correlation, realized covariance and correlation have been studied in a remarkably less extent. To fix ideas, consider that the two assets $\{X_t\}$ and $\{Y_t\}$ follow a standard Brownian motion B^X and B^Y with a drift μ_t^X and μ_t^Y and instantaneous variance $\sigma_t^{2,X}$ and $\sigma_t^{2,Y}$, respectively. The integrated covariance is then given by:

$$\left\langle X,Y\right\rangle_T = \int_0^T \sigma_t^X \sigma_t^Y d\left\langle B^X, B^Y\right\rangle_t \tag{2.16}$$

Under the limit theorem for stochastic processes, the realized covariance, defined below as the sampling frequency tends to infinity, is a consistent estimator of the true latent process (see Barndorff-Nielsen and Shepherd, 2002; Mykland and Zhang, 2006).

$$RCov_T = \sum_{i:\tau_i \in [0,T]} (X_{T_i} - X_{T_{i-1}})(Y_{T_i} - Y_{T_{i-1}}) \tag{2.17}$$

It follows that the realized correlation is obtained by dividing the realized covariance by the product of variances of assets *X* and *Y* as follows:

$$RC_T = \frac{RCov_T}{\sqrt{\sum_{i:\tau_i \in [0,T]} (X_{T_i} - X_{T_{i-1}})^2 \sum_{i:\tau_i \in [0,T]} (Y_{T_i} - Y_{T_{i-1}})^2}} \tag{2.18}$$

In practice, the estimation of realized covariance encounters several problems. The first, also present in the calculation of realized volatility, is the market microstructure noise that arises from price discreteness, bid-ask spreads or the lack of liquidity. Aït-

Sahalia et al. (2005) and Zhang et al. (2005) showed that the bias increases as the sampling frequency increases. Thereafter, the optimal sampling frequency has been at the core of several studies that discuss the trade-off between market microstructure effects and loss of information from sampling too sparsely (see Zhou, 1996; Bandi and Russell, 2008; Oomen, 2006 and De Pooter, Martens, and van Dijk, 2008). The second problem stems from the asynchronous trading of the assets included in the covariance matrix. The so called "Epps effect" (Epps, 1979) states that realized covariance tends to zero as sampling frequency increases. Barndorff-Nielsen et al. (2011) proposed the use of realized kernels and utilized the refresh time to synchronize the timing of observations while also incorporated leads and lagged autocovariance terms. Alternative methods of addressing the two problems include subsampling (Zhang et al., 2005) and pre-averaging (Jacod et al., 2009). Zhang et al. (2005) proposed the two-time scales estimator that essentially utilizes realized variance estimates from a low and a higher frequency sampled returns and applies the subsampling method to reduce the variance of the low frequency realized variance.

Several studies have assessed the information content of realized measures within a portfolio optimization and hedging allocation context. Fleming, Kirby and Ostdiek (2003) examine the use of intraday data in the context of tactical asset-allocation decisions and suggest that the improvement in terms of economic gains can be substantial. Yeh, Huang and Hsu (2008) suggest that trading strategies based on intraday returns result to an almost perfect hedge, with the risk reduction being equal to nearly 100 percent. In contrast, Liu (2009) show an investor who aims to track the S&P 500 index with 30 stocks of the Dow Jones Industrial Average Index shall switch from daily to intraday data only when the investment horizon is shorter than six months or portfolio rebalance occurs on a daily basis.

Following the proliferation of alternative measures of volatility and correlation, namely realized, implied and conditional measures, a number of studies compare the informational efficiency and forecasting accuracy. Martens (2002) uses data on the S&P 500 futures markets and find that when overnight returns are excluded, realized volatility is the best proxy of daily volatility. Martens and Zein (2004) compare the three alternatives measures of volatility on the S&P 500 index, the YEN/USD FX rate and the Crude Oil and provide evidence of better forecasting performance of the realized alternative. Koopman, Jungbacker, and Hol (2005) attain similar results for the S&P 100 index. Pong et al. (2004) conduct the same comparative analysis for foreign exchange rates and find that realized volatility, closely followed by the implied measure, yield better forecasting performance for one-month and three-month horizons.

Andersen et al. (2006), extended their previous work on realized volatility and correlation, and defined realized beta as the ratio of realized covariance and variance. They find that, in contrary to the highly persistent and fractionally integrated variance and covariances series, realized betas are significantly less persistent and best modeled as stationary processes. Morana (2007) also uses daily returns to estimate factor betas. The findings of the study support the time-variability and the predictability of the series, especially in the short run. Patton and Verrado (2012) examine the information flow and beta reaction to earnings announcements. They show that realized betas increase significantly on announcement days while the days that follow they revert to their long-run average levels.

2.5 Correlation risk

Correlation risk arises from fluctuations in the correlation structure between assets or across countries. As discussed in Section 2.2, correlation is expected to increase during periods of high volatility and low returns, resulting in moderated diversification benefits when they are most needed. It is therefore natural for investors to seek protection and to pay an extra premium for obtaining assets that perform well and offer higher payouts in states of high correlation.

Krishnan, Petkova and Ritchken (2009) analyze the inter-asset correlation structure of 25 Fama and French portfolios for an extended period and examine whether pairwise correlation is priced in the cross section of returns. They find significant evidence of negatively priced correlation risk even after controlling for macroeconomic factors, higher moments and different test assets. Results suggest that investors are willing to pay an extra premium for stocks that are performing well in states of diminished diversification; thus increasing the demand and diminishing the expected returns. Mueller, Stathopoulos and Vedolin (2012), using data on FX rates and options, also find that correlation risk is a significant factor that carries a negative premium, while currency portfolios with low or negative exposure to correlation risk perform well. In an asset allocation context, Buraschi, Porchia and Trojani (2010) suggest that the hedging demand for correlation risk substantially affects the optimal portfolio weights.

Prompted by the seminal paper of Driessen, Maenhout and Vilkov (2009), several studies have investigated the pricing of correlation risk through the decomposition of the variance risk premium of the index into the variance risk premium of the constituent stocks and the correlation structure. In specific, Driessen, Maenhout and Vilkov (2009) provide evidence of priced volatility risk premium on the index options in contrast to individual options where the price is not significant. They argue that the differential pattern is mainly attributed to the negative price of correlation risk that is only present in the index options. Intuitively, index options are expensive and earn lower returns because they offer a valuable hedge against correlation increases and insure against diminished diversification benefits. Schürhoff and Ziegler (2010) and Chen and Petkova (2012) extend the analysis and decompose the volatility risk premium in systematic and idiosyncratic components. Schürhoff and Ziegler (2010) find that systematic (idiosyncratic) risk is negatively (positively) priced whereas Chen and Petkova (2012) find that only systematic risk carries a significant price. In a similar context, Cosemans (2011) further decomposes the volatility and correlation risk premia into the short and long run components. In line with previous literature, he finds that, on market level, both components obtain a negative price whereas in case of individual options, the short (long) term component is positively (negatively) priced.

The Volatility Risk Premium (VRP), defined as the difference of implied and realized volatility, is often regarded a measure of market-implied risk aversion while several studies have provided evidence of negatively priced premium. Jackwerth and Rubinstein (1996) provided evidence that at-the-money Black–Scholes implied volatilities are systematically and consistently higher than realized volatilities measured as the square root of the empirical quadratic variations of the price sampled intradaily at small time intervals. Duarte and Jones (2007) observed that the volatility risk premium varies positively with the implied volatility of the S&P 500 index options. Bakshi and Kapadia (2003a) provided evidence of negative volatility risk premium in index option markets by creating a delta hedged portfolio and observed that volatility affects delta hedges gains even after controlling for jump fears. In their work that followed, Bakshi and Kapadia (2003b) also show that the volatility risk premium embedded in individual options is negatively priced, though much smaller to index VRP.

The information content of volatility and correlation risk premia for future market returns has been at the core of limited studies. Bollerslev, Tauchen, and Zhou (2009) measure the volatility risk premium as the difference of implied volatility (following the model-free methodology suggested by Britten-Jones and Neuberger, 2000) and realized volatility, from high-frequency intraday data, and find that the VRP is able to predict future equity premium. Cosemans (2011) finds that the predictive power of market VRP might be attributed to the correlation risk premium and to the market component of the VRP in individual stock options. Driessen, Maenhout and Vilkov (2012) assess the predictive power of option-implied correlation for future aggregate stock returns and find that, when combined with the VRP, they jointly explain 15% of future market variance.

Evidence of priced correlation risk induced academics and practitioners to develop trading strategies that exploit directly the correlation risk. Intrigued by the finding that the premium of the index is relatively higher than the premia of constituent stocks, Nelken (2006) introduced the “dispersion strategy”, usually implemented through options and straddles or variance swaps. The strategy involves short/(long) positions in index and long/(short) positions in the constituent stocks,. The strategy will result in profits when either the difference of implied volatility of the index and individual stocks eliminates or the options expire and the earnings from written stocks are greater than the loss from the long positions. Jacquier and Slaoui (2010) suggest that the profit/(loss) of the dispersion trading strategy is equal to the spread between implied and realized correlation.

Two alternative lines of reasoning have been set forward to explain the profitability of the dispersion trading strategy. The first one is based on the extensively discussed price of correlation risk that induces index options to be more expensive than individual options. The second hypothesis is related to market demand and supply forces that lead option prices to deviate from their theoretical values. Deng (2008) uses data on the S&P 500 options and studies the period of late 1999 and 2000 when the options market has undergone structural changes (e.g. cross-listing of options and reduced bid-ask spreads amongst others). He finds that the profitability of the dispersion trading strategy eliminated in 2000, suggesting market inefficiency as the institutional changes should not have affected the pricing of correlation risk. In contrast, Härdle and Silyakova (2012) employ the dynamic semiparametric factor model (DSFM) to model the implied correlation of options on the DAX index and find that a dispersion trading strategy with variance swaps outperforms alternative strategies.

2.6 Market efficiency and predictability of moments

An enormous literature has been devoted in examining the ability of alternative measures of volatility and correlation to provide accurate forecasts of future realized moments. In Sections 2.3 and 2.4, various studies that address that issue are presented. Another stream of research that has received notably less attention is the predictability of the moments per se, i.e. whether the dynamics that govern the evolution of the series contain any predictable pattern. The above question constitutes a direct test of the efficient market hypothesis.

Efficient Market Hypothesis (EMH) evolved as a fundamental concept of modern finance. Fama (1970) was the first to provide a rather vague definition of market efficiency as follows: “A market in which prices at any time “fully reflect” available information is called efficient”. Some years later, Jensen (1978) followed and different approach and defined a market as efficient “[…] with respect to information set θ_t if it is impossible to make economic profits by trading on the basis of information set θ_t”.

A significant amount of studies has assessed the efficiency of option markets through the predictability of the volatility dynamics while the results remain mixed. Chiras and Manaster (1978) conclude that the Black-Scholes implied variance, compared to historical variance, is a better predictor of future realized variance. Consecutively, they employ a trading strategy that yields significant abnormal returns suggesting that the CBOE market was inefficient. Galai (1978) assesses the hypothesis of options market efficiency through the investigation of boundary conditions of the options and underlying spot market and finds that positive profits can be attained. Goyal and Saretto (2009) find evidence of economically significant predictability in the cross-section of stock options returns.

In contrast, several studies find evidence of a statistically significant predictability, which however cannot be economically exploited. Harvey and Whaley (1992) use near-the-money options on the S&P 100 market and find that implied volatility fails to produce significantly positive profits when transaction costs are taken into account. Noh, Engle and Kane (1994) compare implied volatility forecasts to those obtained from a GARCH model and find that the latter results in greater profits. Additionally, Gonçalves and Guidolin (2006), Bernales and Guidolin (2010) and Chalamandaris and Tsekrekos (2010) study the predictability in the implied volatility surface and find conclude that the predictability supported by the statistical measures cannot to attain abnormal profits. Konstantinidi, Skiadopoulos and Tsagkaraki (2008) reach the same conclusion for several European and US indices volatility indices. Finally, Neumann & Skiadopoulos (2013) discuss the predictability of higher-order moments and find that economic gains, after transaction costs, are eliminated thus suggesting that the efficient market hypothesis for the S&P 500 market holds.

To the best of my knowledge, despite the extensive literature on the predictability of volatility, the predictability of correlation has remained unexplored. Buraschi, Trojani and Vedolin (2013) provide a theoretical explanation of why correlation process might contain predictable patterns. In specific, they develop an equilibrium model of two heterogeneous agents with different beliefs and study the effect of aggregate economic uncertainty and diversity in beliefs across investors on volatility and correlation risk premia. Chapter 4 addresses the issue of predictability of options-implied correlation.

Market efficiency in futures markets has been initially assessed through regression analysis, which however produces unreliable results if the series are non-stationary, a common trait of equity returns. To overcome the ambiguities deriving from such an assumption, an alternative approach based on the cointegrating relationship of the spot and the futures markets has emerged (see for example MacDonald and Taylor, 1988; Baillie, 1989; Chowdhury, 1991; Lai and Lai, 1991; Brenner and Kroner, 1995). Based on the fundamental principles underlying the pricing of futures, it is expected that the spot and the futures market co-move closely. Divergence from the long-run cointegrating relationship calls the market efficiency hypothesis into question.

The predictability of futures prices may derive from several types of risk premia. Bessembinder and Chan (1992) find that the several macroeconomic factors have limited predictive power over movements in agricultural, metal, and currency futures prices. Kho (1996) and Hong and Yogo (2012) suggest that the predictability of foreign exchange rate and commodity futures prices, respectively, is mainly attributed to time-varying risk premia. Miffre (2001) finds that time-varying risk premia, driven by investors' expected returns along with variation in betas eliminate the predictability of the FTSE 100 Index futures. Recently, Szymanowska et al. (2014) identified spot returns premium and term structure premia related to changes of the basis can account for the commodities futures risk premia.

From an economic perspective, several papers have provided evidence of exploitable patterns that can yield profit to market participants (indicatively, Yoo and Maddala, 1991 for commodity and currency futures; Klitgaard and Weir, 2004; and Wang, 2004 for currency markets). In contrast, Konstantinidi and Skiadopoulos (2011) employ a trading strategy to explore the predictability on futures of the VIX index and find that no abnormal returns can be attained.

2.7 Optimal Hedge Ratio

A trader is willing to enter a transaction with derivatives products in order to minimize the risk deriving from changes in the spot market or to enhance plausible profits. Hedging with futures is among the most widespread risk management practices. The hedge ratio denotes the number of futures contracts that the investor holds in order to offset changes in the underlying spot market. The prevailing methodology of obtaining the hedge ratio originates in minimizing the portfolio risk (Johnson, 1960; Ederington, 1979). The basic underlying assumption is that all investors embrace the minimum variance scope and take positions accordingly. In contrast, another stream of literature has recognized the discrepancy of investors' risk aversion. A number of alternative hedge ratios that seek to maximize the risk-return function have been developed. Cheung, Kwan and Yip (1990), Lien and Luo (1993), Kolb and Okunev (1993) and Shalit (1995) and introduced the Mean-Extended-Gini-coefficient (M-MEG) that considers the investor's level of risk aversion. In the same context, the Mean Generalized Semivariance (M-GSV) focuses on the downside risk, i.e. returns falling below a target level (De Jong, De Roon, and Veld, 1997; Lien and Tse, 1998, 2000; Chen et al. 2001;).

However, due to the simplicity of calculation, the minimum variance hedge ratio has been at the core of the hedging problem and has attracted the vast majority of research interest. Several methods of estimating the minimum variance hedge ratio have been proposed. The first method, the naïve hedge ratio, has considered the hedge ratio to be equal to unity during the entire investment horizon, assuming that spot and future returns are perfectly correlated. Ederington (1979) suggested that the optimal hedge ratio could be obtained from regressing the spot and the futures returns via Ordinary Least Squares (OLS). The static OLS method recognizes the imperfect correlation but accepts that the joint distribution remains constant. The OLS static hedge ratio is defined as follows:

$$\Delta S_t = \alpha + \beta \Delta F_t + \varepsilon_t \tag{2.19}$$

where the slope of the regression, β, is the minimum variance hedge ratio.

Nevertheless, the above-mentioned methods fail to consider the widely documented time variation of the asset returns, and, consequently of the hedge ratio. In an attempt to depart from the constant hedge ratios and following the emergence of the GARCH family models, several studies supported the time variability of the hedge ratio and introduced the GARCH models in the estimation procedure of the spot and futures (co)variance matrix. For example, the estimation procedure of obtaining a hedge ratio from a GARCH (1,1) model is outlined below. First, the spot and futures variance, h_{ss} and h_{ff}, are modelled from a GARCH(1,1) model as follows:

$$\begin{aligned} h_{ss,t} &= c_{ss} + \alpha_1 \varepsilon_{s,t-1}^2 + \beta_1 h_{s,t-1} \\ h_{ff,t} &= c_{ff} + \alpha_1 \varepsilon_{f,t-1}^2 + \beta_1 h_{f,t-1} \\ h_{sf,t} &= \rho \sqrt{h_{ss,t} h_{ff,t}} \end{aligned} \tag{2.20}$$

where h_{sf} is the covariance and ρ is the correlation coefficient assumed constant. The optimal GARCH(1,1) hedge ratio is then calculated as

$$HR_t = \frac{h_{sf,t}}{h_{ff,t}} \tag{2.21}$$

Despite the well-documented time variation of the covariance matrix, several studies suggest that the simple OLS model continues to yield at least equal or superior hedging effectiveness when compared to more sophisticated GARCH models. Myers (1991) employ rolling moving average estimates of variance and covariance, and a GARCH model for modelling the spot and futures returns and find only marginal improvement over the traditional constant methods of hedge ratio estimation. Miffre (2004) introduce a conditional OLS model than incorporates a set of informational variables and test whether the basis risk is reduced. They find that, especially traders with long-term horizon, gain substantial benefit in terms of portfolio risk reduction, and the conditional OLS model performs better than the naïve, the static OLS and the GARCH(1,1) model. Cotter and Hanley (2012) compare the out-of-sample hedging performance of the GARCH and Asymmetric GARCH with the OLS hedge ratio and find that the latter perform consistently well. In contrast, the integration of regime switching dynamics in the GARCH process has been documented to yield better results. Lee (2009) develops a GARCH model that incorporates the impact of jumps and regime switching states in the evolution of the variance process and suggests that the out-of-sample hedging effectiveness is significantly improved. Lee and Yoder (2007) suggest a Markov Regime Switching time-varying correlation GARCH model and find similar results.

More recently, the availability of high-frequency data has enabled the direct estimation of the variance and covariance processes, which essentially become observable. Additionally, it has been well documented that, the realized measures exhibit superior forecasting ability when compared to estimators from daily or lower frequency data. While the notions of realized variance and covariance have been extensively studied, not much attention has been paid on the properties, forecasting performance and hedging effectiveness of the realized hedge ratio, while the empirical evidence is mixed. Lai and Sheu (2008; 2010) argue that the use of realized volatility measures provides substantial benefit to the hedger in terms of risk reduction and economic value. Harris, Shen and Stoja (2010) employ four models of conditional variance (namely, the exponentially weighted moving average, EWMA, the Diagonal VECH, the constant correlation and the S-GARCH models) and one dynamic model of the realized hedge ratio. They use data on three foreign exchange rates and find that the realized hedge ratio yields superior hedging effectiveness whereas the conditional models provide only marginal improvement over the unconditional equivalent. In contrast, McMillan & Garcia (2010) advocate that the portfolio variance is minimized when the hedge ratio is estimated from daily returns while the realized hedge ratio yields superior Sharpe ratio.

Chapter 3

The Dynamic Evolution of the Correlation Risk Premium

Chapter Abstract: In this Chapter, I study the distributional properties of the Correlation Risk Premium (CRP), defined as the difference of realized and implied correlation. To this end, an improved model-free measure of implied correlation is proposed. The new measure is firstly inferred from currently observed option prices without relying on any specific option pricing model, and secondly, corrected for the early exercise premium of the American options. An extensive dataset of fifteen years, from 1996 to 2010 is used to investigate potential differential properties during alternative periods of financial distress. Results from the distributional analysis of CRP suggest that the series are far from normal while exhibit high persistence and long memory. Interestingly, the negative correlation risk premium is no longer significant during the subprime mortgage crisis suggesting that the investors were not able to exploit previously reported arbitrage opportunities, arising from the discrepant pricing of options on the index and options on the constituent stocks. Finally, a set of macroeconomic and market-specific variables fails to provide accurate forecasts of the CRP series. However, I find that forecasting ability of the volatility risk premium of the index for future CRP levels is significant.

3.1 Introduction

Understanding the dynamics of risk, embodied in the time-variation of price returns, is of paramount importance for investors when managing risk, allocating assets, pricing and hedging derivative securities. The importance of correlation as a priced risk factor has been the subject of only a few recent studies. Krishnan, Petkova and Ritchken (2009) use historical prices of US stocks and, after controlling for asset volatility, risk factors and higher-order moments, find a significant correlation risk premium tested under different specifications. Driessen, Maenhout and Vilkov (2009) create a correlation trading strategy that aims to exploit the correlation risk premium and find that the compensation for bearing correlation risk is substantial. The strategy sells index straddles and buys individual straddles and stocks in order to hedge individual variance risk and stock market risk, respectively. In a similar context, Mueller et al. (2012) provide evidence of priced correlation risk in currency markets using option-implied correlations. Pollet & Wilson (2010) examine correlation risk using historical price data and find that differences in exposures to correlation risk justify differences in expected returns, while volatility risk is negatively related or unrelated to the stock market risk premium. Kelly, Lustig, and Nieuwerburgh (2011) examine the differential pricing of put options on the finance sector index and the individual firms of the sector during the recent financial turmoil, through the price of correlation risk and find that the industry specific risk has been partially eliminated by the government bailout guarantee.

While volatility risk premium has stimulated the academic interest over the past years, correlation risk premium has attracted less attention. The few relevant studies have been intrigued by prior empirical findings that the volatility risk premium of an index option is higher than the corresponding premium of a portfolio consisting of the options on the constituent stocks of the index. The discrepancy has been initially attributed to market frictions and specifically, the differential demand and supply dynamics for the index and individual stocks. Bollen and Whaley (2004) proved that daily changes of implied volatility can be attributed to increased buying pressure, especially in OTM puts for the index options and in, much less extent, for calls options on the individual stocks. Gârleanu, Pedersen, and Poteshman (2009) model the option equilibrium prices as a function of demand pressure and find that end users take long (short) positions in index (individual) options, leading the options' price to increase (decrease). From an alternative perspective, Driessen, Maenhout and Vilkov (2009) provided a risk-based explanation for the contrast of volatility risk premium between index and individual options. Essentially, they argue that index options are expensive and earn low returns, unlike individual options, because they offer a valuable hedge against correlation increases and insure against the risk of a loss in diversification benefits.

Motivated by the initial finding of Jackwerth and Rubinstein (1996) that at-the-money implied volatility is higher than its realized counterparty, volatility risk premium has stimulated the interest of the academics and practitioners over the past years, with the vast majority of the studies focusing on the study of index options. Negative price of volatility risk premium suggests that investors are willing to pay an extra premium so as to be hedged against downward market movements in stock returns. Bakshi and Kapadia (2003a) amongst others, provide evidence of negatively priced volatility risk, suggesting that investors are willing to pay an increased premium for the options, to protect themselves from increased volatility and lower returns in the market. On the contrary, only a few papers study the pricing of individual equity options. Bakshi and Kapadia (2003b) noted that individual equity option prices embed a negative market volatility risk premium, although much smaller than for the index option.

Driessen, Maenhout and Vilkov (2009) have provided evidence of negatively priced volatility risk premium in index options, whereas no such evidence is found for individual options. He suggested that the discrepancy is triggered by the fact that, by definition, index options are exposed to a risk factor lacking in the individual process, namely correlation risk. To this end, they provided solid theoretical background by decomposing the volatility risk premium of the S&P 100 index into individual stocks' volatility risk premium and inter-asset correlation structure and, through the implementation of a dispersion trading strategy provide empirical evidence of priced correlation risk. Additionally, Cosemans (2011) shows that the aggregate volatility risk premium is driven by the changes in the correlation risk premium structure as well as the market risk component inherited in the idiosyncratic risk. Moreover, both volatility and correlation processes are separated into short and long run components and finds that even after accounting for size, value, momentum, and liquidity factors, the price remains significant.

In this Chapter, the distributional properties of the correlation risk premium, defined as the difference of implied and realized correlation, are examined. The first contribution of the Chapter lies in the derivation of the Model-Free Implied Correlation measure. In specific, an adjustment of the model-free implied volatility, firstly proposed by Britten-Jones and Neuberger (2000), is proposed so as to account for the Early Exercise Premium inherent in the American options. An extensive dataset of fifteen years is used, from 1996 to 2010, which allows the examination of the time series behavior during several recent periods of financial distress with scrutiny. Motivated by the decomposition proposed by Driessen, Maenhout and Vilkov (2009), the distributional properties of the alternative measures of volatility and correlation, namely implied and realized, along with the properties of the volatility risk premia for the S&P 100 Index and the portfolio of constituent stocks and correlation risk premia are presented. Following, prompted by several stylized facts about correlation, reported in the literature, and analyzed extensively in Chapter 2, the distribution of the correlation risk premium is investigated for long memory and asymmetric traits. Finally, Intrigued by previous studies that have documented superior information content of economic and financial variables in shaping equity returns, the importance of macroeconomic and market-specific variables in modelling and forecasting the correlation risk premium is examined.

The remaining Chapter is organized as follows: Section 2 describes thoroughly the methodology followed for the derivation of the Model-Free Implied Correlation series. Section 3 and 4 present the data and discuss the time evolution of the series. In Section 5, the CRP series is examined for seasonality patterns, long memory and asymmetric response to index returns and volatility risk premium. Section 6 examines the informational content of several variables in forecasting the CRP. Section 7 concludes.

3.2 Construction of Model-Free Implied Correlation Index – Methodology

Intuitively, fluctuations of the index options' prices are associated with fluctuations of prices of individual options as well as fluctuations in the correlation structure. Thus, the variance of an index, consisting of N stocks, can be defined as follows:

$$\sigma_{p,t}^2 = \sum_{i=1}^{n} w_{i,t}^2 \sigma_{i,t}^2 + 2\sum_{i=1}^{n-1}\sum_{j>i} w_{i,t} w_{j,t} \rho_{ij,t} \sigma_{\iota,t} \sigma_{j,t} \tag{3.1}$$

where σ_p is the variance of the index, σ_i is the standard deviation of the asset *i*, ρ_{ij} is the correlation between asset *i* and *j*, and w_i is the relative weight of each constituent asset on the portfolio.

The correlation measure examined in this study is a weighted average of all pair-wise correlations of the constituents of the index and an average measure of the degree of diversification in the market represented by the index (see Skintzi and Refenes, 2005), defined as follows:

$$\rho_t = \frac{\sigma_{p,t}^2 - \sum_{i=1}^{n} w_{i,t}^2 \sigma_{i,t}^2}{2\sum_{i=1}^{n-1}\sum_{j\neq i} w_{i,t} w_{j,t} \sigma_{\iota,t} \sigma_{j,t}} \tag{3.2}$$

Historical volatility and correlation estimators rely on the assumption that the future dynamics of the series will be similar to the past. To overcome the ambiguities deriving from the above assumption, stock return moments are inferred from currently traded option prices. Naturally, option prices reflect the current market view of future price movements of the underlying stock. Therefore, the estimation of stock price distribution moments from option prices is widely considered to outperform the estimation based on historical performance in terms of informational efficiency.

Unlike the traditional, model-based Black-Scholes implied volatility, Britten-Jones and Neuberger (2000) have developed a non-parametric specification of implied variance under the assumption that price process is continuous while the volatility is stochastic. Jiang and Tian (2005) extended the methodology so as to account for the presence of jumps in the price process of the underlying asset. Specifically, they provided evidence that, under the risk-neutral measure, Q, the integrated variance between T_1 and T_2 is fully specified by a set of call options, expiring on the specified dates:

$$E^Q\left[\int_{T_1}^{T_2}\left(\frac{dS_t}{S_t}\right)^2\right]=2\int_0^{\infty}\frac{C(T_2,K)-C(T_1,K)}{K^2}dK\approx \sum_{i=1}^{m}\frac{C^F(T,K_i)-\max(0,F_0-K_i)}{K_i^2}+\frac{C^F(T,K_{i-1})-\max(0,F_0-K_{i-1})}{K_{i-1}^2}\Delta K \tag{3.3}$$

where $C^F(T_i,K_i)$ is the forward price of a call option, with a strike price equal to K_i expiring at time T_i, F_0 is the forward asset price, $\Delta K=(K_{\max}-K_{\min})/m$ and $K_i=K_{\min}+i\Delta K$, for $0\le i\le m$. Assuming deterministic interest rates and dividends, the forward option price is the future value of an option expiring at time T_i, defined as $C^F(T_i,K_i)=C(T_i,K_i)\times e^{rT_i}$. Accordingly, the forward asset price is equal to $F_0=S_0\times e^{rT_1}$, where S_0 is the current asset price minus the present value of all the expected dividends within the time-to-maturity of the option, T_i.

The above equation gives rise to two main implementation issues. First, the trapezoid rule for numerical integration leads to discretization errors. However, according to Jiang and Tian (2005), the choice of a large number of integral points, m, leads to negligible errors. Second, only a finite number of strike prices, [K_{min}, K_{max}], is available within a trading date. In order to overcome the limitations, following Jiang and Tian (2005), a fine grid, of m=1000 points, with option prices within the moneyness range of [0.3, 3] is created. Tests under alternative number of integral points, m, and moneyness range, have produced insignificant changes.

In detail, the implementation procedure is as follows. First, the moneyness level for every option in the dataset is calculated. Taking into account that the options of the constituent stocks of the S&P100 index are American style, the current underlying stock price is adjusted by subtracting the present value of known dividends due to payment within the expiration of the option. In case of duplicate moneyness levels, i.e. a call and a put option with the same strike price, the average of implied volatilities is computed. In order to calculate the mode-free implied volatility every day, at least three unique implied volatility points of an option with a specific time-to-maturity are required. In the next step, I fit a Piecewise Cubic Hermite Interpolating Polynomial to create a fine grid of implied volatilities corresponding to the range of moneyness previously defined. For any moneyness level beyond the existing bounds, the last known implied volatility value on this boundary is used. Thereafter, 1000 points of implied volatility are translated into call prices using the Black-Scholes model and the equation (3.3) is applied. Finally, model-free implied volatility is annualized on a 30/365 basis.

The model-free implied volatility methodology, as developed by Britten-Jones and Neuberger (2000), relies on the assumption of European style options. In contrast, the dataset consists of the options on the S&P 100 index as well as the options on the constituent stocks, all of which are American style, thus, their incorporating the Early Exercise Premium (EEP). Among the prevailing methodologies of extracting the EEP is the Barone-Adesi and Whaley (1987) approximation. However, their methodology could be only applied to the options on the S&P100 index, which have a continuous dividend yield. Recently, Tian (2011) proposed a methodology of extracting the risk-neutral density and, subsequently, the EEP from American options with either discrete or continuous dividends. He proposed an iterative Implied Binomial tree (*i*IB), as a modification of the standard implied binomial tree proposed by Jackwerth and Rubinstein (1996). The main steps of the methodology are outlined below.

At the first step, the European option price is considered equal to the observed American prices. Next, a standard implied binomial tree is built and the ending nodal probabilities are calibrated to the initial option prices. Due to the limited availability of strike prices from traded options and with the view of maintaining the non-parametric feature of the methodology, I follow Jackwerth and Rubinstein (1996) and choose the ending nodal probabilities through the following optimization procedure that maximizes the smoothness of the implied risk-neutral density.

$$\min_{P} \sum_{j=1}^{n-1} (P_{j-1} - 2P_j + P_{j+1})^2 + \alpha \sum_{i=1}^{m} (V_i^{\text{model}} - V_i^{\text{mkt}})^2 \qquad (3.4)$$

subject to:

a) $\sum_{j=0}^{n} P_j = 1$ where $P_j \geq 0$ and

b) $S_0 = \exp\left[-(r-q)T\right]\sum_{j=0}^{n} P_j S(j)$ (In case of the index options with a continuous dividend yield equal to q), or

$S_0 - D_0 = \exp(-rT)\sum_{j=0}^{n} P_j S(j)$ (In case of the options of the constituent stocks that pay discrete dividends equal to D_0)

where P_j is the ending nodal probability, V_i^{model} is the price of the option calculated from the implied binomial tree model, V_i^{mkt} is the observed option price, r is the risk-free rate, S_0 and S_j is the current asset price and the asset price at the ending node, respectively.

After the refined European option price from the implied binomial tree is obtained, the early exercise premium is calculated as the difference of the American and the European option price. Next, the refined estimate of European option price is defined as the difference between the initial European price and the early exercise premium, and used as an input for the second iteration where the above optimization procedure is replicated. The iterative procedure continues until the early exercise premium converges.

Additionally, a daily average realized correlation measure based on equation (3.2) using forward-looking realized volatility estimates for σ_i, at time t is calculated. Realized volatility is calculated as a simple measure of standard deviation of squared returns, calculated from closing prices, as follows:

$$RV_{t,T} = \sqrt{\frac{365}{T}\sum_{t=1}^{T}\log(\frac{S_t}{S_{t-1}})^2} \qquad (3.5)$$

where $S_{i,t}$ is the closing price of asset *i*, *i = 1, 2, …, N, I* on day t and T is set equal to 30 calendar days so as to match the time horizon of model-free implied volatility.

The Variance Risk Premium (VRP) is defined as the raw difference between the model-free implied volatility, a proxy for risk-neutral measure, and realized volatility, the volatility under the physical measure, i.e.

$$VRP_{t,i} = MFIV_{t,i} - RV_{t,i} \tag{3.6}$$

where MFIV is the Model-Free Implied Volatility, as computed from equation 3.3, and RV is the Realized Volatility, computed from equation 3.5.

By analogy, the Correlation Risk Premium (*CRP*) is thus defined as:

$$CRP_t = RC_t - MFIC_t \tag{3.7}$$

where RC_t is the Realized Correlation at day t calculated from equation (3.2) using realized volatility estimates and $MFIC_t$ is the Model-Free Implied Correlation at day t calculated also from equation (3.2) using model-free implied volatility estimates.

3.3 Data

For the purposes of the study, daily closing quotes for the S&P100 index options and for the constituent stocks index are retrieved from OptionMetrics. As a subset of S&P 500, S&P 100 is a capitalization-weighted index including the largest and most established companies across several industries with traded options. Additions/deletions to the index occur on an as-needed basis, whenever any of the S&P's inclusion criteria are violated. Data on underlying asset price and company distributions is obtained from OptionMetrics database. The zero-coupon interest rate curve, also provided by OptionMetrics database, is used as the risk-free interest rate. For the purpose of calculating the realized volatility, which requires stock prices adjusted for splits and distributions, stock adjusted closing prices are obtained from Datastream. The sample extends from January 1996 to October 2010. The list of constituents stocks is retrieved from Datastream and the composition of the S&P 100 Index is replicated on a daily basis; thus, the sample includes 148 stocks throughout the fifteen years of the sample. On each day, the actual daily weight of stock *i* is calculated based on its daily market capitalization divided by the total market capitalization of the *N* stocks included in the index.

During the sample period, the world economy experienced several periods of financial turmoil. Starting with the Asian and Russian crisis in late 1997 and 1998, respectively, the 9/11, the Argentina financial turmoil and the “dot.com” bubble followed during the first years of the decade. The next and far more severe period of financial distress that the global economy experienced was the subprime mortgage crisis that started in 2007. The partition of the whole sample period of fifteen years into three subsets of five years each, allows the study of differential pattern of VRP and CRP across distinct periods. To this end, the three sub-periods extends from 1996 to 2000, from 2001 to 2005 and from 2006 to 2010.

Several filtering rules are applied. First, in order to ensure trading activity, options with bid price greater than zero and positive open interest are selected. Second, options with less than one-week remaining time to maturity are eliminated; such options are more vulnerable to liquidity and microstructure issues. Third, options with implied volatility either greater than one or missing implied volatility are discarded[2]. Additionally, in order to eliminate options with extreme moneyness levels, only call options with delta greater than 0.15 and smaller than 0.5 and put options with delta greater than -0.5 and less than -0.05 are included in the analysis. Finally, any option that violates arbitrage bounds is excluded.

Before proceeding with the properties of the correlation risk premium, Table 3.1 reports the average values of early exercise premium throughout the sample period, as derived from the methodology outlined in Section 3.2. I find that the average EEP is equal to 0.51%, while put option report an average value of EEP equal to 0.6%. Obtained results are consistent with previous literature that suggests that options with longer maturity incorporate large value of early exercise premium. Dueker and Miller (2003) record the early exercise premium for a period of 02/04/1986, through 20/06/1986 and find evidence of a substantial early exercise premium, with an average value of 5% for

[2] OptionMetrics does not report implied volatility for an option in case of a special settlement, failure of the implied volatility process to convergence, unavailability of the underlying asset price, vega value being below 0.5 or in cases where the midpoint of bid/ask price is below intrinsic value.

calls and 9.5% for put options. The pronounced discrepancy with results of this analysis may be stemming from the different periods of the employed dataset.

Table 3.1: Early Exercise Premium (EEP)

	T<=30	T<=60	T<=180	T<=365	Total
Put options	0.23%	0.31%	0.43%	0.49%	0.60%
Call options	0.28%	0.32%	0.33%	0.34%	0.36%
All options	0.24%	0.31%	0.39%	0.43%	0.51%

Note: The Columns present the Early Exercise Premium during the whole sample period for call and put options of different maturity (T=30, 60, 180, 365).

3.4 Time evolution and statistical properties

In this Section, the time evolution of Model-Free Implied Correlation (MFIC), Realized Correlation (RC) and Correlation Risk Premium (CRP) for the whole sample period spanning from January 1996 to October 2010 is presented. Figure 3.1 plots the risk neutral (model-free) and physical (realized) measures of correlation in conjunction with the S&P 100 index returns. Notably, the risk-neutral correlation trends higher than the realized correlation. Consistent with the well-documented fact of increased correlation during bear markets, both series present several spikes throughout the whole sample period coinciding with periods of high volatility markets, though MFIC presents greater jumps. Specifically, resulting from the Asian and the Russian financial crises, the MFIC series reached the value of 0.9, in late 1997 and later in August 1998, while the RC was significantly lower, around 0.65. During the first years of the ‘00s, the low returns market stemming from the 9/11 terrorist attack, the South-American financial turmoil and the internet bubble burst, resulted in increased levels of correlation, though lower than the aforementioned spikes of the 90s. Over the subsequent years, MFIC trended lower, with the exception of the period of the multi-market sell-off in May 2006. Finally, the MFIC presents several peaks from 2007 until the end of the sample, as a result of the global financial crunch, which commenced in 2007 and culminated in September 2008 with the announcement of the bankruptcy of Lehman Brothers. MFIC series reached the highest value of 0.97 on October 24, 2008, while RC reached its peak of 0.77 on November 11, 2008.

Figure 3.1: Time evolution of Model-Free Implied Correlation, Realized Correlation and S&P 100 Index Returns

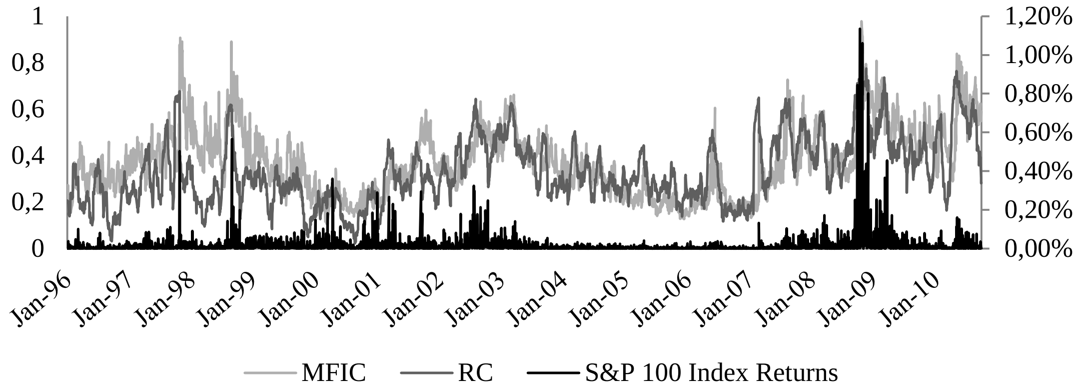

Table 3.1 reports the summary statistics of the model-free implied and realized correlation series for three sample periods described above. First, a significant difference between the value of correlation estimates under the risk-neutral and the physical measure is apparent. Over the first sub-sample, the average MFIC is substantially greater than the realized correlation, resulting in a significant value of correlation risk premium (CRP) equal to -0.11. Over the subsequent period, Realized Correlation increased sharply and reached the value of 0.37, while the average value of MFIC did not change substantially. The MFIC and the RC series are positively skewed with fat tails in contrast to the CRP that is negatively skewed, except the sub-period of 2006-2010. Additionally, the first-order autocorrelation for the CRP is significantly high and above 90%. However, the coefficients for the tenth and the fifteenth lag indicate rather fast decaying persistence, in contrast to MFIC and RC series where the coefficients decay slower. Further investigation of the autocorrelation pattern is conducted in the Section that follows.

The higher moments and the Jarque-Bera test reject the hypothesis of normality for the MFIC series for all the periods under consideration. Table 3.1 also reports the test statistics of the Augmented Dickey-Fuller (ADF) test, the Phillips-Perron (PP) test and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test statistic. According to the PP and the ADF test, the null hypothesis of unit root is strongly rejected for all three subsets at a significance level of 1%. In contrast, the KPSS test provides evidence of non-stationary time series for all three periods under consideration. These mixed results, along with the slow decaying and significant autocorrelation coefficients, suggest the existence of long memory properties in the CRP similar to model-free implied volatility (see Brooks & Oozeer, 2002; Ahoniemi, 2008). To formally test the existence of long memory properties an ARFIMA model is employed. The detailed specification and results are provided in Section 3.5.2. Results suggest that the fractional differencing parameter is significant and within the range of [0, 0.5] only for the whole sample period indicating that the CRP is stationary for the three sub-periods while exhibiting long-memory traits during the extended sample period of fifteen years.

Panel B of Table 3.2 displays the cross-correlation structure between the changes in the Correlation Risk Premium and the log returns of the S&P 100 Index. To simplify the practical interpretation of the negative price of Correlation Risk Premium, an increase (decrease) of the returns signifies that the premium becomes more negative (positive), i.e. the difference between realized and model-free implied correlation widens (narrows). All cross-correlation coefficients are significant for the first lagged and the contemporaneous return. Additionally, the negative relationship for the lagged CRP and index returns suggests that a decrease in index returns at time *t-1* will induce CRP to increase. Interestingly, the relationship reverses in the contemporaneous setting. To fix ideas, an observed decrease of the index return at *t-1* will result in a decrease of the CRP at time *t*, whereas a decrease of the index returns at time t will induce the CRP at time *t* to decline. The positive relationship for the contemporaneous returns is in line with previous literature of volatility risk premium with the following line of reasoning: an increase of index returns will decrease volatility (and correlation) based on the long documented leverage effect, which will in turn, increase the volatility risk premium, i.e. become more negative. By analogy, consider that the investors pay a negative price for the correlation risk premium as an insurance against unexpected correlation increases, while, they capitalize their gains during states that correlation is low. Given that an increase of index returns will decrease correlation, the investors are willing to uptake an extra compensation so as to benefit from the low correlation state. Finally, the CRP will increase and obtain values that are more negative.

Turning the focus to the Correlation Risk Premium, as defined in Equation 3.7, it is noteworthy that the difference of correlation under the two measures is eliminated leading the CRP to its minimum value during the period that encompasses the subprime mortgage crisis. Moreover, the last row of the Table reports the p-values for the null hypothesis that implied and realized correlation are on average equal, i.e. CRP is zero. Importantly, the correlation risk premium is significant only in the first period of 1996-2000 and the whole sample period of 1996-2005, whereas it is insignificant for the two sub periods that encompass the South-American crisis, the internet bubble burst and the subprime crisis. The finding is not surprising. As outlined above, correlation is expected to increase during periods of financial distress whereas the Correlation Risk Premium is expected to decrease and reach less negative values.

Table 3.2: Descriptive Statistics of Model-Free Implied Correlation (MFIC), Realized Correlation (RC) and Correlation Risk Premium (CRP)

	1996 – 2000			2001 – 2005			2006 - 2010			1996 - 2010		
	MFIC	RC	CRP	MFIC	RC	CRP	MFIC	RC	CRP	MFIC	RC	CRP
Mean	0.360	0.248	-0.111	0.341	0.337	-0.004	0.419	0.415	-0.002	0.373	0.332	-0.040
Median	0.349	0.239	-0.104	0.334	0.311	0.001	0.420	0.422	-0.011	0.356	0.304	-0.034
Maximum	0.907	0.677	0.324	0.662	0.644	0.250	0.977	0.774	0.446	0.977	0.774	0.446
Minimum	0.106	0.021	-0.602	0.126	0.084	-0.295	0.132	0.117	-0.335	0.106	0.021	-0.602
Std. Dev.	0.135	0.114	0.126	0.112	0.102	0.096	0.167	0.156	0.128	0.143	0.143	0.128
Skewness	0.808	1.091	-0.245	0.393	0.534	-0.214	0.167	0.037	0.560	0.594	0.571	-0.090
Kurtosis	3.868	5.097	3.965	2.581	2.799	3.010	2.276	2.239	3.811	3.043	2.887	4.069
Jarque-Bera	176.120***	479.859***	61.375***	41.530***	61.902***	9.556***	32.182***	29.065***	95.027***	219.376***	203.281***	181.558***
ρ_1	0.949***	0.976***	0.926***	0.964***	0.977***	0.931***	0.959***	0.982***	0.919***	0.960***	0.984***	0.937***
ρ_{10}	0.778***	0.725***	0.490***	0.847***	0.747***	0.558***	0.834***	0.819***	0.544***	0.836***	0.831***	0.607***
ρ_{15}	0.701***	0.571***	0.252***	0.821***	0.599***	0.413***	0.789***	0.688***	0.360***	0.792***	0.723***	0.443***
ADF	-3.167**	-3.894***	-5.859***	-1.897	-3.776***	-5.472***	-3.350**	-2.999**	-5.225***	-4.410***	-5.393***	-8.847***
PP	-4.964***	-4.460***	-7.001***	-3.341**	-4.169***	-6.315***	-3.869***	-3.655***	-6.199***	-7.477***	-6.244***	-10.346***
KPSS	1.192***	0.613**	0.424*	1.702***	0.783***	0.837***	2.818***	1.356***	1.070***	0.767***	2.391***	2.424***
D			0.243			0.000			0.040			0.224***
p-value for H0: CRP= 0			0.000***			0.797			0.901			0.002***
Panel B												
						CRP Difference						
S&P 100 Log Returns												
-2			-0.004*			-0.057			-0.011			-0.020
-1			-0.165**			-0.114*			-0.125*			-0.135*
0			0.436***			0.304**			0.356 **			0.368**
1			-0.014			-0.016			-0.054			-0.032
2			0.015			-0.002			-0.054			-0.018

Note. The Table reports summary statistics of the MFIC, RC and CRP for three subsets defined in the first row and the whole sample period. The values ρ_1, ρ_{10} and ρ_{15} are the autocorrelation function (ACF) coefficients for the 1st, 10th and the 15th lag. The Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) test for the presence of unit root. The parameter *d* refers to the fractional integration parameter form the ARFIMA specification. The p-value is for the null hypothesis of CRP=RV – MFIV = 0 estimated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

To analyze further the CRP, I take an insight view at the volatility risk premium embedded in individual and index options. Figure 3.2 depicts the time series of the cross-sectional weighted average of implied and realized volatility for the individual equity options while Figure 3.3 plots the implied and realized volatility for the S&P 100 index options. The well-established fact that implied volatility is higher than realized volatility also holds throughout the sample. Visual inspection also suggests that the difference of the two volatility measures is more pronounced for the index returns.

Figure 3.2: Time evolution of Weighted Average MFIV and RV of individual stocks.

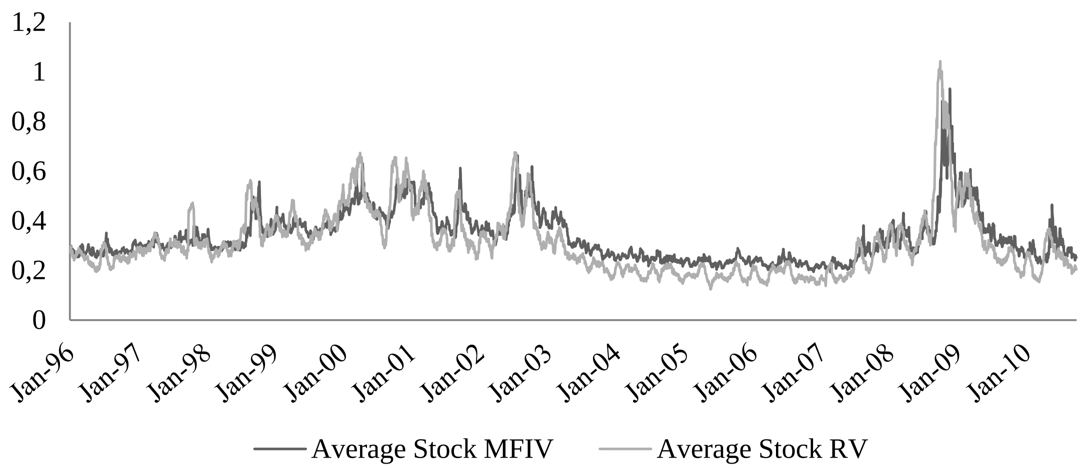

Figure 3.3: Time evolution of Weighted Average MFIV and RV of the S&P 100 Index.

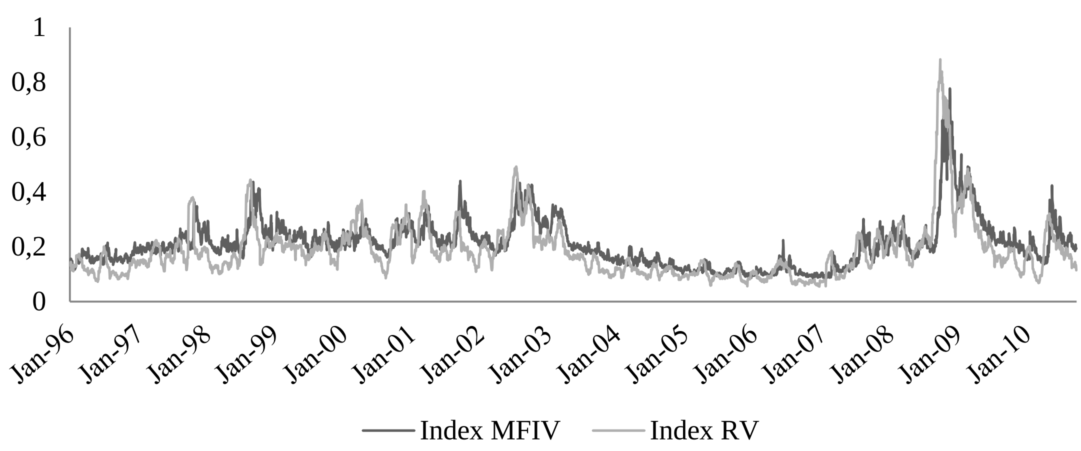

Table 3.3 reports the average value of the model-free implied volatility and realized volatility for the S&P 100 as well as the average values for the cross-sectional weighted average of the model-free implied and realized volatility of the individual

stocks. During the first part of the sample, the difference between realized and implied volatility for the index options and the weighted average of the individual options is equal to -2.9% and -3.1%, respectively. During the second sub-period of 2001 – 2010, the realized correlation of the portfolio of individual options decreased by 20% while the changes in the index volatility measures are less pronounced. Turning to the weighted average of the individual stock options over the period of 2001 – 2010, the average implied and realized volatility reduced to the level of 32.5% and 28.8%, respectively. Although the volatility of individual options, under both measures, increased drastically during the last quarter of 2008 reaching a maximum value of approximately 100%, the diminished average value suggests that the jump in the series stemming from the turbulent times of 2007-2009, faded out quickly.

In line with the obtained results of Bakshi & Kapadia (2003b), a lower value for the volatility risk premium of the individual options when compared to the index volatility risk premium is observed. Consistent with the findings for the CRP, the null hypothesis of equal implied and realized volatility of the portfolio of individual options is accepted during the subprime mortgage crisis period, both for the index and the individual equity options. Interestingly, the null hypothesis of zero risk premium is also accepted for the individual options during the first sub-period

Table 3.3: Volatility Risk Premium

	1996 - 2000	2001-2005	2006 - 2010	1996 - 2010
IV_m	0.215	0.200	0.213	0.210
RV_m	0.187	0.170	0.201	0.186
VRP_m	-0.029	-0.030	-0.013	-0.024
p-value $H0: VRP_m = 0$	0.000	0.000	0.415	0.000
IV_i	0.356	0.331	0.317	0.335
RV_i	0.359	0.282	0.295	0.312
VRP_i	-0.031	-0.049	-0.027	-0.045
p-value $H0: VRP_i = 0$	0.736	0.000	0.198	0.003

Note: Subscript *m* denotes the S&P 100 index and subscript *i* the weighted average of the constituent stocks. The reported *p-values* are for the null hypothesis of zero volatility risk premium, and have been estimated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

3.5 Distributional properties of CRP

3.5.1 Intraweek patterns

The intraweek variation has been a well-documented pattern for the returns and implied volatilities. It is thus natural to examine the presence of such patterns in the evolution of the correlation risk premium. To this end, the following specification is employed, where one lag of CRP has been added to control for autocorrelation.

$$CRP_t = c + \sum_{i=1}^{4} a_i D_{i,t} + CRP_{t-1} + \varepsilon_t \tag{3.8}$$

Results reported in Table 3.4 provide evidence of strong intraweek pattern for every sub-period under examination, although less intense during the period of 2006 – 2010.

Table 3.4: Intraweek patterns of CRP

	1996 - 2000	2001 - 2005	2006 - 2010	1996 - 2010
Constant	0.007**	0.011***	-0.005*	0.007***
Monday	-0.038***	-0.026***	0.006	-0.019***
Tuesday	-0.013***	-0.013***	0.007*	-0.008***
Wednesday	-0.010**	-0.009***	0.002	-0.006***
Thursday	-0.011**	-0.008***	0.008*	-0.004*
CRP_{t-1}	0.931***	0.936***	0.924***	-0.146***

Note: The regression has been estimated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

3.5.2 Long Memory

Results from the ADF, PP and KPSS test have provided indirect evidence of long memory in the CRP series. In addition, several studies have suggested the presence of persistence and long-memory in correlation (see Dacorogna, 1999; ABDL, 2001; ABDE, 2001). To this end, an ARFIMA (p,d,q) model suggested by Granger and Joyeux (1980) is implemented. The specification is as follows:

$$(1-\sum_{r=1}^{r}\varphi_r L^r)(1-L)^d(y_t-\mu) = (1+\sum_{m=1}^{m}\vartheta_m L^m)\varepsilon_t \tag{3.9}$$

The model is estimated for all plausible combination of AR and MA terms up to lag five. The ARFIMA (1,d,1) model, for all subsets and the whole sample period, is the best performing model that minimizes the BIC criterion. The fractional differential parameter is significant only for the whole sample period, with the value lying between [0, 0.5], thus signifying the presence of long-memory and positive dependence between distant observations. In the case of sub-periods, *d* is not significantly different from zero suggesting that the series is stationary and a simple ARMA model is able to accurately model the short-run dynamics of the series. For the rest of the Chapter, the notation "CRP" will refer to the first difference, in case of the whole sample period, and to the levels, for every sub-period.

As an additional test, I also plot the autocorrelation function for the first 100 lags in Figure 3.4. The long-term persistence is confirmed by the high and significant autocorrelation coefficients that, although decaying fast up to the 30th lag, remain stable for the lags that follow. A possible explanation for the significant ACF values for the first almost 30 lags is that by construction, the model-free implied correlation and the realized correlation, are annualized on a calendar month basis.

Figure 3.4: Autocorrelation function of CRP (level)

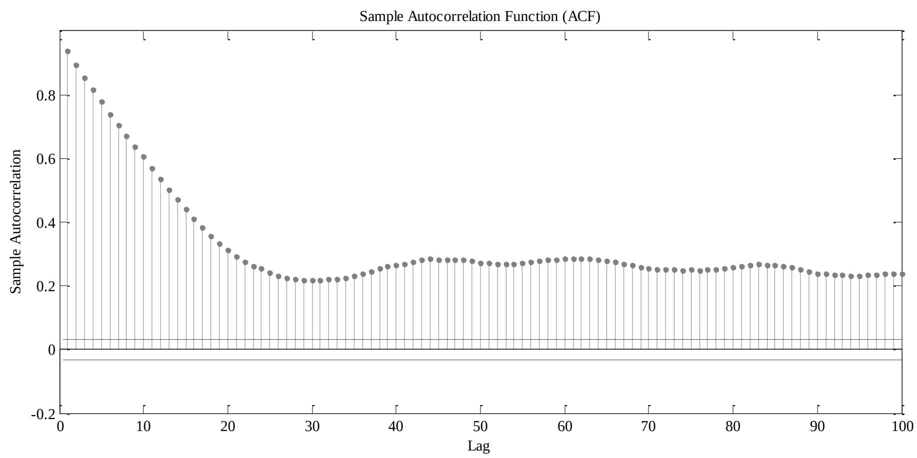

Note: The Figure plots the autocorrelation function for the CRP for the entire sample period (1/1/1996 – 29/10/2010). The dotted lines represent the upper and lower bound of the Bartlett confidence interval, at 95% confidence level.

Based on the findings from the ARFIMA model, I also plot the autocorrelation function for the first difference of the CRP. Figure 3.5 confirms that the first differencing eliminates the persistence and the long run dependence of the series.

Figure 3.5: Autocorrelation function of CRP (first difference)

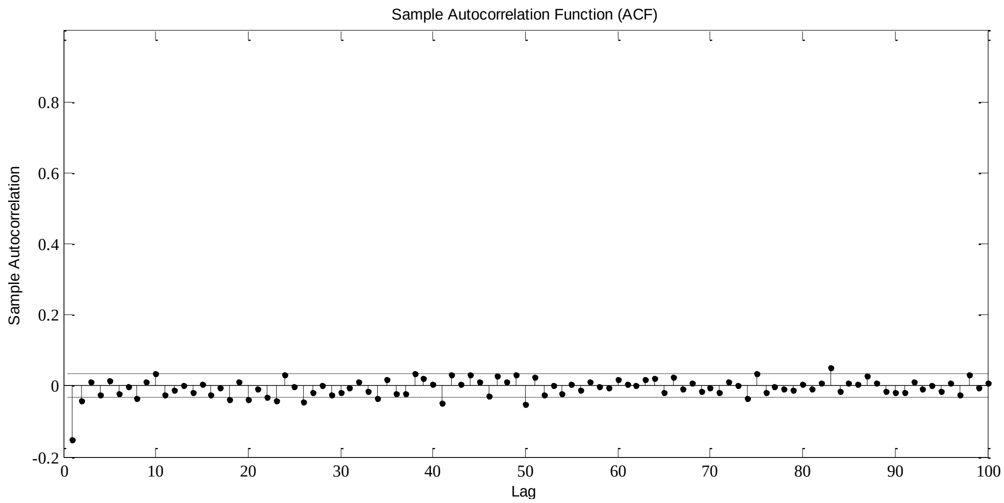

Note: The Figure plots the autocorrelation function for the first difference of the CRP for the entire sample period (1/1/1996 – 29/10/2010). The dotted lines represent the upper and lower bound of the Bartlett confidence interval, at 95% confidence level.

3.5.3 Asymmetry

A number of studies have provided evidence of asymmetric response of correlation to positive and negative returns (see Section 2.2 for the relevant literature). To assess the impact of positive and negative returns to the correlation risk premium the following specification is employed:

$$CRP_t = c + \sum_{i=1}^{3} \beta_i CRP_{t-i} + \sum_{k=0}^{3} \gamma_k R_{t-k} + \sum_{k=0}^{3} \delta_k \left| R_{t-k} \right| \qquad (3.10)$$

where R_{t+k} is the lagged index return and $|R_{t+k}|$ is the absolute index return over the same time interval. The lagged values of CRP are included to account for remaining autocorrelation.

Table 3.5: Asymmetric Relationship of CRP and S&P 100 Index Returns

	1996 - 2000	2001 - 2005	2006 - 2010	1996 - 2010
c	0.0021	0.0105 ***	0.0074 ***	0.0086 ***
β_1	0.8560 ***	0.8942 ***	0.8399 ***	-0.1284 ***
β_2	0.1175 ***	0.0495	0.0889 *	-0.0265
β_3	-0.0120	0.0068	0.0318	0.0056
γ_0	0.0027 ***	0.0017 ***	0.0025 ***	0.0024 ***
γ_1	-0.0006 ***	-0.0005 ***	-0.0004 *	-0.0005 ***
γ_2	-0.0003 **	-0.0005 ***	-0.0002	-0.0003 **
γ_3	-0.0001	-0.0003 **	-0.0002	-0.0002 *
δ_0	-0.0001	-0.0004 *	-0.0002	-0.0003 *
δ_1	-0.0021 ***	-0.0025 ***	-0.0024 ***	-0.0023 ***
δ_2	0.0000	0.0004 *	0.0007 **	0.0004 **
δ_3	0.0009 ***	0.0002	0.0004	0.0005 ***
Adj. R^2	0.8898	0.8952	0.8839	0.2322
LL	2201.2067	2581.2833	2053.5159	6713.7470
BIC	-3.4452	-4.0421	-3.3714	-3.6005
DW	2.0073	2.0132	2.0046	2.0102

Note: The Table reports results from equation 3.10, estimated with a Newey-West heteroskedasticity and autocorrelation consistent standard errors. The last four rows report the adjusted R^2, the Log Likelihood (LL), the Schwarz criterion (BIC) and the Durbin-Watson statistic. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

Consistent with the results from the cross-correlogram reported in Panel B of Table 3.2, the contemporaneous coefficient γ_0 has a statistically significant positive in contrast to the lagged index returns that carry a negative sign. The contemporaneous coefficient of absolute returns is insignificant and negative in almost all subsamples suggesting that the size of market return does not affect the CRP.

3.5.4 Correlation and Volatility Risk Premia

A number of studies have provided evidence of positive relationship between volatility and correlation (ABDL 2001; ABDE 2001). The analysis is extended to study the relationship between the correlation and the volatility risk premium as well as between the CRP and the VIX, which is widely acknowledged as the “investor fear gauge”. The scatterplots in Figure 3.6 provide an initial estimate of the expected relationship between CRP, VIX and the Variance Risk Premium of the S&P 100 Index (VRP_m). Consistent with previous findings, visual inspection of scatterplot suggests that the CRP is negatively related to VIX and positively related to VRP of the index.

Figure 3.6: Scatterplot of changes in CRP against changes in VIX and the Variance Risk Premium (VRP) of the index options

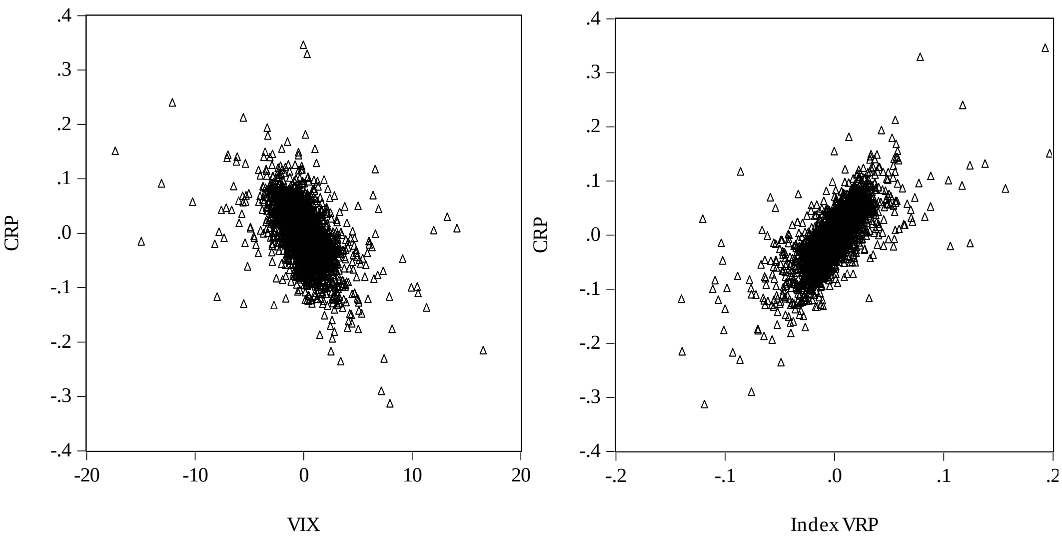

Note: Both variables are in first differences.

To formally test the relationship between the three variables, first, the equations 3.11 and 3.12 are employed to assess the impact of each variable on CRP separately, while, equation 3.13 estimates the joint information content of VIX and VRP on the CRP.

$$\Delta CRP_t = c + \sum_{k=1}^{3} \beta_k \Delta CRP_{t-k} + \sum_{k=0}^{3} \gamma_k \Delta VIX_{t-k} \tag{3.11}$$

$$\Delta CRP_t = c + \sum_{k=1}^{3} \beta_k \Delta CRP_{t-k} + \sum_{k=0}^{3} \gamma_k \Delta VRP_{m,t-k} \tag{3.12}$$

$$\Delta CRP_t = c + \sum_{k=1}^{3} \beta_k \Delta CRP_{t-k} + \sum_{k=0}^{3} \gamma_k \Delta VRP_{m,t-k} + \sum_{k=0}^{3} \delta_k \Delta VIX_{m,t-k} \tag{3.13}$$

Results from Table 3.6 suggest that concurrent values of CRP and VIX are negatively correlated while for greater lags of VIX, an increase of VIX result in increase in CRP, suggesting that the negative effect is absorbed within the certain timeframe. Notably, compared to the regression that includes VIX as an independent variable, the specification with the Volatility Risk Premium obtains significantly greater values of adjusted R^2 and log likelihood (LL), suggesting greater explanatory power of VRP for the CRP.

Table 3.6: Relationship of CRP and VIX changes

	1996 - 2000	2001 - 2005	2006 - 2010	1996 - 2010
c	0.000	0.000	0.000	0.000
β_1	-0.142***	-0.062**	-0.153***	-0.132***
β_2	-0.040	-0.033	-0.079**	-0.053**
β_3	0.024	-0.013	0.005	0.007
γ_0	-0.022***	-0.015***	-0.011***	-0.014***
γ_1	0.001	0.001	0.000	0.000
γ_2	-0.002	0.002**	-0.001	-0.001
γ_3	0.000	0.000	0.000	0.000
Adj. R^2	0.408	0.263	0.276	0.298
LL	2334.896	2612.866	2064.557	6862.289
BIC	-3.711	-4.115	-3.414	-3.700
DW	1.994	2.003	1.999	1.998

Note: The regression has been estimated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. The last four rows report values for the Adjusted R^2, the Log Likelihood (LL), the Schwartz criterion (BIC) and the Durbin-Watson (DW) statistic. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

Consistent with previous literature that suggests positive relationship of correlation and volatility, I find that the relationship also holds for the respective risk premia. Throughout the sample and for every sub-period, the correlation risk premium is significantly and positively correlated with the volatility risk premium while the obtained coefficients suggest strong relationship. This finding might be also attributed to the decomposition of the correlation risk premium into the variance risk premium of the index and the individual stocks, suggested by Driessen, Maenhout and Vilkov (2009).

Table 3.7: Relationship of CRP and Index VRP changes

	1996 - 2000	2001 - 2005	2006 - 2010	1996 - 2010
c	0.000	0.000	0.000	0.000
β_1	-0.115***	-0.048*	-0.108***	-0.095***
β_2	-0.104***	-0.099***	-0.112***	-0.109***
β_3	-0.032	-0.074**	0.006	-0.026
γ_0	2.118***	1.614***	1.400***	1.649***
γ_1	-0.003	-0.126*	-0.049	-0.054
γ_2	0.194**	0.115*	0.190**	0.196***
γ_3	0.017	0.066	0.009	0.025
Adj. R^2	0.654	0.499	0.501	0.542
LL	2687.892	2855.018	2286.348	7668.079
BIC	-4.245	-4.501	-3.785	-4.125
DW	1.998	2.002	1.998	1.999

Note: The Table presents the estimated coefficients from Equation 3.12. The regression has been estimated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

Finally, equation 3.13 allows testing the joint effect of VIX and VRP on CRP changes. By adding the VIX in the specification, current values of VIX are no longer significant whereas current and lagged values of VRP maintain their statistical significance. Notably, the adjusted R^2 and the value of log likelihood value are marginally increased when compared to equation 3.12 indicating that the addition of VIX does not provide any additional useful information.

Table 3.8: Relationship of CRP, VIX and Index VRP changes

	1996 - 2000	2001 - 2005	2006 - 2010	1996 - 2010
c	0.000	0.000	0.000	0.000
β_1	-0.120***	-0.053*	-0.109***	-0.100***
β_2	-0.103***	-0.094***	-0.112***	-0.111***
β_3	-0.006	-0.071**	0.006	-0.020
γ_0	1.800***	1.503***	1.493***	1.597***
γ_1	0.107	-0.075	-0.029	0.002
γ_2	0.229***	0.241***	0.235***	0.271***
γ_3	0.124	0.187***	0.088	0.113
δ_0	-0.007***	-0.002	0.001	-0.001
δ_1	0.002***	0.001	0.000	0.001
δ_2	0.000	0.003***	0.001	0.001***
δ_3	0.003***	0.002**	0.001	0.002***
Adj. R^2	0.680	0.508	0.502	0.545
LL	2718.539	2868.129	2289.736	7663.610
BIC	-4.305	-4.499	-3.767	-4.125
DW	2.009	2.014	1.996	2.003

Note: The Table presents the estimated coefficients from Equation 3.13. The regression has been estimated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

Finally, I apply a Granger causality test to address the question of intertemporal causality relationship between the variables under examination. According to the test, a significant Granger causality implies that the variable driving the changes contains information that could be useful in forecasting the "dependent" variable.

$$\begin{aligned}\Delta CRP_t &= \sum_{k=1}^{5}\beta_{1k}\Delta CRP_{t-k} + \sum_{k=1}^{5}\gamma_{1k}\Delta VRP_{m,t-k} + \varepsilon_{1t} \\ \Delta VRP_t &= \sum_{k=1}^{5}\beta_{2k}\Delta VRP_{t-k} + \sum_{k=1}^{5}\gamma_{2k}\Delta CRP_{m,t-k} + \varepsilon_{2t}\end{aligned} \tag{3.14}$$

Obtained results in Table 3.9 suggest an interesting change of causality direction over the subsamples. For the first two sub-periods, in essence from 1996 – 2005, the Volatility Risk Premium of the Index was driving the Correlation Risk Premium whereas the relationship reversed for the remaining years until 2010 as well as the whole sample period. The hypothesis of CRP driving the changes in the VRP of the index is in line with the proposition set forward by Driessen, Maenhout and Vilkov (2009) who explain that the index options are more expensive because they offer a valuable hedge against correlation.

Table 3.9: Granger causality test for changes in CRP and Index VRP

Coefficient	Null Hypothesis:	1996 - 2000	2001 - 2005	2006 - 2010	1996 - 2010
γ_1	VRP Index changes does not Granger Cause CRP changes	2.043*	2.293**	1.364	0.641
γ_2	CRP changes does not Granger Cause VRP Index changes	1.035	1.671	2.510**	5.506***

Note: For every sub-period and the whole sample period, the F-statistic is reported in the respective column. The lag length is selected using the Schwartz criterion and set to five. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

3.6 Economic Determinants and Time Variation of CRP

A number of studies have examined the predictive power of several economic variables for equity returns and second moment of distribution (e.g. Harvey and Whaley, 1992 for the S&P 100 market; Schwert, 1989, Glosten, Jagannathan and Runkle, 1993, Perez-Quiros and Timmermann, 2001 for the volatility; Bakshi and Madan, 2006, Konstantinidi and Skiadopoulos, 2013 and Feunou et al., 2014 for the volatility risk premium; Sheppard, 2008 and Palandri, 2009 for correlations). Based on these findings, the ability of factors related to general macroeconomic as well as stock market specific conditions to explain the time variation and predict the correlation risk premium is explored. The first set of factors, which capture the broad economic conditions, are the default spread

(*DEF*), defined as the difference of an Aaa bond yield from a Baa bond yield, the yield curve (*TERM*), calculated as the difference 10-year US government bond and the yield of the 3-month Treasury bill, and the three-month USD Libor (*Libor*). Additionally, the vector of factors that are closely related to the stock market conditions is the return of the S&P 100 Index (*R*) as well as the trading volume of the S&P 500 Index (*Vol*), considered as a proxy for the information flow in the financial markets. All variables are obtained from Datastream. Konstantinidi and Skiadopoulos (2013) propose the incorporation of the TED spread measured as the difference of the three-month Libor and the three-month Treasury bill. However, tests for multicollinearity of dependent values have suggested, as expected, high correlation of the TED spread with the yield curve slope and the Libor. In addition, the VIX and the Model Free Implied Volatility of the S&P 100 Index are found to exhibit multicollinearity with the dependent variables. I thus refrain from including these variables in the specification. Finally, with a view to out-of-sample forecasting, the lagged values of independent variables are used and the specification is augmented with lagged values of the CRP to account for autocorrelation.

$$CRP_t = c + CRP_{t-1} + R_{t-1} + TERM_{t-1} + Vol_{t-1} + DEF_{t-1} + Libor_{t-1} + \varepsilon_t \qquad (3.15)$$

The equation 3.15 is estimated with a stepwise regression, which allows only the variables with some explanatory power to the dependent variable to be included in the proposed specification. For every sub-period of the analysis, I keep the first two years for the in-sample analysis and produce one-step forecasts using a rolling window for the following three years for out-of-sample evolution. Respectively, for the whole sample period, the in-sample analysis is based on evidence from the period of 1996 – 2000 while the rest ten years are maintained for the out-of-sample evaluation.

Table 3.10: In-sample evidence of the Economic Determinants model

	1996 - 1997	2001 - 2002	2006 - 2007	1996 - 2000
C	-0.007***	0.000	0.001	-0.001
CRP_{t-1}	0.950***	0.946***	0.953***	-0.146***
R_{t-1}	2.058***	0.741***	3.606***	1.716***
$TERM_{t-1}$	9.368**	-1.815	3.698*	--
Vol_{t-1}	0.015*	0.007	--	0.009
Def_{t-1}	26.858	--	--	--
$Libor_{t-1}$	--	0.230**	--	--

Adj. R^2	0.894	0.892	0.875	0.212
LL	831.379	986.664	901.769	2173.032
BIC	-3.231	-3.872	-3.543	-3.440
DW	2.003	2.111	2.185	1.921

Note: The Table reports results from the in-sample estimation of equation 3.15 using a stepwise regression. "-" indicates that the variable does not augment the explanatory power of the model and thus, has not been added to the proposed specification. The last four rows report values for the Adjusted R^2, the Log Likelihood (LL), the Schwartz criterion (BIC) and the Durbin-Watson (DW) statistic. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

Notably, the only variable that has consistently significant explanatory power to the CRP is the S&P 100 index returns. The yield curve, *TERM*, is significant and positive for the first and the third sub-period. The term structure is related to short term business cycles, signaling information for the general economic activity (see Chordia and Shivakumar, 2002; Fernandes, 2014); small values of term structure (i.e. more flat term structure) is also considered a proxy for recession. The positive sign indicates that an increase in the term spread, coinciding with bull markets, will increase the CRP, leading to values that are more negative. Similarly, the increase of the interest rate, *LIBOR*, has a positive effect on the CRP. The default risk premium is not significant in any period under examination whereas the trading volume is only significant in the first period. Finally but of outstanding importance, for the last in-sample period of 1996 – 2000 none of the macroeconomic variables are significant.

Out-of-sample evaluation

In this Section, the forecasting performance of the economic determinants model, as presented in equation 3.15, is presented. For the out-of-sample forecasting, only the variables found significant at 10% significance level shall be used. Table 3.11 presents the root mean squared error (RMSE) and the mean absolute error (MAE). RMSE is calculated as the square root of the average squared deviations of the forecasted values from the actual series, while MAE is measured by the average of the absolute value of forecast errors. For comparison purposes, the respective statistical measures for the random walk model for the CRP are calculated. Notably, the obtained statistics do not vary substantially between the two models, while the forecasting error is always smaller in the case of random walk.

Additionally, for the purposes of directly testing the presence of market efficiency, the forecasting accuracy of the alternative model to forecasts obtained from a benchmark model, the random walk is compared. To this end, the Modified Diebold-Mariano test (see Harvey et al. 1997) is employed. The loss differential function is defined as $d_{jt} = [g(e_{jt}) - g(e_{0t})]$, where $g(e_{jt})$ is the loss function, for the economic determinants model, and $g(e_{0t})$ is the loss function for the benchmark model of random walk. The null hypothesis of equal forecasting accuracy is tested against the alternative that the forecasting model performs better than the benchmark model, i.e. $E(d_{jt}) < 0$. The loss function is computed both in terms of the mean square error of the forecast and of the mean absolute error. The Diebold-Mariano test statistic is defined as:

$$DM = \frac{\bar{d}}{\sqrt{\frac{2\pi \hat{f}_d(0)}{T}}} \sim N(0,1) \qquad (3.16)$$

where $\bar{d}$ is the sample average of the loss differential, $\hat{f}_d(0)$ is the estimate of the spectral density at frequency zero and T is the number of observations. For h-step-ahead forecasts, the Modified Diebold-Mariano test statistic corrects for small sample sizes and autocorrelation of the loss differential following a Student-t distribution with T-1 degrees of freedom and equals to

$$DM_{\text{mod}} = \left[\frac{T + 1 - 2h + h(h-1)/T}{T} \right] DM \qquad (3.17)$$

Implementation of the DM test to the sample suggests that, both under the RMSE and the MAE metrics, the null hypothesis is strongly accepted for the three sub-periods, proposing that the random walk model produce smaller forecasting error compared to the economic determinants model.

Table 3.11: Out-of-sample performance

	1996 - 2000		2001 - 2005		2006 - 2010		1996 - 2010	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
Economic Variables	0.1144	0.0901	0.0901	0.0716	0.1298	0.1060	0.1276	0.0978
Random Walk	0.1139	0.0899	0.0876	0.0698	0.1292	0.1029	0.0684	0.0532
DM test statistic	2.1724	1.0338	3.2321	2.2475	-0.6167	0.9018	21.1645	25.5733

3.7 Conclusions

Extending the methodology proposed by Britten-Jones and Neuberger (2000) on model-free implied volatility and correcting for the early exercise premium inherent in American option prices I propose a new correlation index, as a modification to the misspecified S&P 500 ICI published by the CBOE. The new measure is inferred by prices of currently traded options, thus providing a model-free estimate of market-wide correlation and diversification benefit. An extensive dataset allows assessing the impact of periods of financial turbulence associated with lower asset returns and increased volatility. First, this study contributes to the existing literature of volatility and correlation risk premium. Additional evidence of a negatively priced volatility risk premium, both for the index and the constituent individual stocks is provided for the entire sample period. In addition, model-free implied correlation is consistently higher than realized correlation resulting in a negative price of correlation risk premium throughout the sample period. Nevertheless, the correlation risk premium is no longer statistically significant during turbulent periods, and the correlation under the risk neutral and the physical measure do not differ substantially. A plausible explanation for the decreased negative value of correlation risk premium is that the market dynamics during the crisis period eliminated any previous mispricing of index options, which was mainly attributed either to investors' irrationality or to lack of arbitrage opportunities due to increased transaction costs and margin requirements.

Understanding the dynamics of correlation risk premium is of paramount importance in asset pricing theory and other financial applications. The study of the time series dynamics during alternative sample periods suggest that the distribution is far from normal, leptokurtic and negatively skewed. Over the entire sample period, the CRP exhibits strong positive persistence and long memory traits along with strong intraweek pattern. Furthermore, the asymmetric response to positive and negative returns of the S&P 100 is examined. Results suggest that CRP and index returns are positively correlated, suggesting that an increase in returns will induce CRP to more negative values. Finally, as expected, the volatility and correlation risk premia are positively correlated and interestingly the direction of causality changes during the last five years of the sample

Finally, the informational content of several market-specific and macroeconomic variables on future changes of the CRP is assessed. To this end, variables that capture the wide economic state, namely the term structure, the default spread structure and the USD Libor, along with market-specific variables, namely the S&P 100 index returns and the trading volume of the S&P 500 index are employed. The in-sample evidence suggest that the index return are consistently significant throughout the alternative sample periods while the term structure obtains a significant value only for the periods of 1996 – 2000 and 2006 – 2010. The out-of-sample evaluation suggests that the proposed model fails to produce superior forecasting performance when compared to the benchmark model of random walk.

Chapter 4

On the predictability of model-free implied correlation

Chapter abstract. In this Chapter, the existence of predictable patterns in the dynamic evolution of the model-free implied correlation is assessed. To this end, alternative time-series specifications are employed to capture different aspects of series distribution. The out-of-sample significance of obtained forecasts is examined through statistical and economic criteria. Combination forecasts obtain the minimum forecast error and the maximum efficiency in terms of predicting accurately the direction of change of the actual series. The statistical measures provide strong evidence in favour of existing predictable pattern in the S&P 100 option market. The economic significance of the out-of-sample forecasts is assessed based on their ability to yield abnormal profits. To this end, I employ an innovative trading strategy that exploits daily changes of the MFIC series. Obtained results suggest that the existence of predictable patterns in the evolution of the series, supported by statistical measures, can be further exploited to attain trading profits. The proposed trading strategy yields significant economic gains, with the AR(I)MA-GARCH model and the combination forecasts of implied correlation generating significant profits. However, when transaction costs are considered, profitability is eliminated suggesting that the efficient market hypothesis cannot be rejected. The results remain robust across different sample windows and forecast periods.

4.1 Introduction

Existing literature has focused its attention mainly on the ability of risk-neutral moments, extracted from option prices, to provide accurate forecasts of the corresponding realized moments or to improve asset allocation and pricing decisions[3]. An equally important yet distinct question is whether risk-neutral moments are predictable per se. The existence of predictable patterns in the dynamics of risk-neutral moments raises questions on the efficiency of the option markets and provides opportunities for profitable trading strategies. Thus far, only a limited number of studies have dealt with the predictability of option-implied measures, and mainly volatility. This Chapter fills this gap by undertaking a comprehensive study of the dynamics and unexplored predictable patterns in the evolution of the implied correlation series per se.

Understanding the dynamics of risk-neutral correlation is important for both academics and practitioners. In this study, implied correlation is derived from observed prices of index and individual equity options without assuming any implicit option-pricing model. Inherently, implied correlation is a forward-looking measure of the aggregate stock market diversification while, according to Skintzi and Refenes (2005), it also quantifies the difference of portfolio variance from its minimum and maximum values. Apparently, the risk neutral measure of correlation is of paramount importance in an asset-pricing and asset allocation framework. Cosemans (2011) and Driessen, Maenhout and Vilkov (2012) provide evidence that risk-neutral correlation measures can explain how expected returns change over time. From a more practical perspective, forecasts of risk-neutral correlation can be used by market participants to form profitable trading strategies. Volatility and correlation trading strategies have stimulated the interest of investors, especially after the dramatic increase in stock market volatilities and correlations, following the 2008 financial crisis. Interestingly, in July 2009, the CBOE

[3] Several studies have used implied correlation measures to forecast information on expected correlations for various asset classes (e.g. Campa & Chang, 1998, and Lopez & Walter, 2000, for exchange rates, Han, 2007, for interest rates, Skintzi & Refenes, 2005, for equities). Moreover, DeMiguel et al. (2012) examine whether option-implied correlations can improve asset allocation decisions, while Chang et al. (2012) and Buss & Vilkov (2012) use option-implied information to derive equity betas.

launched the S&P 500 implied correlation index to measure the expected average correlation of price returns of index components implied through S&P 500 index option prices and prices of options on the 50 largest components of S&P 500.[4]

In the context of volatility, David & Veronesi (2002) and Guidolin & Timmerman (2003) provide a theoretical explanation why implied volatility may by predictable by linking option prices and implied volatility to economic uncertainty. In the context of correlation, Buraschi, Trojani and Vedolin (2013) develop a structural equilibrium model that links the differential pricing of index and individual equity options to aggregate economic uncertainty and diversity in beliefs across investors.

A number of studies have investigated whether predictable dynamics exist in the dynamics of implied volatility. Harvey & Whaley (1992), Guo (2000), Brooks & Oozeer (2002) and Konstantinidi et al. (2008) provide evidence of significant predictability in implied volatility using time-series models and economic determinants as predictors. However, trading strategies designed to exploit predictable patterns do not achieve significant economic profits. In contrast, Goyal & Saretto (2009) provide evidence of statistically and economically significant predictability in the dynamics of volatility implied in at-the-money stock option prices. Goncalves & Guidolin (2006) and Bernales & Guidolin (2010) address the question of whether the implied volatility surface contains any exploitable patterns. They both find evidence of statistically significant predictable patterns that, however, do not yield significant trading profits when transaction costs are taken into account. Neumann & Skiadopoulos (2013) exploit predictable patterns in the dynamics of higher-order moments. To the best of my knowledge, the only related study that also explores the dynamics of risk-neutral correlation is conducted by Härdle & Silyakova (2012). However, Härdle & Silyakova use only one model, a dynamic semi-parametric factor model, to capture the dynamics of the implied correlation surface and forecast future implied correlation for the German market.

[4] The CBOE S&P500 Implied Correlation Index was not used in this study, as it is a constant maturity index, based on specific option pricing model. Moreover, it suffers from liquidity and microstructure issues since no filtering rules are applied on the options involved in its calculation. Similar studies resort also to the construction of the implied correlation index from scratch (e.g. Driessen et al, 2012).

This study makes a number of important contributions on the predictability of risk-neutral correlation. Firstly, I examine whether the dynamics of the implied correlation series, as a measure of market-wide correlation, contain any predictable pattern. To this end, I fit alternative time-series specifications to model and forecast the dynamics of the series. Out-of-sample forecasts are obtained and their forecasting performance is assessed based on a variety of statistical evaluation criteria. Secondly, I investigate whether the predictability of the implied correlation series can be exploited in the context of a correlation trading strategy. Based on the notion of dispersion trade, implemented by practitioners to trade correlation risk, I build a trading strategy exposed to correlation risk that trades inverse positions in index options and stock component options. Thirdly, an extensive dataset consisting of the S&P 100 index options and of the individual options of the S&P 100 constituent stocks for an extended period is used. The sample period extends from January 1996 to October 2010, thus encompassing several turbulent periods associated with high volatility and low returns, such as the 9/11 subsequent increased market volatility conditions as well as the U.S. subprime mortgage crisis that resulted in the on-going global financial turmoil.

The risk-neutral implied correlation measure examined in this study is based on the notion of 'equicorrelation', i.e. constant correlations for each pair of assets. Although this assumption may appear too restrictive, it has a number of important advantages. In an early study of asset allocation, Elton & Gruber (1973) found that the assumption of equicorrelation reduces estimation error and provides superior portfolio selection. More recently, Pollet & Wilson (2010) provide a theoretical explanation and empirical evidence that average realized stock market correlation predicts future stock market returns. Furthermore, Engle & Kelly (2012) show that multivariate models based on the equicorrelation assumption improve portfolio allocation compared to unrestricted models. In the context of implied correlation, Skintzi & Refenes (2005) interpret the implied correlation index as the market view of future stock market diversification and provide evidence that implied correlation measures outperform historical ones in predicting future correlation. Moreover, Driessen, Maenhout and Vilkov (2012) find that average implied correlation has significant predictive power for future stock market returns.

The empirical results of this study, based on alternative time-series specifications, suggest the existence of a strong predictability pattern in the implied correlation structure. Turning the attention to the economic significance of the resulting forecasting values, I find that the implemented trading strategy can yield abnormal profits. The existence of profitable trading patterns raises doubt on the efficiency of the index and stock option market. However, after accounting for transaction costs, the abnormal profits fade out. The reported results are robust across different in-sample sizes and forecast periods.

The remainder of this Chapter is organized as follows: Section 2 describes the Model-Free Implied Correlation (MFIC), the dataset used for the calculation of MFIC series along with the descriptive statistics of the series under examination. Section 3 presents the alternative model specifications used for forecasting purposes. Finally, Sections 4 and 5 discuss the in sample and the out-of-sample evidence from the models under consideration and the trading strategy, respectively. The robustness of the results to different sample sizes and forecasting horizons and the findings are presented in Section 6. Section 7 concludes.

4.2 Methodology and data

Intuitively, index variance is associated with the variance of the constituent stocks as well as the pairwise correlations. Thus, the variance of an index, consisting of N stocks, is defined as follows:

$$\sigma_{\mathrm{I},t}^2 = \sum_{i=1}^{\mathrm{N}} w_{i,t}^2 \sigma_{i,t}^2 + 2\sum_{i=1}^{\mathrm{N}} \sum_{j>i} w_{i,t} w_{j,t} \rho_{ij,t} \sigma_{i,t} \sigma_{j,t} \tag{4.1}$$

where $\sigma_{I,t}^2$ is the index variance, $\sigma_{i,t}$ is the standard deviation of the asset i, i = 1,...,N, $\rho_{ij,t}$ is the correlation between assets i and j, and $w_{i,t}$ is the relative weight of each index component i, all variables at time t.

Assuming stable correlation across assets, the average correlation measure $\bar{\rho}_t$ is defined as a factor that captures any arising difference between the variance of the index and the variance of a portfolio, formed as a weighted average of the constituent stocks of the index (see Skintzi & Refenes, 2005).

$$\bar{\rho}_t = \frac{\sigma_{I,t}^2 - \sum_{i=1}^{N} w_{i,t}^2 \sigma_{i,t}^2}{2\sum_{i=1}^{N-1}\sum_{j\neq i} w_{i,t} w_{j,t} \sigma_{i,t} \sigma_{j,t}} \tag{4.2}$$

As Engle & Kelly (2012) pointed out, since equal pair-wise correlations are assumed, the average correlation measure will lie in the [-1/(n-1), 1] interval. Historical volatility and correlation estimators rely on the assumption that the future dynamics of the series will be similar to the past. To overcome the ambiguities deriving from the above assumption, stock return moments are inferred from currently traded option prices.

The Model-Free Implied Volatility (MFIV) measure proposed in the recent literature (e.g. Jiang and Tian, 2005) is employed to compute the Model-Free Implied Correlation from equation 4.2. The implementation procedure is thoroughly described in Section 3.2.

For the purposes of this study, daily closing quotes for options on the S&P 100 index and the constituent stocks are obtained from OptionMetrics. The sample extends from January 1996 to October 2010, thus including 179 stocks.[5] On each day, the actual daily weight of stock *i* is computed based on its daily market capitalization divided by the total market capitalization of the 100 stocks included in the index, at the specific day. The subset of 04/01/1996 to 31/12/2000 shall be used for the in-sample estimation, whereas the remaining period will be retained for the out-of-sample evaluation of forecasting methods. As in Chapter 3, the calculation of model-free implied volatility is based on observed prices from currently traded options. Data on option prices, underlying asset prices, and company distributions are obtained from OptionMetrics database. The zero-coupon interest rate curve, provided by OptionMetrics database, is used as the risk-free interest rate. Several filtering rules are described in Section 3.3.

[5] The analysis does not suffer from survivorship bias since it does not include only options written on stocks that have been continuously traded throughout the sample period. Instead, the MFIC index, replicates the composition of the S&P100 index on a daily basis. The main criteria for a company to be included in the S&P 100 index are market capitalization, liquidity and the availability of individual stock options. As a result, heavily traded and liquid options introduced during the sample periods are expected to be included in the index.

Figure 1 shows the daily evolution of the Model-Free Implied Correlation (MFIC) from January 1996 to October 2010. Consistent with the well-documented fact of increased correlation during bear markets, the series present several spikes throughout the whole sample period coinciding with periods of low returns and high volatility. Specifically, resulting from the Asian and the Russian financial crises, the MFIC series reached the value of 0.9, in late 1997 and later in August 1998. During the first years of the ‘00s, the low returns market stemming from the 9/11 terrorist attack, the South-American financial turmoil and the internet bubble burst, resulted in increased levels of correlation, though lower than the aforementioned spikes of the ‘90s. Over the subsequent years, MFIC trended lower, with the exception of the period of a multi-market sell-off in May 2006. Finally, as a result of the global financial crunch, which commenced in 2007 and culminated in September 2008 with the announcement of the bankruptcy of Lehman Brothers, the MFIC presents several peaks, while reaching the highest value of 0.97 on October 24, 2008.

Figure 4.1: Model-Free Implied Correlation (MFIC), Model-Free Implied Volatility (MFIV) series for the Index and the individual stocks.

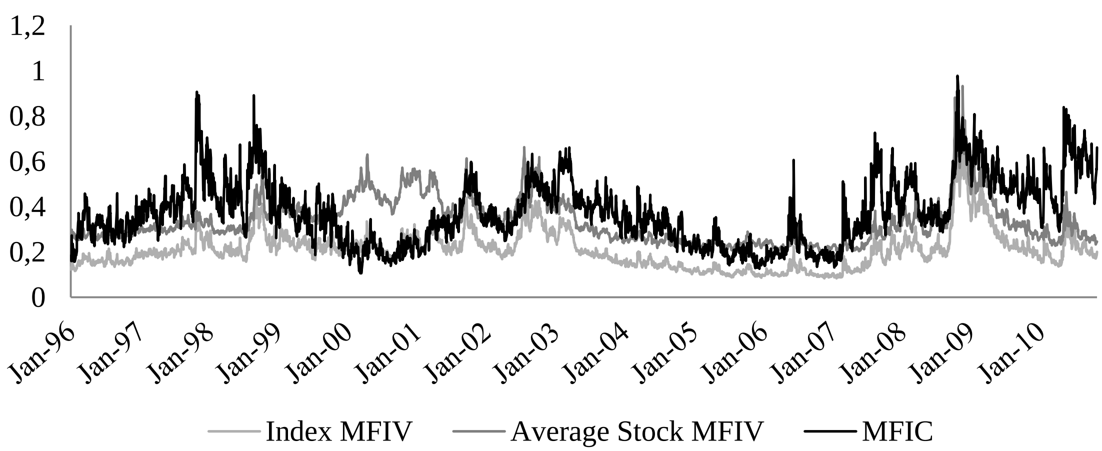

Table 4.1 reports the summary statistics of the model-free implied correlation series for three sample periods. The first period of 1996 – 2000 corresponds to the later used in-sample estimation period; the second is the out-of-sample period January 2001 – December 2010 while the last includes the whole sample period. The higher moments and the Jarque-Bera test reject the hypothesis of normality for the MFIC series for all the periods under consideration. Table 4.1 also reports the test statistics of the Augmented

Dickey-Fuller (ADF) test, the Phillips-Perron (PP) test, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test statistic. According to the PP test, the null hypothesis of unit root is strongly rejected for all three subsets at a significance level of 1%. Comparably, the ADF test suggests that the existence of unit root shall be rejected throughout the whole sample period. In contrast, the KPSS test provides evidence of non-stationary time series for all three periods under consideration. These mixed results, along with the slow decaying and significant autocorrelation coefficients, suggest the existence of long memory properties in the MFIC similar to model-free implied volatility (see Brooks & Oozeer, 2002). These features are further discussed in the following Section.

Table 4.1: Summary Statistics

	1996 - 2000		2001 – 2010		1996 - 2010	
	MFIC	Logit MFIC	MFIC	Logit MFIC	MFIC	Logit MFIC
Mean	0.360	-0.617	0.379	-0.537	0.373	-0.564
Median	0.349	-0.622	0.359	-0.578	0.356	-0.594
Max	0.907	2.281	0.977	3.756	0.977	3.756
Min	0.106	-2.132	0.126	-1.934	0.106	-2.132
Std. Dev.	0.135	0.622	0.147	0.672	0.143	0.656
Skewness	0.808	0.609	0.492	0.396	0.594	0.468
Kurtosis	3.868	4.192	2.746	3.509	3.043	3.698
JB	176.12***	152.062***	106.372***	91.298***	219.376***	211.625***
$\rho 1$	0.949***	0.944***	0.964***	0.961***	0.96***	0.956***
$\rho 10$	0.778***	0.777***	0.857***	0.853***	0.836***	0.832***
$\rho 15$	0.701***	0.699***	0.824***	0.823***	0.792***	0.789***
ADF	-3.167**	-3.129**	-3.07**	-3.077**	-4.41***	-4.06***
PP	-4.964***	-5.129***	-5.389***	-5.571***	-7.477***	-7.88***
KPSS	1.192***	1.277***	1.477***	1.408***	0.767***	0.72**

Note: The Table reports summary statistics for the Model-Free Implied Correlation (MFIC) and the inverse logistic transformation (transformed MFIC). In addition, the autocorrelation coefficient (ρ) for the 1st, 10th and 15th lag of the autocorrelation structure as well as the Augmented Dickey-Fuller (ADF), Phillips-Perron (PP) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test statistics are reported. The number of lags for the ADF test is selected according to the modified Schwarz criterion. The bandwidth for the PP and KPSS tests is automatically selected using the Newey-West lag selection parameter. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

4.3 Forecasting models

As previously documented, by definition, MFIC is bounded within the range of [-1/(n-1), 1]. With the main goal of this study lying within the area of predictability of the MFIC series, and to alleviate arising issues in the forecasting procedure, the inverse of the logit transformation of the original correlation series is employed and defined as

$$\rho_t^* = -\frac{1}{n-1} + (1+\frac{1}{n-1})\frac{1}{1+\exp\left(-\rho_t\right)} \tag{4.3}$$

where ρ_t^* is the MFIC series, n is the sample size and ρ_t is the logit transformation of the series. The logit transformation ensures that the series will lie at the interval [-1/(n-1),1].

Table 4.1 also presents the summary statistics for the logit transformation of the series. Unit root (ADF and PP) and stationarity tests (KPSS) provide mixed results regarding the stationarity of the logit transformation, similar to those obtained for the original MFIC series. To further investigate the stationarity and long memory properties of the series, an ARFIMA (p,d,q) model to the logit transformation is estimated. A statistically significant estimate of d equal to 0.5 is obtained, thus supporting the non-stationarity of the series. Thus, the econometric analysis that follows will be conducted on the differentiated logit transformation of the MFIC series. In the rest of the Chapter, for notation reasons, I will refer to the first difference of the logit transformation as the MFIC series.

Stock returns and implied volatility have been extensively examined for the presence of seasonality effects. Monday (Friday) tends to be a day in which traders open (close) positions for the week and excess buying (selling) pressure may result in higher (lower) implied volatility (see Harvey & Whaley, 1992). To this end, the presence of day-of-the-week effects is examined through the following regression.

$$MFIC_t = \sum_{i=1}^{5} \gamma_i D_{i,t} + \delta MFIC_{t-1} + u_t \tag{4.4}$$

where i takes values from 1 to 5 for Monday to Friday, respectively.

Additionally, the presence of the January effect on the MFIC is tested by the following specification:

$$MFIC_t = \sum_{i=1}^{12} \gamma_i D_{i,t} + \delta MFIC_{t-1} + u_t \tag{4.5}$$

where *i* takes values from 1 to 12 for January to December, respectively. A lagged term of the dependent variable is included in both equations (4) and (5) to eliminate the effect of autocorrelated errors.

Following, the modelling of the evolution and the time series properties of the MFIC under different forecasting models is described.

AR(I)MA, AR(I)MA - GARCH and ARFIMA models

First, an autoregressive moving average process (AR(I)MA) is fitted to the evolution of the series by adding lags of the error term as well as lags of the series under examination. The AR(I)MA(r,d,m) model can be specified by:

$$\Phi(L)\Delta^d (X_t - \mu) = \Theta(L)u_t \tag{4.6}$$

where Φ and Θ are polynomials of order r and m respectively, such that $\Phi(L) = 1 - \varphi_1 L - \varphi_2 L^2 - \ldots - \varphi_r L^r$, and $\Theta(L) = 1 - \theta_1 L - \theta_2 L^2 - \ldots - \theta_m L^m$, X_t is the MFIC, μ is the expectation of X_t, ε_t is the white noise error term, Δ is the difference operator and *d* is equal to one, representing the order of integration.

Furthermore, an AR(I)MA(r,d,m) – GARCH(p,q) model is employed to account for the remained heteroskedasticity in the error terms,. The error terms are assumed to follow a normal distribution with zero mean and variance equal to h_t^2, where:

$$h_t^2 = b_0 + \sum_{q=1}^{q} b_q u_{t-q}^2 + \sum_{p=1}^{p} c_p h_{t-p}^2 \tag{4.7}$$

Taking the AR(I)MA model one step forward by allowing *d* to take fractional values, the specification extends to an ARFIMA (r,d,m) model, given by:

$$\Phi(L)(1-L)^d (X_t - \mu) = \Theta(L)u_t \tag{4.8}$$

where $(1-L)^d$ is the fractional integration operator. The ARFIMA model captures both the short-run component of the series, through the autoregressive and the moving average terms, and the long-run dependence through the fractional differencing parameter *d*. The process is stationary under the condition that $-0.5 < d < 0.5$. The series exhibits long memory if $0 < d < 0.5$, suggesting positive dependence between distant observations, while in case of $-0.5 < d < 0$, the series presents negative dependence between distant observations, known as anti-persistence. Finally, for $d= 0$ (the general AR(I)MA process) the process exhibits short memory.

The AR(I)MA, AR(I)MA-GARCH and ARFIMA models have been estimated for all plausible combinations of autoregressive (AR) and moving average (MA) order terms both in the mean and the variance specification, where applicable, up to the fifth lag. The model selection is based on the Schwarz information criteria, which in contrast to Akaike information criteria, includes an extra term that penalizes data over fitting.

Regime Switching Model

In order to capture potential asymmetries in the correlation process a dynamic Regime Switching (RS) model is employed. More specifically, the transition between the regimes is governed by a Markov chain, two regimes are assumed and the coefficient on the lagged dependent variable is allowed to be regime-varying, i.e.

$$\sum_{r=1}^{R} \Phi_{s_t}^{r}\left(L\right)\Delta^{d}\left(X_t - \mu_{s_t}^{r}\right) = u_t \quad (4.9)$$

The transitions between the regimes $s_t = 1$ and $s_t = 2$ are given by a Markov chain with transition probabilities $p_{ij}= P(s_t=j|s_{t-1} = i)$ for i, j = 1,2 and $\sum_{j=1}^{2} p_{ij} = 1$ for i=1,2. The selection of the lag order is based on the Schwarz information criterion.

Heterogeneous Autoregressive (HAR) Model

The theory of heterogeneous market hypothesis is based on the empirical finding that volatility dynamics are affected by the different investment horizons of traders. Specifically, traders with short-term investment horizon rapidly incorporate any arriving information in their strategy, thus affecting directly the shorter-term volatility. In contrast, longer-term investors rebalance their position less frequently, disregarding daily information flow, thus empowering the long memory characteristic of volatility. Corsi

(2009) identified and modelled this asymmetry through an autoregressive process by aggregating daily, weekly and monthly volatilities.

The heterogeneous autoregressive (HAR) model is employed to capture the long-memory property of the model-free implied correlation. The model is given by:

$$MFIC_t^{(d)} = \alpha_0 + \alpha_{(d)} MFIC_{t-1}^{(d)} + \alpha_{(w)} MFIC_{t-1}^{(w)} + \alpha_{(m)} MFIC_{t-1}^{(m)} + u_t \qquad (4.10)$$

where $MFIC_{t-1}^{(w)} = \frac{1}{5}(\sum_{h=1}^{5} MFIC_{t-h}^{(d)})$ is the lagged weekly MFIC, and $MFIC_{t-1}^{(m)} = \frac{1}{22}(\sum_{h=1}^{22} MFIC_{t-h}^{(d)})$ is the lagged monthly MFIC.

Economic determinants

A number of studies have examined the role of economic variables in determining asset return correlations (see Erb et al., 1994, Moskowitz, 2003, Sheppard, 2008, Palandri, 2009, amongst others). Based on their findings, financial and macroeconomic variables, such as the short-term interest rates or the slope of the term structure, are expected to determine the systematic risk of equity portfolios and, consequently affect the time variation of equity correlations. Data on economic variables are obtained from Datastream.

Following Harvey & Whaley (1992) and Sheppard (2008), three interest rate variables, which have been widely used in predictability studies and are documented to influence the volatility and correlation process, are included in the specification; namely, the one-month USD LIBOR interest rate (INT_t), the slope of the yield curve (or else the term structure, $TERM_t$) defined as the difference between the yield of the 10-year US government bond and the yield of the 3-month Treasury bill as well as the slope of the junk bond spread ($JUNK_t$) defined as the difference of an Aaa bond yield from a Baa bond yield. The short-term interest rate is considered to accurately proxy the shocks to expected real economic activity. The junk bond spread and the term structure have been previously documented to capture the long-term and the short-term business cycle conditions, respectively. The lagged return and the trading volume (VOL_t) of the underlying security are included as control variables for leverage and information flow effects (see Bollen & Whaley, 2004). Two dummy variables are ealso included to assess the asymmetric response of the MFIC to positive (R_t^+) and negative index returns (R_t^-). Increased trading volume signals arrival of information to investors, thus inducing fluctuations on both returns and implied volatility. In addition, the Brent Crude Oil price (WTI_t) has been considered as a proxy for the fluctuations of an alternative asset class market. Furthermore, dividend yield of the S&P 100 index (DIV_t), as a proxy of time varying expected returns, and the EURO/USD exchange rate (FX_t) are included as explanatory variables. Each of the aforementioned variables is first differenced to represent innovations, and ensure stationarity.

Additionally, I have tested the economic variables model for a number of inference issues that also arise when regressing returns and volatility on macroeconomic variables (see Paye, 2008). Firstly, the MFIV changes of the index are not included to avoid multicollinearity issues based on the correlation matrix of the dependent variables. Secondly, the model has been tested for endogeneity due to reverse causality between the dependent and independent variables using Granger causality tests. The only dependent variable that exhibits reverse causality with the dependent variable is the trading volume. However, no significant contemporaneous correlation between MFIC and trading volume is found. With a view to minimizing bias of the forecasted values, the lagged values of the macroeconomic variables are used. Finally, the specification is supplemented with the inclusion of lagged values of model-free implied correlation so as to account for autocorrelation patterns. The model has been estimated with the inclusion of up to three lags of the MFIC series. Finally, the model that minimizes the Schwarz information criterion is selected.

$$\begin{aligned} MFIC_t &= b_0 + b_1^+ R_{t-1}^+ + b_2^- R_{t-1}^- + b_3 FX_{t-1} + b_4 INT_{t-1} + b_5 DIV_{t-1} \\ &\quad + b_6 TERM_{t-1} + b_7 JUNK_{t-1} + b_8 WTI_{t-1} + b_9 VOL_{t-1} + u_t \end{aligned} \tag{4.11}$$

Table 4.2 presents summary statistics for the economic determinants. While stationarity is rejected when variables are measured in levels (Panel A), unit root tests on the first difference indicate that all series are stationary (Panel B). The first order autocorrelation coefficient is significant for all variables measured in levels, whereas in most of the cases, no evidence of significant autocorrelation is found when variables are measured in first differences.

Table 4.2: Descriptive statistics for the economic determinants

Panel A: Raw series

	Pi	DIV	Volume (*VOL*)	FX	JUNK	Libor (*r*)	TERM	WTI
Mean	544.85	0.015	864,752,943.52	1.103	0.007	0.057	0.009	21.361
Median	531.97	0.014	860,160,000.00	1.103	0.007	0.056	0.009	20.68
Maximum	832.65	0.022	2,405,100,000.00	1.312	0.011	0.068	0.019	37.22
Minimum	285.77	0.01	136,580,000.00	0.825	0.005	0.049	-0.007	10.82
Std. Dev.	166.758	0.004	208,760,374.44	0.118	0.001	0.005	0.006	5.885
Skewness	0.088	0.33	0.38	-0.258	0.956	0.82	-0.464	0.42
Kurtosis	1.636	1.731	5.313	2.437	3.437	3.111	2.592	2.554
Jarque-Bera	99.091 ***	107.097 ***	310.45 ***	30.532 ***	201.575 ***	141.573 ***	53.77 ***	47.419 ***
$\rho 1$	0.998 ***	0.997 ***	0.727 ***	0.997 ***	0.993 ***	0.994 ***	0.991 ***	0.995 ***
ADF test	-1.272	-1.487	-3.087	-1.16	-1.879	-1.108	-0.713	-1.644

Panel B: First difference of the series

	P_I	DIV	Volume (*VOL*)	FX	JUNK	Libor (*r*)	TERM	WTI
Mean	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Median	0.001	0.000	0.001	0.000	0.000	0.000	0.000	0.001
Maximum	0.056	0.001	2.031	0.042	0.001	0.144	0.005	0.145
Minimum	-0.075	-0.001	-1.917	-0.023	-0.001	-0.108	-0.004	-0.150
Std. Dev.	0.012	0.000	0.221	0.006	0.000	0.008	0.001	0.026
Skewness	-0.331	-0.020	0.290	0.615	-0.750	2.713	0.859	-0.148
Kurtosis	6.299	20.352	27.379	5.781	11.537	183.555	12.764	7.172
Jarque-Bera	592.631 ***	15757.618 ***	31121.145 ***	483.816 ***	3931.412 ***	1707612.708 ***	5143.726 ***	915.461 ***
$\rho 1$	-0.018	-0.002	-0.360 ***	-0.039	-0.056 **	0.027	0.021 ***	-0.002
ADF test	-35.973 ***	-35.479 ***	-49.516 ***	-6.664 ***	-6.523 ***	-7.108 ***	-34.683	-35.237 ***

Note: P_I, DIV and VOL refer to the closing price, the dividend yield and the trading volume of the S&P 100 Index, correspondingly. FX is the EUR/USD exchange rate, JUNK is the slope of the junk bond spread defined as the difference of an Aaa bond yield from a Baa bond yield, Libor (r) is the one-month USD LIBOR interest rate, TERM is the slope of the yield curve (or else the term structure) defined as the yield of the 10-year US government bond minus the yield of the 3-month Treasury bill, and WTI is the price of the Brent Crude Oil. In addition, the autocorrelation coefficient (ρ) for the 1st lag of the autocorrelation structure as well as the Augmented Dickey-Fuller (ADF) test statistics are reported. The number of lags for the ADF test is selected according to the modified Schwarz criterion. The null hypothesis of the ADF test and the Jarque-Bera test is the presence of the unit and normal distribution, respectively. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively.

Combination forecasts

Each different forecasting model captures different dynamics of the underlying asset properties, based on alternative information sets. Empirical results from previous studies (see Becker & Clements, 2008) suggest that linear combinations of individual forecasts may produce superior forecasts compared to the individuals. In the present Section, a comparative analysis of the various methods of combining forecasts proposed to the literature is proposed.

Denoting by $f_t^{\,j}$, (where j=1 for AR(I)MA, 2 for AR(I)MA – GARCH, 3 for ARFIMA, 4 for Regime Switching model, 5 for the HAR and 6 for the economic determinants model), six different forecasted values of MFIC at time t are obtained, based on the time-series models described above. For notation reasons, superscripts of forecasting models are kept the same throughout the methods described below.

The simplest method is to combine forecasts by assigning equal weight to each individual forecast. Thus, the equally weighted combination forecast is a simple average of the individual forecasts, i.e.

$$f_t^{EW} = \frac{1}{6}\sum_{j=1}^{6} f_t^{\,j} \tag{4.12}$$

In addition, existing literature proposes the usage of the Schwarz model selection criterion as an alternative approach for obtaining combination weights. The methodology suggests that each model should be weighted according to the difference of its Schwarz criterion value with the "best" model, which attains the minimum value of the criterion. Following the methodology of Kolassa (2011), the Schwarz weighted combination forecast is obtained by the following equation

$$f_t^{BIC} = \sum_{j=1}^{6} w_{BIC}(j)\cdot f_t^{\,j} \tag{4.13}$$

where $w_{BIC}(j) = \dfrac{\exp(-\frac{1}{2}\Delta_{BIC}(j))}{\sum_{j=1}^{6}\exp(-\frac{1}{2}\Delta_{BIC}(j))}$, $\Delta_{BIC}(j) = BIC(j) - BIC(k)$, BIC (j) is the Schwarz criterion of the j^{th} model, j =1,2,…,6, and BIC (k) is the minimum Schwarz criterion associated with model *k*.

Above mentioned forecasts are calculated with constant weights throughout the out-of-sample period, thus failing to capture the dynamics of the series under consideration. To overcome this restriction, I also consider the derivation of combined forecasts, where the weights are time varying. In essence, the weights are obtained by minimizing the mean squared error of the following regression:

$$MFIC_t = \beta_0 + \beta_1 f^1_{t|t-1} + \beta_2 f^2_{t|t-1} + \beta_3 f^3_{t|t-1} + \beta_4 f^4_{t|t-1} + \beta_5 f^5_{t|t-1} + \beta_6 f^6_{t|t-1} \qquad (4.14)$$

where $MFIC_t$ is the actual value of the model-free implied correlation series and $f^j_{t|t-1}$ are the forecasted values of individual models at time t, calculated at time t-1. Consequently, the combination forecast value is given by:

$$f^W_{t+1|t} = \hat{\beta}_0 + \hat{\beta}_1 f^1_{t+1|t} + \hat{\beta}_2 f^2_{t+1|t} + \hat{\beta}_3 f^3_{t+1|t} + \hat{\beta}_4 f^4_{t+1|t} + \hat{\beta}_5 f^5_{t+1|t} + \hat{\beta}_6 f^6_{t+1|t} \qquad (4.15)$$

The implementation of this two-step approach to obtain the weighted combination forecasts over the out-of-sample period, extending from January 2001 to October 2010, is described below. First, data covering the period from January 1996 until June 1998 is employed to estimate the various models under consideration. Next, one-step-ahead rolling forecasts over the "pseudo" out-of-sample period commencing on July 1998 and ending on December 2000 are obtained. These forecasts are used to obtain the weights by estimating the regression specification described in equation 4.14.

Following Aksu & Gunter (1992), equation 4.14 is estimated under four alternative approaches, by imposing alternative restrictions to the coefficients and the constant term of the regression. As a first test, a constant term is not included, whereas no restriction is applied on the value of the weights (Method A). Next, I estimate the regression again without a constant term but restrict the weights to sum to unity, without imposing any restriction on the sign (Method B). Alternative specification of equation (14) includes a constant term, whereas no restriction is applied on the value of the weights (Method C). Finally, I estimate the equation with a constant term and with imposed restriction on the weights so as to sum to unity and to obtain only positive values (Method D).

Finally, the out-of-sample weighted combination forecasts are obtained by substituting the previously estimated weights and the out-of-sample forecasted values of individual models to equation 4.15.

4.4 In-sample estimation results

Table 4.3 and Table 4.4 present the results from estimated regressions 4.4 and 4.5 employed to test the existence of intra-week and monthly seasonality, respectively. The series clearly present a strong intra-week pattern, as the Monday, and Friday dummy variables are strongly significant. As expected, the Monday (Friday) dummy variable has a significant positive (negative) coefficient, consistent with reported increased buying (selling) activity on the specific days of the week. Obtained estimates from Table 4 suggest that there is no evidence of monthly seasonality for the MFIC series.

Table 4.3: Intra-week pattern of MFIC series

	Coefficient	t-statistic
Monday	0.082 ***	6.999
Tuesday	0.010	0.886
Wednesday	-0.015	-1.61
Thursday	-0.005	-0.357
Friday	-0.067 ***	-5.426
$MFIC_{t-1}$	-0.203 ***	-4.819

Table 4.4: Monthly seasonality of MFIC series

	Coefficient	t-statistic
January	-0.001	-0.036
February	-0.002	-0.140
March	0.013	0.831
April	-0.011	-0.732
May	0.005	0.433
June	-0.015	-1.222
July	0.022	1.279
August	0.004	0.264
September	-0.003	-0.198
October	0.006	0.260
November	-0.012	-0.695
December	-0.005	-0.292
$MFIC_{t-1}$	-0.222 ***	-4.811

Note: Reported t-statistics have been calculated with Newey-West autocorrelation consistent standard errors where the number of lags is automatically selected according to the Schwarz criterion. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively

Next, the four tables that follow present the in-sample estimation results from the AR(I)MA, ARFIMA and AR(I)MA-GARCH, the Regime Switching, the Heterogeneous Autoregressive (HAR) and the economic determinants model, correspondingly. The AR(I)MA(2,1), the ARFIMA(1,d,3), the AR(I)MA(3,5)-GARCH(5,4) and the Regime Switching (4,0) were found to minimize the Schwarz criterion. The previously reported significant day-of-the-week variables are also included in every specification.

Several diagnostic tests are employed to test the goodness-of-fit from the various models. The log likelihood value, the adjusted R^2, the Schwarz criterion, the Ljung-Box statistic for the 20^{th} lag of squared residuals as well as the F-statistic of the ARCH test for remained heteroskedasticity up to the 5^{th} lag are reported. The addition of GARCH specification in the residuals corrects for remained serial autocorrelation and heteroskedasticity. The AR(I)MA-GARCH specification presents the lowest value of the Schwarz criterion and the maximum log-likelihood value, whereas the ARFIMA model produces the highest value of R^2. Both the null hypothesis of no serial autocorrelation and no heteroskedasticity effects are strongly rejected for the AR(I)MA and the ARFIMA model. The value of the fractional differencing parameter *d* in the ARFIMA specification is not significantly different from zero suggesting that the series does not exhibit long memory features and a simple ARMA (1,3) model is able to capture the dynamics of the series.

Table 4.5: In-sample evidence from the AR(I)MA, ARFIMA and AR(I)MA-GARCH models

	AR(I)MA	ARFIMA	AR(I)MA-GARCH
Constant	-0.008*	-0.008*	-0.012***
d	-	0.000	-
φ_1	0.685***	0.901***	0.760***
φ_2	0.118**	-	-0.980***
φ_3	-	-	0.771***
θ_1	-0.934***	-1.144***	-0.956***
θ_2	-	0.139***	1.075***
θ_3	-	0.027	-0.992***
θ_4	-	-	0.075***
θ_5	-	-	-0.040***
b_0	-	-	0.003***
b_1	-	-	0.075***
b_2	-	-	0.291***
b_3	-	-	0.197***
b_4	-	-	-0.211**
b_5	-	-	-0.193***
c_1	-	-	-0.999***
c_2	-	-	0.939***
c_3	-	-	0.887***
c_4	-	-	-0.065***
Monday dummy	0.101***	0.100***	0.104***
Friday dummy	-0.053***	-0.055***	-0.046***
Log likelihood	293.341	299.442	465.785
Adj. R^2	0.122	0.132	0.125
BIC criterion	-0.434	-0.426	-0.624
$Q^2(20)$	258.043***	237.270***	13.300
ARCH test (5)	154.488***	146.620***	2.571

Note: The t-statistics of the AR(I)MA model have been estimated with HAC estimates of standard errors. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

Table 4.6 presents the results from the estimated Regime Switching model. The results from the two regimes differ significantly. The conditional mean term in the first regime is significantly different from zero and equal to 0.87 while the conditional mean in the second regime is not significantly different from zero. The first regime is characterized by high MFIC changes while the second regime is characterized by low MFIC changes and negative mean reversion. Moreover, there is high persistence in the first regime (p_{22}=0.99) while the persistence in the second regime is significantly lower (p_{11}=0.46).

Table 4.6: In-sample evidence from the Regime Switching model.

	Coefficient	t-Statistic
μ_1	0.870***	11.967
φ_1^1	-1.513***	-17.239
φ_1^2	0.868***	3.740
φ_1^3	1.161***	6.141
φ_1^4	0.728***	4.935
μ_2	-0.008	-1.367
φ_2^1	-0.142***	-5.212
φ_2^2	-0.081***	-3.174
φ_2^3	-0.091***	-3.535
φ_2^4	-0.011	-0.425
log(σ)	-1.791***	-85.421
Monday dummy	0.082***	6.421
Friday dummy	-0.057***	-4.505
Log likelihood	418.623	
Adj. R^2	-0.079	
BIC criterion	-0.583	
$Q^2(20)$	527.506***	
ARCH test (5)	377.959***	

Note: One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

Table 4.7 presents the results from the Heterogeneous Autoregressive (HAR) model. The coefficients of all three estimates of implied correlation, corresponding to different time-horizons, are highly significant and negative. The impact of lagged correlation strengthens as the time-horizon of aggregation increases.

Table 4.7: In-sample evidence from the Heterogeneous Autoregressive (HAR) Model

	Coefficient	t-statistic
α_0	-0.024***	-3.430
$\alpha_{(d)}$	-0.129**	-2.534
$\alpha_{(w)}$	-0.307**	-2.456
$\alpha_{(m)}$	-0.902***	-2.694
Monday dummy	0.017	1.296
Friday dummy	0.108***	7.774
Log likelihood	276.132	
Adj. R^2	0.109	
BIC criterion	-0.413	
$Q^2(20)$	264.281***	
ARCH test (5)	163.561***	

Note: The t-statistics of the model have been estimated with HAC estimates of standard errors. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

Table 4.8 reports the estimation results from the economic determinants model. The addition of three lags of the MFIC series minimizes the Schwarz criterion. Interestingly, at 10% significant level, the only statistically significant variables are the first and the second lagged values of MFIC, the lagged negative index returns, the lagged junk spread, the lagged volume of the S&P 100 index and the dummy variables representing the Monday and the Friday effect. Moreover, unreported results confirm the asymmetric response of the MFIC series to negative and positive contemporaneous index returns. The negative coefficients of the lagged default spread and trading volume suggest that an increase in the variables at time t-1 will actually reduce the current MFIC series level. This is in contrast with the widely reported positive contemporaneous relationship between the above-mentioned variables and volatility. A possible explanation for the decreasing effect of the variables to the series might stem from the information flow and their use. Specifically, an increase in the variables under consideration will increase market uncertainty directly and investors will shift their attention on information at the market level, resulting in an increase of the contemporaneous stock market correlations. At the subsequent period, investors will shift their attention on asset specific news and equity correlations are expected to decrease. The attention shift hypothesis has been proposed by Peng (2005) and Peng and Xiong

(2006) while Peng et al (2007) study the effect of this hypothesis on stock return comovements. One-step-ahead forecasts will be obtained from the estimation of the economic determinants model with only the variables that are significant at 10% significance level.

Table 4.8: In-sample evidence from the economic determinants model

	Coefficient	t-statistic
C	0.017	1.498
R^{-}_{t-1}	3.336***	3.391
R^{+}_{t-1}	-1.452	-1.405
FX_{t-1}	-1.055	-1.321
r_{t-1}	-0.038	-0.049
DIV_{t-1}	8.983	0.211
$TERM_{t-1}$	-11.194	-1.415
$JUNK_{t-1}$	-80.318*	-1.932
WTI_{t-1}	-0.033	-0.196
VOL_{t-1}	-0.050*	-1.714
$MFIC_{t-1}$	-0.181***	-3.322
$MFIC_{t-2}$	-0.088**	-2.246
$MFIC_{t-3}$	-0.007	-0.123
Monday dummy	0.084***	5.705
Friday dummy	-0.062***	-4.067
Log likelihood	298.16	
Adj. R^2	0.122	
BIC criterion	-0.391	
$Q^2(20)$	222.427***	
ARCH test (5)	117.069***	

Note: Two lagged values of the MFIC series are also included. The t-statistics of the model have been estimated with HAC estimates of standard errors. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

4.5 Evaluation of out-of-sample forecasting performance

4.5.1 Statistical measures

In this Section, results on the constructed out-of-sample forecasts under the models presented above are presented and their forecasting performance is assessed. For the first out-of-sample forecast, corresponding to 02/01/2001, the model is estimated over the sample period of 04/01/1996 – 29/12/2000, and subsequently, the relevant forecasted value for the following day is obtained. One-step-ahead forecasting values for the out-of-sample period are obtained by employing a rolling sample window.

Panel A of Table 4.9 presents the performance of AR(I)MA, AR(I)MA-GARCH, ARFIMA, Regime Switching, HAR and economic determinants models, while the last three columns present the results from the alternative combination methods employed. Equal and Schwarz weighted combination forecasts are obtained from equations (4.12) and (4.13), respectively. Combination forecasts with time varying weights have been formulated under four alternative approaches described at Section 3. Consistent with the findings of Aksu & Gunter (1992), the obtained forecasts from equation (4.14), estimated with a constant term and under the restriction of positive weights to sum to unity, are found to minimize the RMSE and, thereafter, are chosen to compete with the other specifications. For the rest of the Chapter, I shall refer to the results and the forecasting performance of the weighted combination forecast from Method D as the time varying weighted combination forecast.

To compare the predictive ability of the alternative forecasting models a number of statistical evaluation criteria are used. The accuracy of forecasts is initially evaluated based on the root mean squared error (RMSE) and the mean absolute error (MAE). RMSE is calculated as the square root of the average squared deviations of the forecasted values from the actual series, while MAE is measured by the average of the absolute value of forecast errors. Notably, the obtained results for the RMSE and the MAE do not vary substantially across the employed models. In terms both of RMSE and MAE, the minimum value is attributed to the Schwarz Weighted combination model, however being only 0.0005% and 0.0012% lower than the respective values for the equal weighted combination forecast. The results are in agreement with existing literature suggesting that combination forecasts are able to outperform single models.

The comparison of the forecasting accuracy of each individual model with the benchmark model of random walk, substitutes a direct test for the presence of market efficiency. Taking into account that the random walk model is nested to all alternative models, the McCracken (2007) test statistic, which is valid for comparison of nested models, is employed. The t-statistic is defined as below:

$$MSE - F = (T - h + 1)\frac{MSE_N - MSE_A}{MSE_A} \tag{4.16}$$

where T is the number of out-of-sample forecasts, h=1 is the h-step ahead forecasts, MSE_N is the MSE of the nested model (i.e. the random walk model) and MSE_A is the MSE of the alternative models. The calculated McCracken t-statistics are reported in Panel A of Table 4.9. Results suggest that the null hypothesis of equal forecasting accuracy is strongly rejected proposing that each alternative model produce smaller forecasting error when compared with the benchmark model.

The above-mentioned tests address the question of forecasting accuracy in terms of correct magnitude prediction. However, practitioners, and especially traders, are mostly interested in correctly predicting the direction of change of their portfolio so as to maintain accordingly the appropriate positions. First, the directional predictability of the various models employed is assessed under the mean correct prediction (MCP) measure. The MCP is computed as the percentage of observations for which the forecasting model correctly predicted the realized direction of change of MFIC (see Goncalves & Guidolin, 2006). Secondly, I employ the non-parametric market-timing test introduced by Timmerman & Pesaran (1992), or the PT test. For the purpose of the PT test, a contingency table of realized and forecasted values is created.

		Actual value ($\hat{y}_{t+1}$)	
		Up ($\hat{y}_{t+1}$ =1)	Down ($\hat{y}_{t+1}$ =0)
Forecasted value ($\hat{y}_{t+1}$)	Up ($\hat{y}_{t+1}$)	Hits (N_{uu})	False Alarms (N_{ud})
	Down ($\hat{y}_{t+1}$ =0)	Misses (N_{du})	Correct Rejections (N_{dd})

They showed that the PT test statistic could be expressed as:

$$PT = \frac{\sqrt{T}(H - F)}{\sqrt{\dfrac{\hat{\pi}_f(1-\hat{\pi}_f)}{\hat{\pi}_\alpha(1-\hat{\pi}_\alpha)}}} \tag{4.17}$$

where T is the sample size, $H = N_{uu} / (N_{uu} + N_{du})$ is the portion of correctly predicted "Up" moves, $F = N_{ud} / (N_{ud} + N_{dd})$ is the portion of "false alarms", $\hat{\pi}_{\alpha} = (N_{uu} + N_{du}) / T$ is the probability that the actual series will move upwards and $\hat{\pi}_{f} = (N_{uu} + N_{ud}) / T$ is the probability that the forecasted series will move upwards. The PT test follows the standard normal distribution under the null hypothesis of independence between the actual and forecasted values, i.e. forecasted series are not able to predict the sign of the actual series.

Panel B of Table 4.9 reports the results from the tests of directional accuracy. Remarkably, the obtained MCP values do not vary significantly among competing models, with the equal and the Schwartz combination forecasts successfully predicting the direction of forecasted value 60.54% times within the out-of-sample dataset. The PT test, or the market-timing test, suggests that, in all cases, the forecasted and the actual series are not independently distributed and in fact, there is a predictable pattern in the direction of changes in the MFIC series.

Table 4.9: Evaluation of the out-of-sample performance

Panel A										
	AR(I)MA	AR(I)MA - GARCH	ARFIMA	Regime Switching	HAR	E.D.	Equal Weighted combination	Time varying Weighted combination	Schwarz Weighted combination	Random walk
RMSE	17.828%	17.905%	18.010%	17.976%	17.943%	17.960%	17.716%	17.907%	17.715%	18.578%
MAE	11.851%	11.882%	12.132%	12.005%	12.125%	12.004%	11.804%	11.891%	11.803%	12.468%
McCracken t-statistic	212.427***	189.389***	158.358***	168.501***	177.996***	173.101***	246.508***	188.671***	246.676***	

Panel B									
	AR(I)MA	AR(I)MA - GARCH	ARFIMA	Regime Switching	HAR	E.D.	Equal Weighted combination	Time varying Weighted combination	Schwarz Weighted combination
MCP	59.773%	59.976%	56.860%	59.409%	57.345%	58.357%	60.542%	59.733%	60.542%
PT test	9.585***	9.728***	6.587***	9.199***	7.452***	8.208***	10.372***	9.64***	10.369***

Note: Panel A presents the Root Mean Square Error (RMSE) and the Mean Absolute Error (MAE) in percentage terms, corresponding to the AR(I)MA, AR(I)MA-GARCH, ARFIMA, Regime Switching, HAR, Economic Determinants (E.D.) and the random walk model as well as to the combination forecasts obtained with equal, time varying and Schwarz weights. The null hypothesis that the model produces equal forecasting accuracy with the random walk model is tested against the alternative of a better forecasting accuracy of the model with the McCracken test for MSE. Panel B presents the directional forecasting accuracy of the models under consideration. The Mean Correct Prediction (MCP) measure represents the number of times that the actual change in MFIC is correctly predicted by the forecasted series. The PT test is employed to test the independence between the actual and forecasted values. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

Additionally, for the purposes of pairwise comparison of the forecasting accuracy of the alternative specifications, the Modified Diebold-Mariano test is employed. The loss differential function is computed both in terms of the mean square error of the forecast and of the mean absolute error and is defined as $d_{jt} = [g(e_{jt}) - g(e_{it})], i \neq j$, where $g(e_{jt})$ is the loss function, and i , j = 1 for AR(I)MA, 2 for AR(I)MA – GARCH, 3 for ARFIMA, 4 for Regime Switching, 5 for HAR, 6 for the economic determinants model, 7 for the Schwarz weighted combination forecast, 8 for the equal combination forecast and 9 for the time varying combination forecast. The null hypothesis of equal forecasting accuracy is tested against the alternative that the forecasting model performs better than the benchmark model, i.e. $E(d_{jt}) < 0$. The Diebold-Mariano test statistic is defined as:

$$DM = \frac{\overline{d}}{\sqrt{\frac{2\pi \hat{f}_d(0)}{T}}} \sim N(0,1) \tag{4.18}$$

where $\overline{d}$ is the sample average of the loss differential, $\hat{f}_d(0)$ is the estimate of the spectral density at frequency zero and T is the number of observations. For h-step-ahead forecasts, the Modified Diebold-Mariano test statistic corrects for small sample sizes and autocorrelation of the loss differential following a Student-t distribution with T-1 degrees of freedom and equals to

$$DM_{\text{mod}} = \left[\frac{T+1-2h+h(h-1)/T}{T}\right] DM \tag{4.19}$$

Table 4.10 reports the t-statistics for the null hypothesis that, in terms of root mean squared error (RMSE), the model in row i performs equally well with model in column j[6]. The null hypothesis of equal errors is principally accepted for the majority of pairwise comparisons. Interestingly, however, at 95% confidence level, equal and Schwarz weighted combination forecasts are better than all models except for the AR(I)MA - GARCH, while the AR(I)MA model also outperforms the ARFIMA model.

[6] When the lost function is defined with regard to the Mean Absolute Error, test results are similar.

Table 4.10: Modified Diebold-Mariano tests

	AR(I)MA	AR(I)MA-GARCH	ARFIMA	Regime Switching	HAR	E.D.	Equal	Schwarz	Time - varying
AR(I)MA		-0.503	-2.216 **	-1.181	-0.896	-0.663	1.944	1.986	-0.904
AR(I)MA-GARCH	0.503		-0.588	-0.334	-0.166	-0.178	1.122	1.143	-0.025
ARFIMA	2.216	0.588		0.266	0.529	0.245	4.137	4.150	0.871
Regime Switching	1.181	0.334	-0.266		0.186	0.075	2.375	2.400	0.461
HAR	0.896	0.166	-0.529	-0.186		-0.112	2.359	2.309	0.199
E.D.	0.663	0.178	-0.245	-0.075	0.112		1.482	1.463	0.218
Combination Forecasts									
Equal	-1.944 **	-1.122	-4.137 ***	-2.375 ***	-2.359 ***	-1.482 *		0.177	-1.848 **
Schwarz	-1.986 **	-1.143	-4.150 ***	-2.400 ***	-2.309 **	-1.463 *	-0.177		-1.897 **
Time varying	0.904	0.025	-0.871	-0.461	-0.199	-0.218	1.848	1.897	

Note: The Table presents the t-statistics for the Modified Diebold-Mariano test. The null hypothesis that the model in rows performs equally well with the model in columns is tested against the alternative of a better forecasting accuracy of the model in rows. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively

4.5.2 Economic significance

In addition to the statistical evaluation measures, I investigate whether the predictability of MFIC is significant enough to generate abnormal profits. Following Harvey & Whaley (1992), Guo (2000) and Goncalves & Guidolin (2006), the out-of-sample forecasting performance of different forecasting models is evaluated based on the profitability of a trading strategy. The trading strategy is based on the idea of 'dispersion trading' but, in contrast to relevant studies (e.g. Driessen et al., 2009), is focused on exploiting the daily changes in MFIC (i.e. daily changes in implied volatilities of stock and index options) rather than differences between implied and realized volatilities of stock and index options. A long (short) dispersion trade involves short (long) position on near ATM straddles on S&P 100 index and long (short) position on a portfolio of near ATM straddles on S&P 100 component stocks. Straddles involve buying or selling equal amount of call and put options with the same maturity and strike price and provide an effective way of trading changes in implied volatility (see Guo 2000; Brooks & Oozeer, 2002; Ni et al., 2008). A long dispersion trade will generate profits if the change in the implied volatilities of the stock options is higher than the change in the index option implied volatility i.e. if MFIC decreases.

In detail, the correlation trading strategy proceeds as follows: if on a given day, MFIC is expected to increase/(decrease) on the following day, the investor goes short/(long) the dispersion trade. In order to build the dispersion trade portfolios, on each day for each asset and the index, I choose calls and puts of the shortest maturity with at least one call and one put with the same strike price. If more than one pair satisfies the criterion, I choose the one with moneyness closest to 1. Options with maturity less than 7 days, ask price lower than the bid price, non-positive bid price, moneyness levels higher than 1.15 or lower than 0.85 and zero open interest are discarded. Each day $1,000 worth of options are always bought and sold and the position is liquidated the next day. The funds may be freely invested at the riskless interest rate while the profits are not reinvested next day. The trading exercise is repeated every day in the out-of-sample period.

The value of a portfolio unit at day t is computed as follows:

$$V_t = \begin{cases} -\left(C_{I,t} + P_{I,t}\right) + \sum_{i=1}^{N} n_{i,t}\left(C_{i,t} + P_{i,t}\right) \text{ if long dispersion trade} \\ \left(C_{I,t} + P_{I,t}\right) - \sum_{i=1}^{N} n_{i,t}\left(C_{i,t} + P_{i,t}\right) \text{ if short dispersion trade} \end{cases} \tag{4.20}$$

where $C_{i,t}$ is the call price on stock i, $P_{i,t}$ is the put price on stock *i*, $C_{I,t}$ is the call price on the index, $P_{I,t}$ is the put price on the index, $n_{i,t} = \frac{N_{i,t} S_{I,t}}{\sum_{i=1}^{N} N_{i,t} S_{i,t}}$, $S_{i,t}$ is the closing price of stock *i*, $S_{I,t}$ is the closing price of the index, $N_{i,t}$ is the number of shares outstanding in stock *i*, all variables at day *t*. I assume that each day \$1000 are invested in the portfolio i.e. X_t = 1000 / $|V_t|$ portfolio units are bought/sold. If the portfolio requires funds for its initiation (i.e. V_t>0), the net gain of the portfolio is:

$$NG_{t+1} = X_t (V_{t+1} - V_t) \tag{4.21}$$

If the portfolio generates inflows at is initiation (i.e. V_t<0), the net gain of the portfolio is:

$$NG_{t+1} = X_t \left(V_{t+1} - V_t\right) + 2000\left(e^{r_t/252} - 1\right) \tag{4.22}$$

In order to avoid noisy signals, the weakest signals by applying the following filters: trading occurs only if the forecasted change in MFIC is higher than 0.1%, 0.5% and 1%.

Table 4.11 reports out-of-sample average returns, annualized Sharpe ratio and Leland's alpha. While Sharpe's ratio is an appropriate measure of profitability in the case of normal returns, Leland's alpha allows for non-normal trading strategy returns by taking into account higher-order moments of the return distribution. Using Rubinstein (1976) asset pricing model it is assumed that the returns on the market portfolio are i.i.d. over time and the agent has power utility characterized by a constant risk aversion coefficient γ. Under these assumptions, a marginal utility-adjusted abnormal return measure for the trading strategy is derived as follows:

$$A = E\left[\frac{G_{t+1}}{1{,}000}\right] - r_t - B\left(E\left[r_{m,t}\right] - r\right) \tag{4.23}$$

where $r_{m,t}$ is the return on the market portfolio, $B=\dfrac{Cov\left(E\left[G_{t+1}/1{,}000\right],-\left(1+r_{m,t}\right)^{-\gamma}\right)}{Cov\left(r_{m,t},-\left(1+r_{m,t}\right)^{-\gamma}\right)}$ is the market price of risk and $\gamma=\dfrac{\ln\left(E\left[1+r_{m,t}\right]\right)-\ln\left(1+r_{t}\right)}{Var\left(\ln\left(1+r_{m,t}\right)\right)}$ is a measure of risk aversion. The daily USD Libor and S&P 100 returns are used as the risk-free rate and the market returns, respectively. Finally, since the trading strategy returns are found to be stationary and non-normal (based on unreported results) the statistical significance of the results using bootstrapped 95% confidence intervals is assessed based on Politis & Romano (1994) method. The average block size is set equal to ten and 1000 bootstrap repetitions are used.

The reported results with no filter indicate that the trading strategy based on the AR(I)MA-GARCH, equal and Schwarz weighted combination forecasts produce significant positive average returns over the out-of-sample period. The highest return, Sharpe ratio and Leland's alpha is accomplished by the AR(I)MA-GARCH model. Moreover, when filters are applied, these three models yield significant positive returns across all filters.

Results below question the efficiency of the S&P100 index and stock option market but neglect the effect of transaction cost. Therefore transaction costs are incorporated in the analysis of the profitability of the trading strategies by using bid and ask quotes instead of the mid quotes. More specifically, I assume that the investor buys the options at the ask price and sells at the bid price. Table 4.12 reports out-of-sample average returns, annualized Sharpe ratio and Leland's alpha and the corresponding bootstrapped confidence intervals for the correlation trading strategy based on alternative forecasting models and filters after incorporating transaction cost. Not surprisingly, transactions costs have a major impact on the profitability of the trading strategies. The average daily return, Sharpe ratio and Leland's alpha are significantly negative across all forecasting models and filters applied. The results extend the findings of Driessen, Maenhout and Vilkov (2009), Goyal & Saretto (2009) and Neumann & Skiadopoulos (2013) who also report a significant economic impact of bid-ask spreads in option trading.

Table 4.11: Trading strategy based on out-of-sample forecasts – without transaction costs

Panel A						
	No filter			Filter 0.1%		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	0.195%	0.425	0.002	0.192%	0.417	0.002
C.I. (95%)	*(-0.008%,0.401%)*	*(-0.092,0.918)*	*(0.000,0.004)*	*(-0.021%,0.407%)*	*(-0.073,0.911)*	*(0.000,0.004)*
AR(I)MA-GARCH	0.335%*	0.746*	0.003*	0.335%*	0.746*	0.003*
C.I. (95%)	*(0.108%,0.582%)*	*(0.229,1.197)*	*(0.001,0.006)*	*(0.109%,0.560%)*	*(0.274,1.218)*	*(0.001,0.006)*
ARFIMA	0.188%	0.408	0.002	0.186%	0.405	0.002
C.I. (95%)	*(-0.037%,0.399%)*	*(-0.111,0.973)*	*(0.000,0.004)*	*(-0.037%,0.408%)*	*(-0.094,0.977)*	*(-0.001,0.004)*
RS	0.201%	0.438	0.002	0.193%	0.420	0.002
C.I. (95%)	*(-0.055%,0.441%)*	*(-0.097,1.037)*	*(-0.001,0.004)*	*(-0.042%,0.433%)*	*(-0.099,0.997)*	*(-0.001,0.004)*
E.D.	0.181%	0.393	0.002	0.177%	0.382	0.002
C.I. (95%)	*(-0.054%,0.398%)*	*(-0.107,0.955)*	*(0.000,0.004)*	*(-0.047%,0.378%)*	*(-0.084,0.966)*	*(0.000,0.004)*
HAR	-0.009%	*-0.042*	*0.000*	-0.008%	*-0.040*	*0.000*
C.I. (95%)	*(-0.240%,0.236%)*	*(-0.553,0.531)*	*(-0.002,0.002)*	*(-0.248%,0.222%)*	*(-0.562,0.489)*	*(-0.003,0.002)*
Combination Forecasts						
Schwarz	0.271%*	0.600*	0.003*	0.273%*	0.604*	0.003*
C.I. (95%)	*(0.061%,0.471%)*	*(0.119,1.105)*	*(0.001,0.005)*	*(0.053%,0.470%)*	*(0.139,1.095)*	*(0.001,0.005)*
Equal	0.268%*	0.593*	0.003*	0.272%*	0.602*	0.003*
C.I. (95%)	*(0.061%,0.490%)*	*(0.116,1.090)*	*(0.001,0.005)*	*(0.060%,0.484%)*	*(0.109,1.086)*	*(0.001,0.005)*
Time varying	0.060%	0.115	0.000	0.058%	0.110	0.000
C.I. (95%)	*(-0.157%,0.292%)*	*(-0.348,0.684)*	*(-0.002,0.003)*	*(-0.159%,0.275%)*	*(-0.419,0.631)*	*(-0.002,0.003)*

Panel B						
	Filter 0.5%			Filter 1%		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	0.191%	0.417	0.002	0.188%	0.409	0.002
C.I. (95%)	*(-0.029%,0.394%)*	*(-0.051,0.887)*	*(0.000,0.004)*	*(-0.024%,0.382%)*	*(-0.087,0.898)*	*(0.000,0.004)*
AR(I)MA-GARCH	0.329%*	0.735*	0.003*	0.336%*	0.755*	0.003*
C.I. (95%)	*(0.105%,0.576%)*	*(0.204,1.218)*	*(0.001,0.006)*	*(0.100%,0.573%)*	*(0.263,1.230)*	*(0.001,0.006)*
ARFIMA	0.211%	0.463	0.002	0.219%	0.485	0.002
C.I. (95%)	*(-0.009%,0.434%)*	*(-0.038,1.054)*	*(0.000,0.004)*	*(-0.005%,0.439%)*	*(-0.033,1.112)*	*(0.000,0.004)*
RS	0.194%	0.424	0.002	0.154%	0.336	0.001
C.I. (95%)	*(-0.047%,0.425%)*	*(-0.143,1.003)*	*(-0.001,0.004)*	*(-0.096%,0.374%)*	*(-0.206,0.958)*	*(-0.001,0.004)*
E.D.	0.161%	0.347	0.002	0.179%	0.390	0.002
C.I. (95%)	*(-0.057%,0.364%)*	*(-0.177,0.857)*	*(-0.001,0.003)*	*(-0.037%,0.398%)*	*(-0.142,0.951)*	*(-0.001,0.004)*
HAR	-0.042%	*-0.121*	*-0.001*	-0.042%	*-0.122*	*-0.001*
C.I. (95%)	*(-0.286%,0.200%)*	*(-0.695,0.456)*	*(-0.003,0.002)*	*(-0.298%,0.197%)*	*(-0.642,0.460)*	*(-0.003,0.002)*
Combination Forecasts						
Schwarz	0.284%*	0.631*	0.003*	0.265%*	0.591*	0.002*
C.I. (95%)	*(0.073%,0.502%)*	*(0.144,1.119)*	*(0.001,0.005)*	*(0.036%,0.469%)*	*(0.072,1.096)*	*(0.000,0.005)*
Equal	0.276%*	0.613*	0.003*	0.265%*	0.590*	0.002*
C.I. (95%)	*(0.076%,0.499%)*	*(0.143,1.131)*	*(0.000,0.005)*	*(0.055%,0.480%)*	*(0.109,1.093)*	*(0.000,0.005)*
Time varying	0.056%	0.106	0.000	0.058%	0.111	0.000
C.I. (95%)	*(-0.177%,0.283%)*	*(-0.416,0.633)*	*(-0.002,0.003)*	*(-0.171%,0.278%)*	*(-0.378,0.632)*	*(-0.002,0.003)*

Note: The Table reports the average daily return, the annualized Sharpe ratio, and Leland's alpha corresponding to MFIC forecasts from the models as well as to the combination forecasts obtained with the Schwarz, equal and time varying weights. Results are reported for the trading strategy without filter and with 0.1% (Panel A), 0.5% and 1% (Panel B) filters. The bootstrapped 95% confidence intervals (C.I.) are reported in parentheses. One asterisk indicates the rejection of the null hypothesis of zero average return, Sharpe Ratio or Leland's alpha at significance level 5%.

Table 4.12: Trading strategy based on out-of-sample forecasts with transaction cost

Panel A						
	No filter			Filter 0.1%		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	-12.531%*	-8.530*	-0.125*	-12.528%*	-8.524*	-0.125*
C.I. (95%)	*(-13.910%,-11.288%)*	*(-10.987,-7.007)*	*(-0.138,-0.114)*	*(-13.870%,-11.322%)*	*(-10.973,-6.931)*	*(-0.139,-0.113)*
AR(I)MA-GARCH	-12.234%*	-9.889*	-0.123*	-12.253%*	-9.889*	-0.122*
C.I. (95%)	*(-13.456%,-11.223%)*	*(-11.095,-8.852)*	*(-0.134,-0.112)*	*(-13.344%,-11.232%)*	*(-11.184,-8.870)*	*(-0.133,-0.112)*
ARFIMA	-13.362%*	-3.692*	-0.134*	-13.311%*	-3.672*	-0.133*
C.I. (95%)	*(-16.158%,-11.218%)*	*(-10.836,-2.716)*	*(-0.162,-0.112)*	*(-16.163%,-11.230%)*	*(-10.801,-2.680)*	*(-0.163,-0.112)*
RS	-13.684%*	-3.774*	-0.137*	-13.693%*	-3.767*	-0.137*
C.I. (95%)	*(-16.475%,-11.546%)*	*(-10.804,-2.757)*	*(-0.167,-0.116)*	*(-16.600%,-11.484%)*	*(-10.807,-2.727)*	*(-0.165,-0.115)*
E.D.	-12.930%*	-7.700*	-0.129*	-12.903%*	-7.683*	-0.129*
C.I. (95%)	*(-14.409%,-11.571%)*	*(-10.869,-6.046)*	*(-0.146,-0.116)*	*(-14.408%,-11.604%)*	*(-10.821,-6.146)*	*(-0.144,-0.115)*
HAR	-13.178%*	*-7.701**	*-0.132**	-13.128%*	*-7.664**	*-0.131**
C.I. (95%)	*(-14.898%,-11.786%)*	*(-10.671,-6.130)*	*(-0.148,-0.117)*	*(-14.770%,-11.736%)*	*(-10.581,-6.219)*	*(-0.149,-0.116)*
Combination Forecasts						
Schwarz	-12.585%*	-8.451*	-0.126*	-12.551%*	-8.426*	-0.125*
C.I. (95%)	*(-14.065%,-11.354%)*	*(-10.813,-6.930)*	*(-0.140,-0.114)*	*(-13.967%,-11.322%)*	*(-10.801,-6.856)*	*(-0.139,-0.114)*
Equal	-12.583%*	-8.451*	-0.126*	-12.569%*	-8.435*	-0.126*
C.I. (95%)	*(-14.028%,-11.392%)*	*(-10.851,-6.885)*	*(-0.138,-0.114)*	*(-13.896%,-11.348%)*	*(-10.808,-6.954)*	*(-0.141,-0.113)*
Time varying	-12.966%*	-8.326*	-0.130*	-12.954%*	-8.308*	-0.129*
C.I. (95%)	*(-14.518%,-11.709%)*	*(-10.683,-6.867)*	*(-0.145,-0.117)*	*(-14.311%,-11.690%)*	*(-10.675,-6.934)*	*(-0.144,-0.117)*

Panel B

	Filter 0.5%			Filter 1%		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	-12.412%*	-8.470*	-0.124*	-12.336%*	-8.418*	-0.123*
C.I. (95%)	*(-13.863%,-11.167%)*	*(-10.953,-6.755)*	*(-0.138,-0.113)*	*(-13.694%,-11.171%)*	*(-10.844,-6.913)*	*(-0.136,-0.111)*
AR(I)MA-GARCH	-12.141%*	-9.826*	-0.122*	-11.996%*	-9.726*	-0.120*
C.I. (95%)	*(-13.238%,-11.055%)*	*(-11.035,-8.835)*	*(-0.133,-0.111)*	*(-13.131%,-10.892%)*	*(-10.972,-8.706)*	*(-0.132,-0.109)*
ARFIMA	-13.034%*	-3.607*	-0.131*	-12.823%*	-3.551*	-0.128*
C.I. (95%)	*(-15.992%,-11.107%)*	*(-10.754,-2.690)*	*(-0.158,-0.110)*	*(-15.663%,-10.749%)*	*(-10.617,-2.637)*	*(-0.156,-0.107)*
RS	-13.521%*	-3.730*	-0.136*	-13.400%*	-3.697*	-0.134*
C.I. (95%)	*(-16.500%,-11.446%)*	*(-10.623,-2.774)*	*(-0.165,-0.113)*	*(-16.511%,-11.220%)*	*(-10.587,-2.722)*	*(-0.163,-0.114)*
E.D.	-12.819%*	-7.634*	-0.128*	-12.645%*	-7.541*	-0.126*
C.I. (95%)	*(-14.428%,-11.490%)*	*(-10.761,-6.128)*	*(-0.144,-0.115)*	*(-14.183%,-11.274%)*	*(-10.615,-6.058)*	*(-0.142,-0.112)*
HAR	-13.022%*	-7.616*	-0.130*	-12.859%*	-7.525*	-0.129*
C.I. (95%)	*(-14.785%,-11.650%)*	*(-10.600,-6.165)*	*(-0.147,-0.115)*	*(-14.478%,-11.407%)*	*(-10.421,-5.987)*	*(-0.146,-0.115)*
Combination Forecasts						
Schwarz	-12.373%*	-8.365*	-0.124*	-12.218%*	-8.285*	-0.122*
C.I. (95%)	*(-13.632%,-11.171%)*	*(-10.827,-6.859)*	*(-0.138,-0.112)*	*(-13.635%,-11.087%)*	*(-10.702,-6.691)*	*(-0.135,-0.111)*
Equal	-12.369%*	-8.362*	-0.124*	-12.210%*	-8.280*	-0.122*
C.I. (95%)	*(-13.755%,-11.168%)*	*(-10.810,-6.873)*	*(-0.137,-0.112)*	*(-13.475%,-11.070%)*	*(-10.717,-6.807)*	*(-0.136,-0.111)*
Time varying	-12.740%*	-8.213*	-0.128*	-12.530%*	-8.116*	-0.125*
C.I. (95%)	*(-14.048%,-11.498%)*	*(-10.544,-6.860)*	*(-0.142,-0.115)*	*(-13.928%,-11.325%)*	*(-10.512,-6.723)*	*(-0.139,-0.113)*

Note: The Table reports the average daily return, the annualized Sharpe ratio, and Leland's alpha corresponding to MFIC forecasts from the models as well as to the combination forecasts obtained with the Schwarz, equal and time varying weights, after accounting for transaction cost. Results are reported for the trading strategy without filter and with 0.1% (Panel A), 0.5% and 1% filters (Panel B). The bootstrapped 95% confidence intervals (C.I.) are reported in parentheses. One asterisk indicates the rejection of the null hypothesis of zero average return, Sharpe Ratio or Leland's alpha at significance level 5%.

Most importantly, results imply that only investors who face zero or limited transaction costs (e.g. market makers) can generate abnormal profits by correctly predicting the change in MFIC. The elimination of profitability when accounting for transactions costs can be attributed to the simultaneous trading position of the investor to all constituent stocks of the index. An alternative strategy could be to hold long/short positions on options only on some representative stocks of the index (e.g. those representing the 75% of the index market capitalization). This strategy could significantly reduce transaction costs and enhance profitability.

4.6 Robustness checks

Previous analysis is based on an extensive dataset that covers the period of 1996-2010. To assess the robustness of the results across different samples, I use expanding in-sample windows of three, six, nine and twelve years and produce out-of-sample forecasts for the following three years. In that way, a series of non-overlapping forecasted values is created. Specifically, the in-sample periods are: 01/01/1996 – 12/31/1998, 01/01/1996 – 12/31/2001, 01/01/1996 – 12/31/2004, 01/01/1996 – 12/31/2007 corresponding to out-of-sample forecasting periods: 01/01/1999 – 12/31/2001, 01/01/2002 – 12/31/2004, 01/01/2005 – 12/31/2007 and 01/01/2008 – 10/29/2010, respectively.

4.6.1 Evaluation of out-of-sample performance

Table 4.13 presents the forecasting performance of the models during the different sampling periods as well as the average values throughout the out-of-sample period (01/01/1999 – 10/29/2010). Notably, the RMSE and the MAE do not vary significantly across models in the same period but do vary across the different sampling schemes. Specifically, when compared with the main estimation window results, RMSE is lower during the first two forecasting periods, while the last forecast period that includes the 2008 crisis attains the largest values in terms of both the RMSE and the MAE. In terms of directional accuracy, the results do not vary significantly among the different sample windows and the main estimation window. The PT test across all periods and models, with the exception of the AR(I)MA-GARCH model for the last forecasting period that includes the 2008 crisis, suggests that the direction of the change of the MFIC series can be predicted.

The last column of Table 4.13 presents the average values of the statistical measures for evaluation of the out-of-sample performance of the competing models for the out-of-sample period of 1999-2010, where the forecasted values are obtained as described above. The minimum RMSE and MAE values are obtained from the time varying and the Schwartz combination forecast, respectively. In terms of the best performing model on average, the weighted combination forecast produces superior forecasts throughout the different sample periods, while the Schwarz and the equal combination forecast were the best performing models for the main estimation window. Obtained results are in agreement with vast literature that suggests that combination forecasts produce better forecasting accuracy than individual model specifications.

Obtained t-statistics from the McCracken (2007) test suggest that the null hypothesis of equal forecasting accuracy with the random walk is strongly rejected for all models, except for the Regime Switching model where the null hypothesis is rejected only at 10% significance level. In contrast, the PT test suggests that the models are able to track a predictable pattern in the directional changes of the MFIC series. Overall, results across different forecasting periods suggest the existence of predictability in the MFIC series, consistent with results obtained from the main estimation window.

Table 4.13: Robustness checks. Evaluation of the out-of-sample performance

	1999-2001					2002-2004				
	RMSE	MAE	MSE-F	MCP	PT test	RMSE	MAE	MSE-F	MCP	PT test
ARMA	14.618%	11.191%	74.092***	60.347%	5.620 ***	12.462%	9.515%	100.880***	61.772%	6.396 ***
ARMA-GARCH	14.536%	11.096%	83.411***	61.148%	6.247 ***	12.423%	9.470%	106.243***	62.566%	6.828 ***
ARFIMA	15.290%	11.743%	3.256***	54.206%	2.312 **	13.040%	10.022%	26.579***	54.630%	2.532 **
RS	14.624%	11.178%	73.359***	58.611%	4.668 ***	16.408%	12.415%	-261.749	56.878%	3.740 ***
HAR	15.087%	11.601%	23.681***	56.208%	3.434 ***	12.998%	10.004%	31.599***	56.349%	3.517 ***
E.D.	15.040%	11.459%	28.512***	59.680%	5.296 ***	18.206%	12.284%	-354.536	60.847%	5.869 ***
Equal	14.559%	11.110%	80.716***	60.881%	5.933 ***	12.593%	9.567%	83.172***	60.317%	5.606 ***
Time-varying	14.608%	11.139%	75.130***	60.214%	5.562 ***	12.459%	9.525%	101.241***	61.376%	6.196 ***
Schwarz	14.552%	11.104%	81.570***	60.881%	5.931 ***	12.601%	9.570%	82.021***	60.582%	5.749 ***
Random Walk	15.324%	11.753%				13.267%	10.219%			

	2005 - 2007					2008 - 2010				
	RMSE	MAE	MSE-F	MCP	PT test	RMSE	MAE	MSE-F	MCP	PT test
ARMA	18.189%	12.212%	74.417***	61.406%	6.064 ***	22.978%	14.685%	42.599***	54.278%	2.375 **
ARMA-GARCH	18.251%	12.245%	68.795***	60.743%	5.642 ***	23.336%	14.933%	19.579***	53.156%	1.586
ARFIMA	18.271%	12.461%	67.035***	59.284%	5.097 ***	22.784%	14.506%	55.509***	56.522%	3.428 **
RS	22.836%	15.548%	-228.420	58.090%	4.398 ***	30.053%	20.792%	-271.291	55.259%	2.757 ***
HAR	18.294%	12.395%	64.934***	60.477%	5.841 ***	22.665%	14.491%	63.629***	56.522%	3.653 ***
E.D.	18.206%	12.284%	72.862***	60.477%	5.682 ***	23.138%	14.778%	32.189***	55.680%	2.973 ***
Equal	18.347%	12.328%	60.210***	61.671%	6.335 ***	22.921%	14.697%	46.373***	55.259%	2.853 ***
Time-varying	18.117%	12.227%	81.029***	60.875%	5.787 ***	22.946%	14.572%	44.699***	56.381%	3.553 ***
Schwarz	18.363%	12.340%	58.754***	61.538%	6.267 ***	22.940%	14.719%	45.106***	55.400%	2.924 ***
Random Walk	19.066%	12.918%				23.654%	15.046%			

Table 4.13: Robustness checks. Evaluation of the out-of-sample performance (cont'd)

	Average				
	RMSE	MAE	MSE-F	MCP	PT test
AR(I)MA	17.434%	11.862%	64.727***	59.5222%	10.527 ***
AR(I)MA-GARCH	17.540%	11.894%	54.927***	59.4886%	10.754 ***
ARFIMA	17.647%	12.150%	45.263***	56.1575%	6.124 ***
RS	18.183%	12.118%	-0.909*	57.2342%	10.422 ***
HAR	17.564%	12.090%	52.776***	57.4024%	7.692 ***
E.D.	17.604%	12.002%	49.097***	59.2194%	9.895 ***
Equal	17.418%	11.809%	66.255***	59.5895%	10.799 ***
Time varying	17.387%	11.810%	69.142***	59.7577%	10.827 ***
Schwarz	17.420%	11.809%	66.045***	59.6568%	10.918 ***
RW	18.172%	12.444%			

Note: Panel A presents the Root Mean Square Error (RMSE) and the Mean Absolute Error (MAE) in percentage terms. The t-statistic of the McCracken test for MSE is also reported. The null hypothesis that the model produces equal forecasting accuracy with the random walk model is tested against the alternative of a better forecasting accuracy of the model. The Mean Correct Prediction (MCP) measure represents the number of times that the actual change in MFIC is correctly predicted by the forecasted series. The PT test is employed to test the independence between the actual and forecasted values. The results are reported across different forecast periods and on average. One, two and three asterisks indicate the rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

4.6.2 Economic significance

Table 4.14 presents the results from the correlation trading strategy during the four out-of-sample periods without accounting for transaction cost. During all the forecast periods, except the last one, there is at least one model that produces significantly positive returns although these models differ during the various sub-periods. It is interesting that most of the models attain significant positive returns when forecasting the period before the 2008 crisis while none of the models attains significant returns during the last forecast period that includes the 2008 crisis. The last column in Table 4.14 presents the average return, Sharpe ratio and Leland alpha across the four out-of-sample periods. The highest return is attained by the AR(I)MA-GARCH specification consistent with the results from the main estimation window. The second, third and fourth model with the highest average return are the ARFIMA, the equally weighted and the Schwarz weighted combination forecasts. When transaction cost is taken into account, all of the models produce significant negative returns during the four out-of-sample periods (Table 4.15) similar to the results in Table 4.12 for the main estimation window. Overall, the main economic result supporting the existence of profitable correlation trading strategies based on MFIC forecasts that fade out when transaction costs are taken into account is robust across different in-sample sizes and forecast periods with the exception of the after crisis forecast period.

Table 4.14: Robustness checks. Trading strategy based on out-of-sample forecasts – without transaction costs

	1999 – 2001			2002 – 2004		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	-0.191%	-0.797	-0.002	-0.009%	-0.044	0.000
C.I. (95%)	*(-0.472%,0.101%)*	*(-1.951,0.256)*	*(-0.005,0.001)*	*(-0.310%,0.307%)*	*(-1.014,0.869)*	*(-0.004,0.003)*
AR(I)MA-GARCH	0.162%	0.532	0.001	0.172%	0.506	0.002
C.I. (95%)	*(-0.097%,0.428%)*	*(-0.415,1.525)*	*(-0.001,0.004)*	*(-0.118%,0.471%)*	*(-0.395,1.418)*	*(-0.001,0.005)*
ARFIMA	0.277%	0.965	0.003	0.563%*	1.705*	0.006*
C.I. (95%)	*(-0.031%,0.585%)*	*(-0.168,2.088)*	*(0.000,0.006)*	*(0.231%,0.928%)*	*(0.632,2.673)*	*(0.002,0.009)*
RS	-0.137%	-0.591	-0.002	-0.055%	-0.184	-0.001
C.I. (95%)	*(-0.400%,0.118%)*	*(-1.610,0.379)*	*(-0.004,0.001)*	*(-0.333%,0.238%)*	*(-0.954,0.728)*	*(-0.003,0.002)*
E.D.	0.439%*	1.580*	0.004*	0.122%	0.355	0.001
C.I. (95%)	*(0.140%,0.759%)*	*(0.442,2.687)*	*(0.001,0.007)*	*(-0.173%,0.419%)*	*(-0.522,1.325)*	*(-0.002,0.004)*
HAR	-0.197%	*-0.819*	*-0.002*	-0.016%	*-0.067*	*0.000*
C.I. (95%)	*(-0.467%,0.041%)*	*(-1.718,0.145)*	*(-0.005,0.000)*	*(-0.290%,0.249%)*	*(-0.892,0.778)*	*(-0.003,0.003)*
Combination Forecasts						
Schwarz	-0.064%	-0.320	-0.001	0.144%	0.419	0.001
C.I. (95%)	*(-0.325%,0.201%)*	*(-1.379,0.720)*	*(-0.003,0.002)*	*(-0.119%,0.412%)*	*(-0.420,1.301)*	*(-0.001,0.004)*
Equal	-0.065%	-0.322	-0.001	0.148%	0.432	0.001
C.I. (95%)	*(-0.343%,0.211%)*	*(-1.353,0.621)*	*(-0.004,0.002)*	*(-0.109%,0.422%)*	*(-0.392,1.267)*	*(-0.001,0.004)*
Time varying	0.134%	0.426	0.001	-0.044%	-0.152	0.000
C.I. (95%)	*(-0.136%,0.381%)*	*(-0.491,1.370)*	*(-0.001,0.004)*	*(-0.335%,0.243%)*	*(-1.140,0.694)*	*(-0.003,0.002)*

Table 4.14: Robustness checks. Trading strategy based on out-of-sample forecasts – without transaction costs (cont'd)

	2005 - 2007			2008 – 2010		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	0.493%*	1.394*	0.005*	-0.048%	-0.086	-0.001
C.I. (95%)	*(0.175%,0.840%)*	*(0.639,2.072)*	*(0.002,0.009)*	*(-0.641%,0.532%)*	*(-0.997,0.888)*	*(-0.006,0.005)*
AR(I)MA-GARCH	0.585%*	1.666*	0.006*	0.318%	0.525	0.003
C.I. (95%)	*(0.261%,0.951%)*	*(0.932,2.295)*	*(0.003,0.009)*	*(-0.211%,0.910%)*	*(-0.459,1.340)*	*(-0.002,0.009)*
ARFIMA	0.204%	0.545	0.002	-0.051%	-0.092	-0.001
C.I. (95%)	*(-0.097%,0.487%)*	*(-0.198,1.761)*	*(-0.001,0.005)*	*(-0.601%,0.511%)*	*(-1.001,0.907)*	*(-0.006,0.005)*
RS	0.436%*	1.226*	0.004*	-0.094%	-0.163	-0.001
C.I. (95%)	*(0.106%,0.836%)*	*(0.259,1.961)*	*(0.001,0.008)*	*(-0.673%,0.511%)*	*(-1.045,0.893)*	*(-0.007,0.004)*
E.D.	0.218%	0.584	0.002	-0.310%	-0.524	-0.003
C.I. (95%)	*(-0.057%,0.493%)*	*(-0.172,1.702)*	*(-0.001,0.005)*	*(-0.876%,0.284%)*	*(-1.415,0.406)*	*(-0.009,0.003)*
HAR	0.372%*	*1.037**	*0.004**	-0.018%	*-0.036*	*0.000*
C.I. (95%)	*(0.048%,0.715%)*	*(0.158,1.789)*	*(0.001,0.007)*	*(-0.620%,0.587%)*	*(-0.955,1.058)*	*(-0.007,0.005)*
Combination Forecasts						
Schwarz	0.500%*	1.415*	0.005*	0.030%	0.044	0.000
C.I. (95%)	*(0.195%,0.873%)*	*(0.619,2.128)*	*(0.002,0.009)*	*(-0.525%,0.593%)*	*(-0.867,0.996)*	*(-0.006,0.006)*
Equal	0.514%*	1.458*	0.005*	0.020%	0.026	0.000
C.I. (95%)	*(0.188%,0.910%)*	*(0.659,2.144)*	*(0.002,0.009)*	*(-0.534%,0.566%)*	*(-0.812,1.120)*	*(-0.006,0.006)*
Time varying	0.483%*	1.364*	0.005*	-0.158%	-0.271	-0.002
C.I. (95%)	*(0.159%,0.843%)*	*(0.514,2.067)*	*(0.001,0.008)*	*(-0.751%,0.440%)*	*(-1.187,0.774)*	*(-0.008,0.004)*

Table 4.14: Robustness checks. Trading strategy based on out-of-sample forecasts – without transaction costs (cont'd)

	Average		
Model	Average Daily Return	Annualized Sharpe Ratio	Leland's alpha
AR(I)MA	0.061%	0.117	0.000
AR(I)MA-GARCH	0.309%	0.807	0.003
ARFIMA	0.248%	0.781	0.002
RS	0.038%	0.072	0.000
E.D.	0.117%	0.499	0.001
HAR	0.035%	0.029	0.000
Combination Forecasts			
Schwarz	0.152%	0.390	0.001
Equal	0.154%	0.399	0.001
Time varying	0.104%	0.342	0.001

Note: The Table reports the average daily return, the annualized Sharpe ratio, and Leland's alpha corresponding to MFIC forecasts from the models as well as to the combination forecasts obtained with the Schwarz, equal and time varying weights. Results are reported for the trading strategy without filters. The bootstrapped 95% confidence intervals (C.I.) are reported in parentheses. The results are reported across different forecast periods and on average. One asterisk indicates the rejection of the null hypothesis of zero average return, Sharpe Ratio or Leland's alpha at significance level 5%.

Table 4.15: Robustness checks. Trading strategy based on out-of-sample forecasts with transaction cost.

	1999 – 2001			2002 – 2004		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	-10.398%*	-12.336*	-0.104*	-14.339%*	-12.375*	-0.143*
C.I. (95%)	*(-11.562%,-9.272%)*	*(-12.904,-11.847)*	*(-0.117,-0.093)*	*(-15.626%,-13.166%)*	*(-13.384,-11.523)*	*(-0.158,-0.132)*
AR(I)MA-GARCH	-9.713%*	-12.434*	-0.097*	-13.872%*	-12.606*	-0.139*
C.I. (95%)	*(-10.899%,-8.651%)*	*(-12.995,-11.944)*	*(-0.108,-0.087)*	*(-15.156%,-12.803%)*	*(-13.432,-11.885)*	*(-0.151,-0.128)*
ARFIMA	-9.591%*	-12.144*	-0.096*	-13.123%*	-13.216*	-0.131*
C.I. (95%)	*(-10.769%,-8.514%)*	*(-12.821,-11.662)*	*(-0.108,-0.085)*	*(-14.038%,-12.271%)*	*(-13.934,-12.537)*	*(-0.140,-0.123)*
RS	-10.242%*	-12.454*	-0.102*	-14.341%*	-12.330*	-0.143*
C.I. (95%)	*(-11.421%,-9.214%)*	*(-12.990,-12.006)*	*(-0.115,-0.093)*	*(-15.784%,-12.998%)*	*(-13.113,-11.669)*	*(-0.157,-0.131)*
E.D.	-9.615%*	-12.327*	-0.096*	-14.320%*	-12.172*	-0.143*
C.I. (95%)	*(-10.765%,-8.620%)*	*(-12.981,-11.794)*	*(-0.108,-0.086)*	*(-15.806%,-13.011%)*	*(-13.168,-11.412)*	*(-0.158,-0.130)*
HAR	-10.298%*	*-12.228**	*-0.103**	-14.265%*	*-12.435**	*-0.143**
C.I. (95%)	*(-11.456%,-9.059%)*	*(-12.849,-11.698)*	*(-0.116,-0.092)*	*(-15.623%,-12.960%)*	*(-13.219,-11.804)*	*(-0.156,-0.130)*
Combination Forecasts						
Schwarz	-10.192%*	-12.354*	-0.102*	-14.049%*	-12.560*	-0.141*
C.I. (95%)	*(-11.505%,-9.114%)*	*(-12.952,-11.862)*	*(-0.114,-0.091)*	*(-15.237%,-12.921%)*	*(-13.348,-11.909)*	*(-0.153,-0.130)*
Equal	-10.197%*	-12.362*	-0.102*	-14.041%*	-12.553*	-0.140*
C.I. (95%)	*(-11.388%,-9.082%)*	*(-12.962,-11.912)*	*(-0.114,-0.092)*	*(-15.285%,-12.959%)*	*(-13.344,-11.887)*	*(-0.152,-0.129)*
Time varying	-9.787%*	-12.491*	-0.098*	-14.443%*	-12.474*	-0.144*
C.I. (95%)	*(-10.847%,-8.765%)*	*(-13.020,-12.033)*	*(-0.109,-0.088)*	*(-15.905%,-13.230%)*	*(-13.312,-11.812)*	*(-0.158,-0.133)*

Table 4.15: Robustness checks. Trading strategy based on out-of-sample forecasts with transaction cost (cont'd)

	2005 – 2007			2008 – 2010		
Model	Average Daily Return	Sharpe Ratio	Leland's alpha	Average Daily Return	Sharpe Ratio	Leland's alpha
AR(I)MA	-8.853%*	-9.261*	-0.089*	-16.496%*	-7.335*	-0.165*
C.I. (95%)	*(-10.461%,-7.182%)*	*(-10.452,-8.100)*	*(-0.106,-0.071)*	*(-20.562%,-13.378%)*	*(-11.552,-5.936)*	*(-0.205,-0.134)*
AR(I)MA-GARCH	-8.597%*	-9.516*	-0.086*	-15.392%*	-9.827*	-0.154*
C.I. (95%)	*(-10.205%,-6.953%)*	*(-10.702,-8.308)*	*(-0.104,-0.069)*	*(-18.257%,-12.941%)*	*(-11.398,-8.900)*	*(-0.181,-0.130)*
ARFIMA	-9.042%*	-9.160*	-0.091*	-16.153%*	-7.340*	-0.162*
C.I. (95%)	*(-10.958%,-7.341%)*	*(-10.409,-8.072)*	*(-0.109,-0.072)*	*(-19.813%,-13.359%)*	*(-11.279,-5.957)*	*(-0.196,-0.132)*
RS	-9.104%*	-9.029*	-0.091*	-20.396%*	-3.096*	-0.205*
C.I. (95%)	*(-10.840%,-7.226%)*	*(-10.259,-7.835)*	*(-0.110,-0.071)*	*(-30.058%,-13.554%)*	*(-10.451,-2.614)*	*(-0.298,-0.137)*
E.D.	-9.110%*	-9.214*	-0.091*	-17.385%*	-6.584*	-0.174*
C.I. (95%)	*(-10.919%,-7.178%)*	*(-10.445,-8.074)*	*(-0.109,-0.073)*	*(-22.255%,-13.677%)*	*(-11.208,-5.457)*	*(-0.221,-0.140)*
HAR	-9.096%*	-9.017*	-0.091*	-19.937%*	-3.036*	-0.200*
C.I. (95%)	*(-10.989%,-7.360%)*	*(-10.295,-7.833)*	*(-0.109,-0.072)*	*(-29.766%,-13.380%)*	*(-10.803,-2.563)*	*(-0.300,-0.136)*
Combination Forecasts						
Schwarz	-9.024%*	-9.121*	-0.090*	-16.280%*	-7.326*	-0.163*
C.I. (95%)	*(-10.816%,-7.087%)*	*(-10.402,-8.023)*	*(-0.107,-0.073)*	*(-20.292%,-13.250%)*	*(-11.304,-6.039)*	*(-0.201,-0.131)*
Equal	-9.014%*	-9.115*	-0.090*	-16.310%*	-7.336*	-0.163*
C.I. (95%)	*(-10.769%,-7.323%)*	*(-10.345,-8.034)*	*(-0.108,-0.073)*	*(-20.000%,-13.168%)*	*(-11.176,-6.020)*	*(-0.204,-0.134)*
Time varying	-8.943%*	-9.237*	-0.090*	-20.174%*	-3.073*	-0.203*
C.I. (95%)	*(-10.712%,-7.179%)*	*(-10.413,-8.069)*	*(-0.108,-0.071)*	*(-30.511%,-13.594%)*	*(-11.138,-2.591)*	*(-0.289,-0.137)*

Table 4.15: Robustness checks. Trading strategy based on out-of-sample forecasts with transaction cost (cont'd)

		Average	
Model	Average Daily Return	Annualized Sharpe Ratio	Leland's alpha
AR(I)MA	-12.522%	-10.327	-0.125
AR(I)MA-GARCH	-11.894%	-11.096	-0.119
ARFIMA	-11.977%	-10.465	-0.120
RS	-13.521%	-9.227	-0.136
E.D.	-12.607%	-10.074	-0.126
HAR	-13.399%	-9.179	-0.134
Combination Forecasts			
Schwarz	-12.386%	-10.340	-0.124
Equal	-12.390%	-10.341	-0.124
Time varying	-13.337%	-9.319	-0.134

Note: The Table reports the average daily return, the annualized Sharpe ratio, and Leland's alpha corresponding to MFIC forecasts. Results are reported for the trading strategy without filters. The bootstrapped 95% confidence intervals (C.I.) are reported in parentheses. The results are reported across different forecast periods and on average. One asterisk indicates the rejection of the null hypothesis of zero average return, Sharpe Ratio or Leland's alpha at significance level 5%.

4.7 Conclusions

Understanding the dynamics that govern the evolution of correlation is of vital importance in asset pricing theory and other financial applications. An extensive dataset allows assessing the impact of periods of financial turbulence associated with lower asset returns and increased volatility.

First, this study contributes to the existing literature on the predictability of option-implied measures. Specifically, the predictability of the MFIC series is assessed under alternative specifications (AR(I)MA, AR(I)MA-GARCH, ARFIMA, Regime Switching, Heterogeneous Autoregressive and Economic Determinants models). Additionally, combination forecasts with constant and time varying weights have been estimated. The forecasting performance has been investigated for constructed out-of-sample forecasts, both in terms of statistical and economic significance. The Schwartz combination forecasting values is found to outperform the competing models, in terms of correct magnitude forecast, whereas the equal and Schwartz combination models have the greatest predictive power in terms of successful correct direction prediction. When compared to the random walk model, used as the benchmark, all models produce better forecasting accuracy. Turning to the economic significance of the obtained forecasts, the AR(I)MA-GARCH and combination forecasts of implied correlation are successful at generating profitable strategies. However, when transaction cost is taken into account, no economically significant profits can be attained. The results are robust across different in-sample windows and forecast periods.

In conclusion, the existence of predictable patterns in the S&P 100 market, supported by statistical measures, cannot be confirmed by the implemented trading strategy, and thus, the efficiency of the S&P100 index and stock option market cannot be rejected. Results imply that only investors who face zero or limited transaction costs (e.g. market makers) can generate abnormal profits by correctly predicting the change in MFIC. The elimination in the correlation trading strategy costs can be attributed to the simultaneous trading position of the investor to all constituent stocks of the index.

Chapter 5

Realized Hedge Ratio:
Predictability and Hedging Performance

Chapter Abstract: In this Chapter, I explore the dynamic properties and predictability of the Realized Minimum Variance Hedge Ratio (RMVHR), constructed from five-minute spot and future returns of two stock indices and two exchange rates. Based on previous findings on the realized beta framework, the distributional properties of the RMVHR are compared to realized variance and covariance. The long memory traits of the individual series are less pronounced in the realized hedge ratio series suggesting that common distribution traits of the realized variance and covariance process are neutralized when the realized hedge ratio is deducted. Triggered by the findings, I propose the direct modelling of the RMVHR series under a number of econometric models and assess the out-of-sample performance with statistical measures and economic significance criteria. Results from statistical measures provide evidence of predictable dynamics in the evolution of the realized hedge ratio series, thus suggesting that the efficient market hypothesis for the spot and futures market of the S&P 100 and FTSE 100 indices, as well as for the EUR/USD and GBP/USD foreign exchange rate markets is rejected. In terms of risk reduction, the realized hedge ratio forecasts dominate conventional methods that use daily data while the benefit is pronounced when economic gains are considered. The superior performance of RMVHR methods holds across different asset classes but is more conspicuous in the case of stock indices. Finally, results remain robust for alternative sampling frequencies and the inclusion of transaction costs.

5.1 Introduction

The increased availability of high-frequency data has stimulated the interest of practitioners and academics, giving rise to a new area of research in financial modeling where multi-dimensional high-frequency data are utilized to estimate, model and forecast the second moments of asset returns. Andersen et al. (1998) introduced the notions of realized volatility and covariance as model-free estimators of the true latent process, computed from intraday return data. ABDL (2001), ABDE (2001) and Barndorff-Nielsen & Shephard (2002) proved that, according to the theory of quadratic variation, by allowing the sampling frequency to tend to zero, the realized measures are unbiased and efficient estimators of the integrated processes, which essentially become observable, thus enabling direct estimation. In the forecasting context, Koopman, Jungbacker, and Hol (2005) and Blair et al. (2010) amongst others have provided evidence of superior informational content of the realized measures when compared to estimators derived from daily closing prices.

Andersen et al. (2006) extended their previous work on realized volatility and correlation, and defined realized beta as the ratio of realized covariance between asset and market returns to market variance. Under continuous-time stochastic volatility diffusion process of the price, the realized beta is a consistent estimator of the true integrated beta. Additionally, they advocate that any common persistence trait of the covariance and variance processes could be neutralized when forming the beta ratio. In addition, the study assesses the comparative predictability of realized beta versus the (co)variance predictability and finds that the former is much smaller. Notably, the predictability of the short-run beta, modelled through a simple autoregressive process, is much higher than the predictability of the long-run beta estimated from an ARFIMA specification.

In a similar context, the Realized Minimum Variance Hedge Ratio (RMVHR, hereafter) is defined as the ratio of the realized covariance of futures and spot returns divided by the futures realized variance. This study is motivated by the findings of Andersen et al. (2006) on the differential distribution properties of the realized (co)variance and beta, and the additive value of utilizing intraday information in dynamic hedging. More specifically, I address the question whether forecasting the dynamics of the RMVHR per se results in substantial benefit to the hedger in terms of risk reduction and economic value while the results are compared with those obtained from conventional models that use daily data. The methodology is in contrast to previous studies that employ econometric specifications on daily returns to model the variance-covariance matrix, construct out-of-sample forecasts and ultimately, calculate the hedge ratio.

Only a limited number of studies have examined the information content of intraday data in a dynamic hedging context. Lai and Sheu (2008; 2010) use data on the S&P 500 index and argue that encompassing realized volatility measures in Generalized Conditional Heteroskedasticity (GARCH) models provides substantial benefit to the hedger in terms of risk reduction and economic value. Moreover, Yeh, Huang and Hsu (2008) provide evidence of superior hedging performance of ARMA(1,1) forecasts of realized hedge ratio for the S&P 500 index. Based on intraday data on currency futures, Harris, Shen and Stoja (2010) indicate that, when compared to the RMVHR used as benchmark, the parametric variance-covariance models based on daily data perform poorly in terms of hedging effectiveness. They attribute the low hedging performance of the conditional daily models to the low persistence and, hence, the unpredictability of the RMVHR. Lastly, McMillan & Garcia (2010) utilize data on the Spanish IBEX 35 Index and advocate that the portfolio variance is minimized when the hedge ratio is estimated from daily returns while the realized hedge ratio yields superior Sharpe ratio.

The study in this Chapter makes three contributions to the ongoing discussion about the value of employing high-frequency data to the estimation of the hedge ratio. First, a thorough evaluation of the time-series characteristics of the realized volatility, realized covariance and RMVHR is performed and the differential properties of the distributions are assessed. Second, a horse race among alternative model specifications is performed and the statistical significance of the predictability of the RMVHR series per se is evaluated. Third, I examine the improvement in hedging performance against conventional modelling and forecasting techniques. To this end, one-step-ahead forecasts of the RMVHR under various econometric processes are generated. Finally, an extensive dataset of equity indices and foreign exchange rates is used and differentiated patterns across asset classes are examined.

Overall, empirical results from alternative time-series specifications provide evidence of predictability in the dynamics of the RMVHR series per se. Statistical evaluation criteria suggest that the time-varying weighted combination forecasts and forecasts based on the Heterogeneous Autoregressive model (HAR) are the best performing forecasts for the stock indices and exchange rates, respectively. In addition, alternative econometric specifications predict successfully the directional change of the series approximately 70% of the times throughout the out-of-sample period, while results do not vary significantly across models.

Importantly, the direct forecast of the RMVHR series, based on intraday data, improves hedging portfolio performance in terms of risk reduction, Sharpe ratio and mostly in terms of economic gains. In specific, results from the percentage risk reduction metric suggest that, the improvement, when switching from daily to intraday returns, ranges within 0.1% and 0.6%. Notably, the hedger's benefit is substantial when taking into account both the average return and the variance of the hedge portfolio. In the majority of cases, the use of intraday returns and the direct forecast of the series results in substantial improvement of the Sharpe ratio and economic gains. The results hold across the different asset classes, although the benefits are lower in the case of exchange rates. Lastly, the main results are relatively robust for a range of sampling frequencies and the incorporation of transaction cost.

The remaining Chapter is organized as follows. Section 2 presents the methodology for the derivation of the Realized Minimum Variance Hedge Ratio and the econometric models employed to forecast the hedge ratio. Section 3 describes the dataset used in the study along with the descriptive statistics of realized variance, covariance and hedge ratio. Sections 4 and 5 present the in-sample estimation results and the out-of-sample forecast evaluation under statistical and economic metrics, respectively. At Sections 6 and 7, the effect of alternative sampling frequencies and transaction costs to the out-of-sample economic significance of employed models is assessed. Section 8 concludes.

5.2 The Realized Minimum Variance Hedge Ratio and the Forecasting Models

Consider an investor with a long (short) position in the spot market. The hedge ratio denotes the number of futures contracts that the investor is willing to sell (buy) in order to offset the risk deriving from fluctuations of the spot market. The return of the hedged portfolio at time *t* is given by:

$$r_{p,t} = r_{s,t} - \beta_t r_{f,t} \tag{5.1}$$

where $r_{s,t}$ and $r_{f,t}$ are the logarithmic returns, from *t-1* to *t*, of the cash position in the spot and the futures market, respectively, and β_t is the hedge ratio. The optimal hedge ratio is obtained by minimizing the variance of the portfolio and equals to:

$$\beta^*_{t|\Omega_{t-1}} = \frac{\sigma_{sf,t|\Omega_{t-1}}}{\sigma^2_{f,t|\Omega_{t-1}}} \tag{5.2}$$

where $\sigma_{sf,t}$ and $\sigma_{f,t}$ are the spot and futures returns covariance and futures variance, respectively, conditional on the information set Ω, available at time *t-1*.

The vast majority of previous studies estimate the variance-covariance matrix using daily closing prices. Andersen and Bollerslev (1998) show that the Realized Volatility (RV), defined as the sum of squared intraday returns, sampled at non-overlapping intervals of frequency Δ, is a consistent and efficient estimator of the true latent volatility. In essence, RV is defined as follows:

$$RV_{i,t} = \sum_{m=1}^{1/\Delta} r^2_{i,t-1+m\cdot\Delta,\Delta} \tag{5.3}$$

where *i=s,f* for the spot and futures returns, respectively and *1/Δ* is the number of intraday intervals.

Similarly, the Realized Covariance (RC) can be defined as the cross-product of squared intraday returns.

$$RC_{ij,t} = \sum_{m=1}^{1/\Delta} r_{i,t-1+m\cdot\Delta,\Delta} \cdot r_{j,t-1+m\cdot\Delta,\Delta} \tag{5.4}$$

Andersen et al. (2006) introduced the notion of realized beta, while Harris, Shen and Stoja (2010) defined the Minimum Variance Hedge Ratio as the optimal hedge ratio calculated from intraday data. For the purposes of this study, the Realized Minimum Variance Hedge Ratio (RMVHR) is defined as follows:

$$RMVHR_t = \frac{RC_{sf,t}}{RV_{f,t}} \tag{5.5}$$

Forecasting Models

Vast majority of existing literature supports the presence of intraweek and seasonality patterns in the dynamics of the return distribution from closing prices. With the main purpose of this study laying in the predictability of the daily RMVHR, the impact of such regularities on the evolution of the series is assessed. First, I examine the presence of the day-of-the week effect through the specification:

$$y_t = \sum_{i=1}^{5} \alpha_i D_{i,t} + \gamma D_{roll} + \delta y_{t-1} + \varepsilon_t \tag{5.6}$$

where y is the RMVHR and D_i are the dummy variables corresponding to the five days of the week. A dummy variable, D_{roll}, taking a value of 1 at the day the futures series is rolled to the next contract and 0 otherwise, is added to the specification to control for the possible effects of futures contract rollover.

Moreover, the so-called month effect is assessed through the statistical significance of the obtained parameters from the following regression.

$$y_t = \sum_{i=1}^{12} \alpha_i M_{i,t} + \delta y_{t-1} + \varepsilon_t \tag{5.7}$$

where y is the RMVHR and M_i is a dummy variable that equals to 1 for month i, where i=1,2,…,12 for each month of the year, and 0 otherwise. Equations 5.6 and 5.7 have been augmented with the inclusion of the lagged value of RMVHR to account for the effects of autocorrelation.

Several econometric specifications that capture different traits of the dynamics of the series are employed to model the evolution of the series per se. For notation reasons, in paragraphs that follow, the Realized Minimum Variance Hedge Ratio is represented by the variable *y*.

ARMA, ARFIMA and ARMA- GARCH models

The ARMA and ARFIMA models are employed to investigate the presence of short and long memory dynamics of the hedge ratio series, respectively. The following ARMA(r,m) model is estimated:

$$y_t = c + \sum_{r=1}^{r} \varphi_r y_{t-r} + \sum_{m=1}^{m} \vartheta_m \varepsilon_{t-m} + \varepsilon_t \tag{5.8}$$

where φ and θ are the autoregressive and moving average parameters, respectively. The ARFIMA model is an extension of the ARMA model by allowing the integration order to take fractional values. The specification of the ARFIMA(r,d,m) model is as follows:

$$(1 - \sum_{r=1}^{r} \varphi_r L^r)(1-L)^d (y_t - \mu) = (1 + \sum_{m=1}^{m} \vartheta_m L^m)\varepsilon_t \tag{5.9}$$

where d is the differencing order, L is the backward-shift operator, μ is the expectation of y_t, and ε_t is the white noise error term. In the case of $0 < d < 0.5$ ($-0.5 < d < 0$), the process is stationary while exhibiting long memory characteristics, with positive (negative) dependence.

Furthermore, the ARMA specification is augmented so as to allow time variation in the variance structure of the residuals. The error terms of the resulting ARMA-GARCH model follow a normal distribution with zero mean and variance h_t^2 equal to:

$$h_t^2 = b_0 + \sum_{q=1}^{q} b_q \varepsilon_{t-q}^2 + \sum_{p=1}^{p} c_p h_{t-p}^2 \tag{5.10}$$

Regime Switching Model

In order to capture potential asymmetries in the RMVHR process, a dynamic Regime Switching (RS) model is employed. More specifically, the transition between the regimes is governed by a Markov chain, two regimes are assumed and the coefficient on the lagged dependent variable is allowed to be regime-varying, i.e.

$$\sum_{r=1}^{r}\Phi_{s_t}^{r}\left(L\right)\Delta^{d}\left(y_t-\mu_{s_t}^{r}\right)=\varepsilon_t \tag{5.11}$$

The transitions between the regimes s_t = 1 and s_t = 2 are given by a Markov chain with transition probabilities $p_{ij}= P(s_t = j|s_{t-1} = i)$ for i, j = 1,2.

Heterogeneous Autoregressive Model

Based on the theory of heterogeneous market hypothesis, Corsi (2009) developed the Heterogeneous Autoregressive (HAR) model. The fundamental underlying concept is induced by the empirical finding that investors' investment horizon affect different components of the volatility structure. In other words, dealers and market makers trade on a notable high intraday frequency, whereas, institutional investors have a longer-term investment horizon and rebalance their position less frequently. The HAR model aims to gauge persistence through a simple autoregressive representation, which aggregates daily *(d)*, weekly *(w)* and monthly *(m)* information sets. In an attempt to capture the long run dynamics of the RMVHR, the HAR specification is applied as below described.

$$y_t^{(d)}=\alpha_0+\alpha_{(d)}y_{t-1}^{(d)}+\alpha_{(w)}y_{t-1}^{(w)}+\alpha_{(m)}y_{t-1}^{(m)}+\varepsilon_t \tag{5.12}$$

where $y_{t-1}^{(w)}$ and $y_{t-1}^{(m)}$ is the lagged aggregated RMVHR over weekly and monthly horizon.

Combination Forecasts

In addition to the forecasts based on the above-mentioned models, I assess the forecasting accuracy of combination forecasts, which in essence, accumulate the information provided by individual models. Three alternative methods of combining forecasts are applied. The first, and simplest, method of combining is the arithmetic average of the forecasts from individual models, f_t^E

$$f_t^E=\frac{1}{5}\sum_{j=1}^{5}f_t^j \tag{5.13}$$

where f_t^j, is the forecasted values based on the time-series models described above (where j = 1 for ARMA, 2 for ARMA – GARCH, 3 for ARFIMA, 4 for regime switching model and 5 for HAR model). For notation reasons, the superscripts of forecasting models are maintained throughout the methods described below.

Instead of equal weighting across all models, the theory of combination forecasting suggests that assigning heavier weight to the best performing models can produce better results than the equal counterpart. Following Kolassa (2011), the model with the minimum Schwarz criterion is defined as the best performing model, while the weights of the other models are determined as below.

$$f_t^{BIC} = \sum_{j=1}^{5} w_{BIC}^{j} \cdot f_t^{j} \tag{5.14}$$

where $w_{BIC}^{j} = \dfrac{\exp(-\frac{1}{2}\Delta_{BIC}^{j})}{\sum_{j=1}^{5}\exp(-\frac{1}{2}\Delta_{BIC}^{j})}$, $\Delta_{BIC}^{j} = BIC_j - BIC_k$, *BIC(j)* is the Schwarz criterion of the j^{th} model, and *BIC(k)* is the minimum Schwarz criterion of model *k*.

Both of the methods described above assign constant weights to the alternative models throughout the out-of-sample period. As an alternative approach, the weights can be time-varying and chosen to minimize the mean forecast error of the following regression.

$$y_t = \beta_0 + \sum_{i=1}^{5} \beta_i f_{t|t-1}^{i} + \varepsilon_t \tag{5.15}$$

The two-step implementation procedure is described thoroughly in Section 4.3.

5.3 Data Description

The predictability and hedging performance of the RMVHR is empirically assessed for two asset classes, namely equity indices and foreign exchange rates. Intraday data for the S&P 500 and the FTSE 100 indices as well as the EUR/USD and the GBP/USD foreign exchange rates were obtained from Olsen & Associates. Following ABDL (2001) and ABDE (2001), who study the distribution of the exchange rate and the stock return volatility, the analysis is based on five-minute return series while results from alternative sampling frequencies are discussed in Section 6. The sample period spans from January

01, 2009 to December 31, 2012. Excluding weekends and holidays, the sample size for the S&P 500, the FTSE 100 and the foreign exchange rates is 987, 992 and 982 days, respectively. For forecasting purposes, the in-sample estimation period runs from 1/1/2009 to 30/6/2010 and the out-of-sample forecasting period extends from 1/7/2010 to 31/12/2012. In the case of combination forecasts, the "new" in-sample period runs from 1/1/2009 to 30/09/2009 and the "pseudo" out-of-sample period from 1/10/2009 to 30/06/2010. Continuous futures series are obtained from the nearest-to delivery contract and rolled to the next month contract on the 10^{th} day of the delivery month of the current contract.

Following ABDL (2001), several filtering rules are applied. First, weekend and holidays observations according to the trading calendar of the each exchange are excluded. For FX rates in specific, the weekend period is defined from Friday, 21:05 GMT until Sunday 21:00 GMT while, in case of holidays, observations from 21:05 GMT of the previous day to 21:00 GMT on the same day are excluded. Secondly, series of ten zero or constant returns are removed to avoid thin trading effects or possible lapses in the Reuter's data feed that Olsen uses to extract the return series. Finally, quotes of the first 15 minutes every day are filtered out to remove the open auction effects. The Flash Crash day, May 6, 2010, has also been removed from the S&P 500 series.

Table 5.1 exhibits the descriptive statistics for the realized variance of the spot and futures returns along with the realized covariance measured as the cross-product of the returns. Notably, the equity markets exhibit higher volatility and covariance compared to the foreign exchange rate markets. The distributions are leptokurtic and positively skewed while the Jarque-Bera (JB) statistic also suggests that the normality assumption is strongly rejected. As manifested by the autocorrelation function (ACF) statistic, the serial autocorrelation of the series is high and significant up to the 10^{th} lag, displaying high levels of persistency. According to the Augmented Dickey-Fuller (ADF) and the Phillips–Perron (PP) test statistics, the null hypothesis of unit root is rejected for all series across all assets.

Table 5.1: Summary Statistics of the Spot and Futures (Co)Variance

	S&P 500			FTSE 100			EUR/USD			GBP/USD		
	Spot Variance	Futures Variance	Covariance	Spot Variance	Futures Variance	Covariance	Spot Variance	Futures Variance	Covariance	Spot Variance	Futures Variance	Covariance
Mean	0.896	0.928	0.881	0.915	0.987	0.908	0.501	0.511	0.498	0.498	0.504	0.490
Std. Dev.	1.198	1.263	1.205	0.907	0.961	0.898	0.341	0.344	0.340	0.475	0.486	0.468
Skewness	3.999	4.289	4.108	2.631	2.584	2.662	2.285	2.266	2.277	3.262	3.383	3.345
Kurtosis	27.736	33.072	29.731	12.239	11.899	12.600	9.866	9.752	9.819	18.448	19.831	19.503
JB	2.769***	4.015***	3.209***	0.467***	0.438***	0.498***	0.278***	0.271***	0.275***	1.151***	1.346***	1.297***
ρ1	-0.111***	-0.111***	0.797***	0.004***	-0.010***	0.763***	0.001***	0.003***	0.755***	-0.001***	0.005***	0.848***
ρ10	0.002***	0.000***	0.499***	-0.048***	-0.054***	0.582***	-0.011***	-0.009***	0.517***	-0.095***	-0.097***	0.695***
ADF	-22.266***	-22.282***	-4.640***	-12.282***	-31.756***	-3.969***	-6.291***	-6.266***	-4.973***	-7.490***	-7.492***	-6.727***
PP	-35.023***	-35.053***	-10.432***	-31.309***	-31.759***	-14.662***	-31.257***	-31.208***	-13.568***	-31.384***	-31.184***	-8.628***

Note. The Table reports summary statistics of the daily spot and futures (co)variance, constructed from five-minute intraday returns, for the whole sample period of January 1, 2009 to December 31, 2012. It is noted that for obtaining the covariance matrix, the futures returns after the closure of the spot market are dropped. The Jarque-Bera (JB) tests for normality and the critical value, at 5% significance level, is 5.99. The values ρ1 and ρ10 are the autocorrelation function (ACF) coefficients for the 1st and the 10th lag. The Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) test for the presence of unit root. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively. The mean and standard deviation have been scaled by 10^4 and the Jarque-Bera statistic by 10^{-4}.

Table 5.2 displays the descriptive statistics of the RMVHR while Figure 5.1 plots the series over the whole sample period. On average, the RMVHR for the equity indices are smaller in absolute values while displaying greater volatility compared to the RMVHR of FX rates. The distributions are negatively skewed while the normality assumption is strongly rejected in all cases. The Phillips-Perron and the ADF test of stationarity indicate that the unit-root hypothesis is strongly rejected. Furthermore, the ACF values suggest that autocorrelation is present in the RMVHR series, although the values obtained are significantly lower and decay faster than the corresponding values for the covariance of the series. Moreover, an ARFIMA(p,d,q) model is employed so as to test for the presence of long memory in the series. The estimated fractional differential parameter d is in the range of [0.38, 0.47] for realized variances and covariances (with the exception of GBP/USD) and in the range of [0, 0.32] for realized hedge ratio. Results are in accordance with the findings of Andersen et al. (2006), suggesting lower persistence and lower degree of fractional integration for the ratio series.

Table 5.2: Summary Statistics of Realized Minimum Variance Hedge Ratio (RMVHR)

	S&P 500	FTSE 100	EUR/USD	GBP/USD
Mean	0.934	0.914	0.969	0.972
Std. Dev.	0.058	0.058	0.024	0.020
Skewness	-3.232	-2.673	-2.623	-0.956
Kurtosis	31.737	25.048	16.049	5.880
JB	3.563***	2.127***	0.809***	0.049***
$\rho 1$	0.171***	0.126***	0.414***	0.282***
$\rho 10$	0.148***	0.008***	0.332***	0.140***
ADF	-4.280***	-4.280***	-4.280***	-4.280***
PP	-30.044***	-30.044***	-30.044***	-30.044***

Note. The Table reports the summary statistics for the RMVHR, for the whole sample period of January 1, 2009 to December 31, 2012. The Jarque-Bera (JB) tests for normality and the critical value, at 5% significance level, is 5.99. The values ρ1 and ρ10 are the autocorrelation function (ACF) coefficients for the 1st and the 10th lag. The Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) test for the presence of unit root. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively. The Jarque-Bera statistic has been scaled by 10^{-4}.

Figure 5.1: Time Evolution of Realized Minimum Variance Hedge Ratio for the S&P 500 and the FTSE 100 Indices, the EUR/USD and the GBP/USD FX Rates

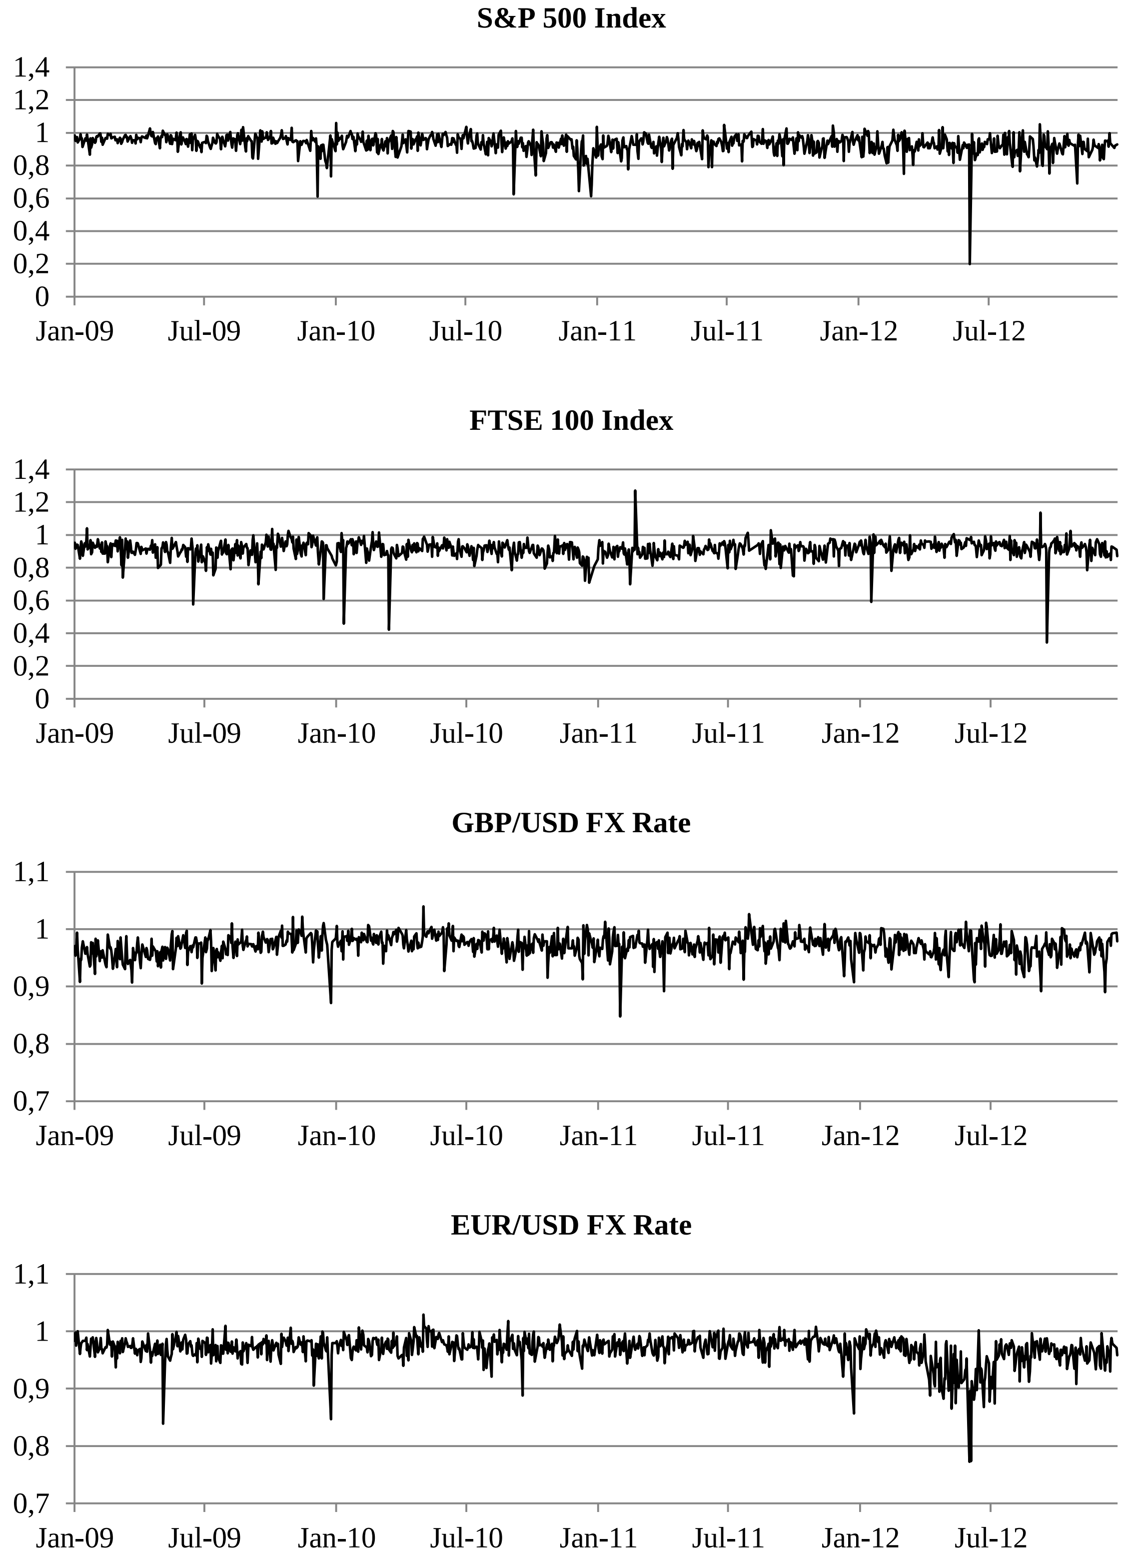

5.4 In-Sample Estimation Results

Table 5.3 reports the estimated in-sample model parameters for the ARMA, ARMA-GARCH, and ARFIMA models, employed on the RMVHR, for the two stock indices and the two foreign exchange rates. Table 5.4 and 5.5 report the estimated in-sample model parameters for the regime switching and the HAR model, respectively. All tables also report a number of goodness-of-fit tests, namely the log-likelihood value, the adjusted R^2 and the Schwarz criterion. The forecasting models have been augmented with the significant dummy variables from the seasonality tests discussed in Section 2, e.g. M_3 and M_{10} take a value of 1 in March and October, respectively, and 0 otherwise. The rollover dummy (Roll D) takes the value of 1 on a futures rollover day, and 0 otherwise. One, two and three asterisks denote rejection of the null hypothesis at 10%, 5% and 1%, respectively. Where applicable, the models have been estimated for all plausible combinations of autoregressive and moving average order terms up to the fifth lag and finally, the number of lags is selected based on the minimum Schwartz criterion.

A noteworthy result in Table 5.3 is the value of the fractional differential parameter *d* in the ARFIMA specification being significantly different from zero and within the range of (0,0.5) in all series, except the S&P 500 index, thus indicating that realized hedge ratio is stationary while exhibiting long memory and positive dependence between distant observations.

The regime switching model attains the maximum log-likelihood value for all assets, except GBP/USD, and the lowest Schwartz criterion for the FTSE 100 index and the EUR/USD. The ARMA-GARCH and ARMA models produce the lowest Schwartz criterion for S&P 500 and GBP/USD realized hedge ratio series, respectively. The ARMA specification attains the highest adjusted R^2 for both stock indices and GBP/USD.

Table 5.3: In-Sample Evidence from the ARMA, ARMA-GARCH and ARFIMA Models

	ARMA				ARMA - GARCH				ARFIMA			
	S&P 500	**FTSE 100**	**EUR/USD**	**GBP/USD**	**S&P 500**	**FTSE 100**	**EUR/USD**	**GBP/USD**	**S&P 500**	**FTSE 100**	**EUR/USD**	**GBP/USD**
c	0.955 ***	0.902 ***	0.977 ***	1.001 ***	0.959 ***	0.913 ***	0.977 ***	0.990 ***	0.956 ***	0.899 ***	0.978 ***	0.957 ***
φ_1	-0.126	-0.075		1.313 ***		0.981 ***		0.994 ***	-0.114			1.000 ***
φ_2	-0.770 ***	-0.036		-0.314 ***								
φ_3	-0.449 **	0.927 ***					0.125 ***					
θ_1	0.103	0.168 ***	0.060	-1.168 ***	0.064	-0.944 ***		-0.825 ***		-0.082	-0.078	-0.970 ***
θ_2	0.935 ***	0.125 **		0.153 ***				-0.149 ***				
θ_3	0.465	-0.932 ***										
θ_4	0.091											
θ_5	-0.036											
d									0.099	0.157 ***	0.126 **	0.109 ***
b_0					0.000 ***	0.000	0.000 **	0.000 ***				
b_1					0.050 ***	-0.022 ***	-0.027 ***	-0.026 ***				
c_1					1.941 ***	1.020 ***	1.024 ***					
c_2					-2.053 ***							
c_3					1.783 ***							
c_4					-0.738 ***							
M_3			-0.002				-0.003				-0.002	
M_4			-0.007 ***				-0.010 ***				-0.008 **	
M_6			-0.003				-0.004 **				-0.005 **	
M_7			-0.008 ***				-0.010 ***				-0.009 ***	
M_8			-0.007				-0.007 ***				-0.006	
M_{10}			-0.002				-0.003				-0.002	
M_{12}	-0.059 ***		-0.012 **		-0.080 ***		-0.005		-0.057 ***		-0.013 ***	
Mon			-0.002				0.000 ***				-0.003	
Tue		0.018 **				0.021 ***				0.025 ***		
Wed		0.025 ***				0.027 ***				0.030 ***		
Thu		0.013				0.026 ***				0.030 ***		
Roll D	-0.064				0.030 ***				-0.069			
$Adj.R^2$	0.230	0.140	0.025	0.294	0.083	0.048	0.008	0.274	0.136	0.057	0.046	0.248
LL	678.971	525.441	957.697	967.339	685.337	557.137	990.653	966.147	667.206	515.414	973.248	969.627
BIC	-3.552	-2.688	-5.114	-5.249	-3.620	-2.876	-5.247	-5.226	-3.520	-2.645	-5.097	-5.189

The reported mean coefficients of the regime switching model in Table 5.4 indicate that the realized hedge ratio is higher in regime 1 than in regime 2. In the case of FTSE 100 series, the maximum difference between the mean coefficients in the two regimes reaches 0.4. For all assets, the estimates of the transition probabilities imply that there is high persistence in the first regime and low persistence in the second regime. The estimated coefficients from the Heterogeneous Autoregressive (HAR) model reported in

Table 5.5, are not significant in almost all cases (except for the GBP/USD realized hedge ratio series), suggesting a poor in-sample performance of the specific model.

Table 5.4: In-Sample Evidence from the Regime Switching Model

	S&P 500	FTSE 100	EUR/USD	GBP/USD
Regime 1				
C	0.958 ***	0.909 ***	0.979 ***	0.974 ***
φ_1	0.054	0.184 ***	0.108 ***	0.350 ***
φ_2				0.199 ***
Regime 2				
C	0.824 ***	0.513 ***	0.940 ***	0.923 ***
φ_1	-0.229	0.880	-2.865 ***	0.709
φ_2				-0.380
Common				
M_3			-0.003	
M_4			-0.009 ***	
M_6			-0.005 **	
M_7			-0.010 ***	
M_8			-0.008 ***	
M_{10}			-0.003	
M_{12}	-0.028 ***		-0.012 ***	
Mon			-0.001	
Tue		0.015 ***		
Wed		0.019 ***		
Thu		0.019 ***		
Roll D	-0.051 ***			
Log(Sigma)	-3.389 ***	-3.079 ***	-4.266 ***	-4.140 ***
Transition Matrix Parameters				
P11-C	3.703 ***	4.515 ***	4.344 ***	-3.532 ***
P21-C	1.505 **	8.682	1.307	-7.530 ***
$Adj.R^2$	0.131	0.049	0.028	0.242
LL	687.376	597.207	1025.185	960.603
BIC	-3.591	-3.043	-5.345	-5.104

Table 5.5: In-Sample Evidence from the HAR Model

	S&P 500	FTSE 100	EUR/USD	GBP/USD
α_0	0.802***	0.495***	0.023	0.149**
$\alpha_{(d)}$	-0.014	0.045	0.066	0.179***
$\alpha_{(w)}$	0.046	0.116	0.352	-0.033
$\alpha_{(m)}$	0.131	0.299		0.702***
M_3			-0.004	
M_4			-0.005**	
M_6			-0.005	
M_7			-0.004**	
M_8			-0.006	
M_{10}			0.000	
M_{12}	-0.062***		-0.011	
Mon			-0.002	
Tue		0.011**		
Wed		0.009		
Thu		-0.026**		
Roll D	-0.032**			
Adj.R^2	0.131	0.049	0.028	0.242
LL	617.584	476.255	906.792	913.734
BIC	-3.468	-2.597	-5.054	-5.229

5.5 Out-Of-Sample Estimation Results

In this Section, I present the out-of-sample forecast evaluation results for the models described in Section 2, in terms of statistical and economic significance. Each of the competing models is estimated using a rolling window of observations (starting with the in-sample period from 01/01/2009 to 30/06/2009) and one-day-ahead forecasts of the hedge ratio for the out-of-sample period are generated. At each iteration, the specification chosen for each model is maintained but the model parameters are re-estimated.

Combination forecasts with time-varying weights have been formulated under four alternative approaches described at Section 4.3. Across all assets, the obtained forecasts from equation (15), estimated with a constant term and under the restriction of positive weights to sum to unity (namely method D), are found to minimize the RMSE and are, thereafter, chosen to compete with the other specifications. For the rest of the

Chapter, the results and the forecasting performance of the weighted combination forecast from method D will be referred as the time-varying combination forecast.

5.5.1 Statistical Evaluation

An out-of-sample statistical evaluation of the ARMA, ARMA-GARCH, ARFIMA, regime switching, HAR and combination forecast models is presented in Table 5.6. The statistical evaluation criteria used to assess the correct magnitude prediction of the models are the Root Mean Squared Error (RMSE), the Mean Absolute Error (MAE) and the McCracken test for equal RMSE in nested models. The results from the random walk model are reported in the last column for comparison purposes. In terms of the RMSE, the best performing model is the time-varying combination forecast for the two stock indices and the HAR model for the two exchange rates. For the S&P 500 index and the two exchange rates, the lowest MAE is attained by the HAR forecast while, for the FTSE 100 index, by the time-varying combination forecast. I further evaluate the predictability of the RMVHR series by comparing the forecasting accuracy of each model to the benchmark model of random walk. To this end, and taking into consideration that all abovementioned models nest the random walk, the McCracken (2007) test is employed (analysis of the test is provided in Section 4.5.1). The estimated McCracken test statistics suggest that all models produce smaller forecasting errors compared to the random walk. The rejection of the null hypothesis of equal forecasting accuracy for all models across all assets suggests the existence of predictable patterns in the evolution of the RMVHR series.

Next, all the forecasts obtained from the alternative models are jointly compared against the forecasted values obtained from the random walk model. To this end, the test proposed by White (2000) test is used. Let f_t^j indicate loss function, defined as $f_t^j = [g(e_{0t}) - g(e_{jt})]$, where $g(e_{jt})$ is the loss function, RMSE, of obtained forecasts from the j^{th} model described in Section 2, and $g(e_{0t})$ is the loss function for the benchmark model of random walk. Under the null hypothesis, no model exhibits superior forecasting performance over the benchmark model. The White test-statistic is then defined as:

$$\overline{V} \equiv \max_{j=1,\dots,n} \left\{ T^{1/2} \overline{f_t^{\,j}} \right\} \tag{5.16}$$

where $\overline{f_t^{\,j}} = \sum_{t=1}^{T} \frac{1}{T} f_t^{\,j}$ and j=1,…,n are the competing forecasting models. Following the implementation procedure described thoroughly in White (2000), the obtained p-values for the S&P 500, the FTSE 100 and the GBP/USD are zero while for the EUR/USD is equal to 0.088, suggesting that the forecasting ability of the benchmark model of random walk is inferior to any other model employed.

The correct prediction of the directional change of the RMVHR is of crucial importance for the asset allocation decision and the adjustment of the existing position of an investor. To this end, the directional forecasting accuracy of the employed models is examined using the Mean Correct Prediction (MCP) measure and market-timing test (PT test) suggested by Timmerman and Pesaran (1992), analyzed explicitly in Section 4.5.1. The results of the PT test, reported in Table 5.6, suggest that the null hypothesis is rejected in almost all cases, i.e. there is a predictable pattern in the direction of changes in the realized hedge ratio series for all assets.

To complete the statistical evaluation of the competing models a pairwise comparison based on the Modified Diebold-Mariano test is conducted. In specific, the null hypothesis of equal forecasting accuracy is compared against the alternative that the forecasts from model i performs better than model j forecasts. Further description of the test statistic is provided in Section 4.5.1. Table 5.7 reports the t-statistics for the null hypothesis that, in terms of RMSE, the model in row *i* performs equally well with model in column *j* against the alternative that the model in rows *i* outperforms the model in column *j*. In most of the cases combination forecasts outperform ARMA, ARFIMA and Regime Switching models. Interestingly, in almost all cases, all models outperform the Regime Switching model while none of the models outperforms the time-varying combination forecast.

Table 5.6: Evaluation of the Out-Of-Sample Performance: Statistical Measures

	ARMA	**ARMA-GARCH**	**ARFIMA**	**RS**	**HAR**	**Equal**	**Schwarz**	**Time-varying**	**Random Walk**
S&P 500									
MAE	4.321%	4.236%	4.175%	4.424%	**4.143%**	4.149%	4.150%	4.146%	5.567%
RMSE	6.428%	6.374%	6.474%	7.184%	6.312%	6.348%	6.351%	**6.233%**	8.409%
MCP	70.665%	71.475%	72.609%	71.961%	72.447%	72.771%	**73.420%**	72.771%	
McCracken test	439.596 ***	457.446 ***	424.495 ***	228.573 ***	478.669 ***	466.487 ***	465.432 ***	506.801 ***	
PT test	-1.730 *	-2.794 ***	-3.102 ***	5.040 ***	-7.044 ***	-3.025 ***	-2.785 ***	-3.102 ***	
FTSE 100									
MAE	3.664%	3.671%	3.812%	4.050%	3.585%	3.650%	3.657%	**3.556%**	4.604%
RMSE	5.460%	5.428%	5.506%	5.811%	5.416%	5.395%	5.401%	**5.383%**	7.105%
MCP	**69.692%**	68.882%	69.206%	68.071%	68.071%	68.395%	69.044%	68.558%	
McCracken test	428.632 ***	441.012 ***	411.280 ***	305.949 ***	445.884 ***	453.819 ***	451.615 ***	458.866 ***	
PT test	-0.921	-2.044 **	-6.159 ***	6.081 ***	-6.157 ***	-2.450 **	-1.964 **	-3.500 ***	
EUR/USD									
MAE	1.595%	1.689%	1.625%	1.812%	**1.564%**	1.567%	1.573%	1.566%	1.963%
RMSE	2.495%	2.669%	2.434%	3.179%	**2.254%**	2.407%	2.422%	2.327%	2.800%
MCP	71.289%	70.962%	70.636%	71.941%	71.452%	72.431%	71.778%	**72.594%**	
McCracken test	159.643 ***	61.758 ***	198.573 ***	-137.478 ***	333.599 ***	217.048 ***	206.764 ***	275.409 ***	
PT test	-3.774 ***	-4.216 ***	-5.740 ***	3.289 ***	-7.341 ***	-3.552 ***	-3.705 ***	-5.263 ***	
GBP/USD									
MAE	1.472%	1.454%	1.472%	1.563%	**1.448%**	1.450%	1.450%	1.469%	1.897%
RMSE	1.986%	1.967%	1.972%	2.127%	**1.955%**	1.969%	1.969%	1.968%	2.489%
MCP	71.452%	71.615%	71.778%	**73.246%**	71.941%	71.452%	72.920%	71.452%	
McCracken test	350.789 ***	369.546 ***	364.575 ***	226.905 ***	381.013 ***	367.098 ***	367.670 ***	368.489 ***	
PT test	-4.644 ***	-5.937 ***	-6.101 ***	6.907 ***	-7.067 ***	-5.129 ***	-5.209 ***	-5.694 ***	

Note. The Table reports the Root Mean Square Error (RMSE), the Mean Absolute Error (MAE) and the Mean Correct Prediction (MCP) in percentage terms for all econometric specifications. The out-of-sample window extends from July 1, 2010 to December 31, 2012. The McCracken statistic for the RMSE tests for the null hypothesis that the model produces equal forecasting accuracy with the random walk model against the alternative of a better forecasting accuracy of the model. The critical values of the test are provided in McCracken (2007). Under the null hypothesis of the PT test, the actual and forecasted values are independent. Bold values indicate the best performing model for each asset and each criterion. One, two and three asterisks indicate rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

Table 5.7: Pairwise Modified Diebold-Mariano Tests

S&P 500	ARMA	ARMA-GARCH	ARFIMA	RS	HAR	Equal	Schwarz	Time-varying
ARMA		0.954	-0.286	-1.076	0.803	0.702	0.672	1.994
ARMA-GARCH	-0.954		-0.681	-1.169	0.486	0.295	0.259	1.652
ARFIMA	0.286	0.681		-1.014	1.846	1.098	1.053	2.245
RS	1.076	1.169	1.014		1.212	1.319	1.318	1.338
HAR	-0.803	-0.486	-1.846 **	-1.212		-0.315	-0.336	0.924
Equal	-0.702	-0.295	-1.098	-1.319 *	0.315		-1.235	1.291
Schwarz	-0.672	-0.259	-1.053	-1.318 *	0.336	1.235		1.299
Time-varying	-1.994 **	-1.652 **	-2.245 **	-1.338 *	-0.924	-1.291 *	-1.299 *	

FTSE 100	ARMA	ARMA-GARCH	ARFIMA	RS	HAR	Equal	Schwarz	Time-varying
ARMA		0.735	-0.766	-3.489 ***	0.785	1.620	1.437	1.786
ARMA-GARCH	-0.735		-1.455 *	-4.202 ***	0.273	1.144	0.927	1.187
ARFIMA	0.766	1.455		-2.754 ***	1.841	2.438	2.209	2.513
RS	3.489	4.202	2.754		4.208	5.228	5.326	4.204
HAR	-0.785	-0.273	-1.841 **	-4.208 ***		0.622	0.429	0.933
Equal	-1.620 *	-1.144	-2.438 ***	-5.228 ***	-0.622		-1.922 **	0.363
Schwarz	-1.437 *	-0.927	-2.209 **	-5.326 ***	-0.429	1.922		0.496
Time-varying	-1.786 **	-1.187	-2.513 ***	-4.204 ***	-0.933	-0.363	-0.496	

EUR/USD	ARMA	ARMA-GARCH	ARFIMA	RS	HAR	Equal	Schwarz	Time-varying
ARMA		-4.703 ***	0.836	-2.489 ***	2.134	2.392	2.016	1.496
ARMA-GARCH	4.703		2.843	-1.962 **	3.089	5.059	4.991	2.624
ARFIMA	-0.836	-2.843 ***		-2.580 ***	2.146	0.568	0.243	1.159
RS	2.489	1.962	2.580		2.943	2.811	2.789	3.216
HAR	-2.134 **	-3.089 ***	-2.146 **	-2.943 ***		-1.599 *	-1.708 **	-1.032
Equal	-2.392 ***	-5.059 ***	-0.568	-2.811 ***	1.599		-3.926 ***	0.832
Schwarz	-2.016 **	-4.991 ***	-0.243	-2.789 ***	1.708	3.926		0.976
Time-varying	-1.496 *	-2.624 ***	-1.159	-3.216 ***	1.032	-0.832	-0.976	

GBP/USD	ARMA	ARMA-GARCH	ARFIMA	RS	HAR	Equal	Schwarz	Time-varying
ARMA		2.291	1.155	-4.190 ***	2.484	1.890	2.004	1.386
ARMA-GARCH	-2.291 **		-0.449	-4.671 ***	1.154	-0.300	-0.237	-0.107
ARFIMA	-1.155	0.449		-4.222 ***	1.812	0.230	0.288	0.362
RS	4.190	4.671	4.222		4.749	5.573	5.535	4.146
HAR	-2.484 ***	-1.154	-1.812 **	-4.749 ***		-1.352 *	-1.328 *	-1.113
Equal	-1.890 **	0.300	-0.230	-5.573 ***	1.352		1.775	0.105
Schwarz	-2.004 **	0.237	-0.288	-5.535 ***	1.328	-1.775 **		0.063
Time-varying	-1.386 *	0.107	-0.362	-4.146 ***	1.113	-0.105	-0.063	

Note: The Table presents the test statistics for the Modified Diebold-Mariano test. The null hypothesis that the model in rows performs equally well with the model in columns is tested against the alternative of a better forecasting accuracy of the model in rows. One, two and three asterisks indicate rejection of the null hypothesis at 10%, 5% and 1% significance level, respectively.

5.5.2 Economic Evaluation

Having established the existence of predictability in the evolution of the RMVHR series through the means of statistical measures, I now turn to evaluate the forecasting performance of the competing models in economic terms. For each day of the out-of-sample period, the investor rebalances the position according to the forecasted hedge ratio. The realized variance of the constructed portfolios, $\sigma^2_{p,t}$, is based on the intraday data so as to be directly related with the realized hedge ratio methodology (similar to Lai & Sheu, 2008; Sheu & Lai, 2013), thus defined as follows:

$$\sigma^2_{p,t} = \sigma^2_{s,t} - 2\hat{\beta}_t \sigma_{sf,t} + \hat{\beta}^2_t \sigma^2_{f,t} \tag{5.17}$$

where σ^2_s, σ^2_f and σ_{sf} are the sum and the cross product of squared intraday five minutes returns, respectively, and $\hat{\beta}_t$ is the forecasted value of the RMVHR at time *t* with the information set of *t-1*.

For comparison purposes, I also evaluate the performance of hedge ratios derived by models that use daily returns, namely the naïve hedge ratio equal to 1, the static and the rolling OLS method, the Error Correction Model and the Dynamic Conditional Correlation (DCC) GARCH model. For the rest of the Chapter, the hedge ratios derived from the above-mentioned models will be referred as conventional hedge ratios.

The widely used for the estimation of hedge ratio Error Correction Model (ECM) takes into account the cointegrating relationship of the spot and futures prices. Following Juhl et al. (2012), the standard ECM approach is augmented with a lagged error correction term ($\varepsilon_{1,t-1}$) along and lagged values of both spot, $r_{s,t}$, and futures, $r_{f,t}$, returns, where the number of lags is selected to minimize the Schwartz criterion. In specific,

$$r_{s,t} = \alpha_0 + \alpha_1 \varepsilon_{1,t-1} + b_{ECM} r_{f,t} + \sum_{i=1}^{I} \gamma_i r_{f,t-i} + \sum_{j=1}^{J} \delta_j r_{s,t-j} + \varepsilon_t \tag{5.18}$$

Furthermore, the DCC-GARCH(1,1) specification is defined as follows:

$$r_{s,t} = \mu_s + H_{s,t}^{1/2}\varepsilon_{s,t}$$
$$r_{f,t} = \mu_f + H_{f,t}^{1/2}\varepsilon_{f,t} \quad (5.19)$$
$$\begin{pmatrix} \varepsilon_{s,t} \\ \varepsilon_{f,t} \end{pmatrix} | \Phi_{t-1} \sim N(0, H_t)$$

while the conditional variance H_t is equal to

$$H_t = \begin{pmatrix} h_{s,t} & h_{sf,t} \\ h_{sf,t} & h_{f,t} \end{pmatrix} = \begin{pmatrix} h_{s,t}^{1/2} & 0 \\ 0 & h_{f,t}^{1/2} \end{pmatrix} \cdot \begin{pmatrix} 1 & q_{sf,t} \\ q_{sf,t} & 1 \end{pmatrix} \cdot \begin{pmatrix} h_{s,t}^{1/2} & 0 \\ 0 & h_{f,t}^{1/2} \end{pmatrix} = D_t R_t D_t$$

where D_t is the diagonal matrix of time-varying standard deviations from univariate GARCH processes and R_t is the conditional correlation matrix of the standardized disturbances ε_t. The correlation structure is identical to a univariate GARCH model with the following dynamics:

$$\rho_t = (1-\alpha-\beta)\bar{\rho} + \alpha\varepsilon_{t-1}\varepsilon'_{t-1} + \beta\rho_{t-1} \quad (5.20)$$

where $\bar{\rho}$ is the unconditional correlation of the spot and the futures returns and α and β are scalars with $\alpha+\beta<1$.

The hedging effectiveness of each method is assessed by the percentage of risk reduction attained by the hedged portfolio compared to the unhedged portfolio and is defined as follows:

$$HE = 1 - \sqrt{\frac{\sigma_{UP,t}^2}{\sigma_{HP,t}^2}} \quad (5.21)$$

where $\sigma_{UP,t}^2$ is the realized variance of the unhedged portfolio, i.e. the variance of the spot position, and $\sigma_{HP,t}^2$ is the realized variance of the hedged portfolio returns.

Results, reported in Table 5.8, indicate that, in terms of hedging effectiveness, the time-varying combination and the HAR forecast of the RMVHR achieve the highest improvement compared to the unhedged positions for the two stock indices and the two currencies, respectively. Overall, the hedging models based on intraday returns attain a higher risk reduction compared to the hedging models based on daily data. The improvement is within the range of 0.1% and 0.6% and is higher in the case of stock indices than foreign exchange rates.

The performance of the competing models is further evaluated in terms of the Sharpe ratio, the Value-at-Risk (VaR) and the Expected Shortfall (ES). The Sharpe ratio is a typical measure of portfolio performance evaluation expressing the average return per unit of risk. The maximum Sharpe ratio is attained by the HAR model for the S&P 500, the ARFIMA model for the FTSE 100, the ARMA-GARCH for the EUR/USD and the regime switching model for the GBP/USD. Coinciding with results from the hedge effectiveness criterion, the comparative improvement of the RMVHR models, in term of the Sharpe ratio, is more pronounced in the case of the stock indices. The highest Sharpe ratio for the conventional hedge ratios is achieved by the DCC model for the two stock indices, being approximately equal to 0.25%, while the highest Sharpe ratio achieved by the realized hedge ratio models is significantly higher and reaches 1.3%.

Value-at-Risk (VaR) is defined as the portfolio loss over a specific time horizon that is not expected to be exceeded with probability α and is calculated as $VaR(\alpha) = -F^{-1}(r_{p,t+1};\alpha)$ where F^{-1} is the inverse of the empirical cumulative distribution function of portfolio returns. The standard normality assumption is adopted and the variance-covariance method is used to compute VaR. An alternative risk measure that satisfies the axiom of a coherent risk measure is Expected Shortfall (ES). The ES measures the expected loss given that the portfolio returns have exceeded the α^{th} empirical percentile i.e. $ES(\alpha) = -F^{-1}(r_{p,t+1} \mid r_{p,t+1} \leq VaR(\alpha))$[7]. Results indicate that modeling directly the RMVHR outperforms conventional daily-return based methods in managing portfolio risk. In addition, the HAR model and the time-varying combination forecasts of realized hedge ratio are the best performing models that minimize the VaR and ES.

[7] Table 5.8 reports the VaR and ES for 95% confidence level. For 99% confidence level, the empirical findings remained unchanged.

To explore further the economic significance of the results, the forecasting models are compared in terms of a mean-variance utility function. I assume that an investor faces the mean-variance expected utility function,

$$EU(r_{p,t+1};\beta_{t+1,\gamma}) = E(r_{p,t+1};\hat{\beta}_{t+1}) - \gamma Var(r_{p,t+1};\hat{\beta}_{t+1}) \tag{5.22}$$

where $r_{p,t+1}$ is the portfolio return and γ is the degree of risk aversion. Based on the relative literature (Sheu & Lai, 2013; Park & Switzer, 1995), different levels of risk aversion, γ=1,3,4,7,10,15,20 are considered. Results are similar for the different levels of risk aversion, thus, for brevity, Table 5.8 reports the average utility levels only for γ=3. First, in almost all cases (except from the regime switching model for S&P 500 and EUR/USD) the models based on intraday returns substantially outperform the conventional models based on daily returns. Moreover, the net utility benefit from hedging with the best model based on intraday returns compared to the best model based on daily returns ranges from 0.002 to 0.84. These results are in line with empirical evidence of Sheu and Lai (2013) and Lai and Sheu (2010); thus supporting and enhancing the findings that exploiting intraday price information can increase hedging performance and utility gains. Results are similar across different degrees of risk aversion implying that the gains from forecasting directly the realized hedge ratio do not differentiate across hedgers with different risk aversions. Moreover, there is no single winner across assets; however, one notable finding is that combination forecasts are always outperformed by at least one of the other models employed to the RMVHR series. For the two stock indices and the GBP/USD foreign exchange rate, the best realized hedge ratio forecast attain positive economic gains, whereas the economic gains of the daily return based methods are negative for all models and assets. To conclude, empirical evidence indicate that forecasting directly the realized hedge ratio based on intraday data improves hedging portfolio performance in terms of risk reduction, managing portfolio risk, Sharpe ratio and markedly, in terms of economic gains.

Table 5.8. Out-Of-Sample Hedging Performance Net of Transaction Costs: Economic Evaluation

	Conventional Hedge Ratio					RMVHR							
	Naïve	**Static OLS**	**Rolling OLS**	**ECM**	**DCC**	**ARMA**	**ARMA-GARCH**	**ARFIMA**	**RS**	**HAR**	**Equal**	**Schwarz**	**Time-varying**
S&P 500													
HE	89.923%	89.955%	89.955%	89.933%	89.912%	90.313%	90.327%	90.334%	90.240%	90.334%	90.332%	90.331%	**90.334%**
Sharpe Ratio	-0.064%	-0.011%	-0.021%	-0.026%	0.244%	0.992%	0.843%	1.209%	0.087%	**1.337%**	0.894%	0.886%	0.983%
VaR	0.372%	0.372%	0.371%	0.372%	0.372%	0.365%	0.365%	0.365%	0.366%	**0.365%**	0.365%	0.365%	0.365%
ES	0.467%	0.466%	0.466%	0.466%	0.467%	0.458%	0.458%	0.457%	0.459%	**0.457%**	0.457%	0.457%	0.457%
Expected Utility	-0.201	-0.187	-0.190	-0.191	-0.124	0.064	0.028	0.117	-0.158	**0.148**	0.040	0.038	0.062
FTSE 100													
HE	90.220%	90.650%	90.480%	90.497%	90.307%	90.988%	90.992%	90.970%	90.953%	90.992%	90.993%	90.993%	**90.994%**
Sharpe Ratio	-0.084%	0.262%	0.298%	0.211%	0.249%	0.727%	0.843%	**1.155%**	0.822%	0.899%	0.890%	0.885%	0.900%
VaR	0.361%	0.352%	0.354%	0.355%	0.360%	0.346%	0.346%	0.346%	0.347%	0.345%	0.346%	0.346%	**0.345%**
ES	0.453%	0.441%	0.444%	0.445%	0.451%	0.434%	0.433%	0.434%	0.435%	0.433%	0.433%	0.433%	**0.433%**
Expected Utility	-0.216	-0.121	-0.115	-0.137	-0.130	-0.002	0.027	**0.103**	0.021	0.040	0.038	0.037	0.041
EUR/USD													
HE	96.385%	96.375%	96.392%	96.397%	96.382%	96.479%	96.473%	96.484%	96.444%	**96.486%**	96.483%	96.482%	96.484%
Sharpe Ratio	0.035%	0.019%	-0.033%	-0.024%	-0.139%	0.227%	**0.247%**	0.062%	-0.327%	-0.062%	0.029%	0.024%	-0.114%
VaR	0.196%	0.196%	0.196%	0.196%	0.196%	0.193%	0.194%	0.193%	0.194%	**0.193%**	0.193%	0.193%	0.193%
ES	0.246%	0.246%	0.245%	0.245%	0.246%	0.243%	0.243%	0.243%	0.243%	**0.243%**	0.243%	0.243%	0.243%
Expected Utility	-0.042	-0.044	-0.050	-0.049	-0.064	-0.017	**-0.015**	-0.037	-0.086	-0.053	-0.041	-0.042	-0.059
GBP/USD													
HE	95.868%	95.824%	95.841%	95.845%	95.788%	95.937%	95.938%	95.937%	95.932%	**95.938%**	95.938%	95.938%	95.937%
Sharpe Ratio	-0.050%	-0.134%	-0.385%	-0.135%	-0.550%	0.071%	0.154%	0.155%	**0.339%**	0.164%	0.176%	0.174%	0.141%
VaR	0.174%	0.175%	0.175%	0.175%	0.176%	0.173%	0.173%	0.173%	0.173%	**0.173%**	0.173%	0.173%	0.173%
ES	0.219%	0.220%	0.219%	0.219%	0.221%	0.217%	0.217%	0.217%	0.217%	**0.217%**	0.217%	0.217%	0.217%
Expected Utility	-0.041	-0.051	-0.078	-0.051	-0.098	-0.028	-0.019	-0.019	**0.001**	-0.018	-0.016	-0.016	-0.020

Note: The Table reports the Hedge Effectiveness (HE), the Sharpe Ratio, the Value-at-Risk (VaR), the Expected Shortfall (ES), both at 95% confidence level, and the expected utility levels (γ=3). Bold values indicate the best performing model for each asset and each criterion.

5.6 Sampling Frequency

The dynamic properties of realized volatilities and covariances are likely to be influenced by the sampling frequency employed (see Lai & Sheu 2008, 2010). To this end, I study the statistical and economic performance of various hedging models against various sampling frequencies. Results from the statistical and economic criteria are reported in Table 5.9 and Table 5.10.

Obtained results from the statistical measures lead to a number of conclusions. Firstly, the ordering of the models based on the RMSE and MAE changes as a function of the sampling frequency. Second, the sampling frequency influences the magnitude of the MAE and the RMSE; however, an overall conclusion across assets and models on whether the two measures increase or decrease as the sampling frequency increases cannot be reached. The rejection of the null hypothesis of equal forecasting accuracy, based on the estimated McCracken test statistics, remains valid across the various sampling frequencies suggesting that all models produce smaller forecasting errors compared to the random walk.

Turning to the economic loss functions described thoroughly in Section 5, results from the 15-minute and 30-minute returns indicate that, with only a few exceptions, the models based on intraday data yield better performance when compared to models based on daily data across all frequencies, consistent with the results for the 5-minute returns. However, the best performing model for each asset varies across the different sampling frequencies. Thereafter, it is concluded that previously reported results from the statistical and economic measures are robust to the sampling frequency used in the derivation of realized variance and covariance.

Table 5.9: Evaluation of the out-of-sample performance for alternative sampling frequencies – Statistical Measures

		ARMA	ARMA - GARCH	ARFIMA	RS	HAR	Equal	Schwarz	Time-varying
S&P500									
15min	MAE	3.787%	3.839%	3.709%	3.707%	3.708%	3.698%	3.702%	**3.661%**
	RMSE	4.994%	5.090%	4.974%	4.970%	4.951%	4.938%	4.942%	**4.885%**
	McCracken test	207.325 ***	191.856 ***	210.656 ***	211.392 ***	214.530 ***	216.766 ***	216.074 ***	225.859 ***
30 min	MAE	3.859%	3.829%	3.819%	3.836%	3.842%	3.816%	3.816%	**3.799%**
	RMSE	5.437%	5.447%	5.416%	5.415%	5.438%	5.410%	5.410%	**5.372%**
	McCracken test	232.076 ***	230.493 ***	235.326 ***	235.552 ***	231.839 ***	236.341 ***	236.305 ***	242.366 ***
FTSE 100									
15min	MAE	3.658%	3.665%	4.214%	3.695%	3.660%	3.664%	3.649%	**3.610%**
	RMSE	**5.395%**	5.421%	5.851%	5.480%	5.482%	5.414%	5.406%	5.442%
	McCracken test	188.593 ***	184.754 ***	125.732 ***	176.055 ***	175.756 ***	185.761 ***	187.008 ***	181.616 ***
30 min	MAE	3.612%	3.527%	3.511%	3.614%	**3.505%**	3.524%	3.528%	3.585%
	RMSE	5.117%	5.038%	5.036%	5.295%	**5.018%**	5.043%	5.046%	5.100%
	McCracken test	232.415 ***	245.779 ***	246.153 ***	203.812 ***	249.216 ***	244.888 ***	244.294 ***	235.279 ***
EUR/USD									
15min	MAE	1.539%	1.448%	1.485%	1.527%	1.439%	1.435%	**1.433%**	1.458%
	RMSE	2.724%	2.612%	2.663%	3.155%	**2.579%**	2.619%	2.631%	2.608%
	McCracken test	131.139 ***	162.953 ***	148.329 ***	29.357 ***	173.175 ***	160.862 ***	157.529 ***	164.210 ***
30 min	MAE	1.478%	1.571%	1.505%	1.526%	1.458%	**1.456%**	1.456%	1.485%
	RMSE	2.362%	2.555%	2.412%	2.555%	**2.345%**	2.361%	2.367%	2.482%
	McCracken test	197.294 ***	136.048 ***	180.405 ***	136.068 ***	203.263 ***	197.710 ***	195.725 ***	158.209 ***
GBP/USD									
15min	MAE	1.456%	1.443%	3.933%	1.455%	1.443%	1.434%	1.434%	**1.433%**
	RMSE	2.049%	2.030%	10.701%	2.042%	2.037%	2.023%	2.024%	**2.020%**
	McCracken test	214.543 ***	222.536 ***	-455.335 ***	217.566 ***	219.591 ***	225.124 ***	224.871 ***	226.531 ***
30 min	MAE	4.341%	4.243%	**1.910%**	5.352%	4.443%	4.257%	4.396%	4.363%
	RMSE	10.266%	10.245%	**2.437%**	12.694%	10.388%	10.352%	10.515%	10.293%
	McCracken test	274.657 ***	276.467 ***	3130.217 ***	104.724 ***	264.206 ***	267.293 ***	253.663 ***	272.362 ***

Note: Bold values indicate the best performing model for each asset and each criterion.

Table 5.10: Evaluation of the out-of-sample performance for alternative sampling frequencies – Economic Measures

		Conventional Hedge Ratios					Realized Hedge Ratio							
		Naïve	Static OLS	Rolling OLS	ECM	DCC	ARMA	ARMA-GARCH	ARFIMA	RS	HAR	Equal	Schwarz	Time-varying
S&P500														
15 min	HE	96.215%	96.230%	96.229%	96.220%	96.191%	96.305%	96.303%	96.311%	96.309%	**96.314%**	96.312%	96.312%	96.313%
	S.R.	-0.105%	-0.018%	-0.036%	-0.043%	0.403%	0.985%	1.196%	0.878%	1.261%	**1.429%**	1.151%	1.177%	1.174%
	VaR	0.221%	0.220%	0.220%	0.221%	0.221%	0.218%	0.218%	0.218%	0.218%	**0.218%**	0.218%	0.218%	0.218%
	ES	0.277%	0.276%	0.276%	0.277%	0.278%	0.273%	0.274%	0.273%	0.273%	**0.273%**	0.273%	0.273%	0.273%
	Utility	-8.325%	-6.986%	-7.254%	-7.380%	-0.724%	8.009%	11.136%	6.428%	12.099%	**14.583%**	10.458%	10.845%	10.811%
FTSE 100														
15 min	HE	95.846%	96.038%	95.976%	95.982%	95.872%	96.070%	**96.072%**	96.012%	96.061%	96.063%	96.067%	96.069%	96.064%
	S.R.	-0.130%	0.409%	0.466%	0.329%	0.388%	0.285%	0.234%	**1.453%**	0.711%	0.614%	0.662%	0.610%	0.393%
	VaR	0.230%	0.224%	0.226%	0.226%	0.230%	0.224%	0.224%	0.226%	0.224%	0.224%	0.224%	0.224%	**0.224%**
	ES	0.289%	0.281%	0.283%	0.283%	0.288%	0.281%	0.281%	0.284%	0.281%	0.281%	0.281%	0.281%	**0.281%**
	Utility	-10.160%	-1.124%	-0.281%	-2.472%	-1.645%	-3.061%	-3.862%	**15.584%**	3.717%	2.163%	2.931%	2.118%	-1.344%
30 min	HE	97.443%	97.563%	97.525%	97.534%	97.435%	97.559%	97.563%	97.563%	97.543%	**97.564%**	97.564%	97.563%	97.555%
	S.R.	-0.172%	0.540%	0.615%	0.435%	0.510%	1.020%	0.967%	0.615%	**1.067%**	0.796%	0.894%	0.917%	0.684%
	VaR	0.169%	**0.164%**	0.165%	0.165%	0.169%	0.165%	0.164%	0.164%	0.165%	0.164%	0.164%	0.164%	0.165%
	ES	0.212%	**0.206%**	0.207%	0.207%	0.212%	0.206%	0.206%	0.206%	0.207%	0.206%	0.206%	0.206%	0.206%
	Utility	-6.739%	2.140%	3.036%	0.849%	1.709%	7.953%	7.303%	3.050%	**8.535%**	5.235%	6.425%	6.701%	3.881%
EUR/USD														
15 min	HE	98.428%	98.424%	98.431%	98.434%	98.425%	98.444%	**98.450%**	98.447%	98.430%	98.450%	98.449%	98.448%	98.449%
	S.R.	0.054%	0.029%	-0.051%	-0.036%	-0.215%	-0.191%	**0.233%**	0.095%	0.199%	-0.034%	0.061%	0.060%	0.035%
	VaR	0.124%	0.124%	0.123%	0.123%	0.124%	0.123%	**0.123%**	0.123%	0.123%	0.123%	0.123%	0.123%	0.123%
	ES	0.155%	0.155%	0.155%	0.155%	0.155%	0.154%	**0.154%**	0.154%	0.154%	0.154%	0.154%	0.154%	0.154%
	Utility	-1.510%	-1.716%	-2.356%	-2.230%	-3.681%	-3.459%	**-0.053%**	-1.162%	-0.339%	-2.193%	-1.435%	-1.437%	-1.643%
30 min	HE	98.947%	98.944%	98.949%	98.951%	98.945%	98.956%	98.950%	98.955%	98.950%	**98.957%**	98.956%	98.956%	98.953%
	S.R.	0.066%	0.036%	-0.063%	-0.044%	-0.263%	-0.229%	-0.098%	**0.301%**	0.015%	0.000%	-0.002%	-0.042%	0.177%
	VaR	0.098%	0.099%	0.098%	0.098%	0.099%	0.098%	0.098%	0.098%	**0.098%**	0.098%	0.098%	0.098%	0.098%
	ES	0.123%	0.124%	0.123%	0.123%	0.124%	0.123%	0.123%	0.123%	**0.123%**	0.123%	0.123%	0.123%	0.123%
	Utility	-0.859%	-1.064%	-1.707%	-1.582%	-3.029%	-2.782%	-1.933%	**0.685%**	-1.191%	-1.282%	-1.297%	-1.558%	-0.131%

GBP/USD														
15 min	HE	98.121%	98.092%	98.102%	98.106%	98.052%	98.143%	98.143%	98.143%	98.143%	98.143%	98.144%	98.144%	**98.144%**
	S.R.	-0.075%	-0.202%	-0.581%	-0.838%	-0.827%	0.230%	**0.265%**	0.227%	0.235%	0.259%	0.243%	0.245%	0.216%
	VaR	0.114%	0.114%	0.114%	0.114%	0.115%	0.113%	0.113%	0.113%	0.113%	0.113%	0.113%	0.113%	**0.113%**
	ES	0.142%	0.143%	0.143%	0.143%	0.145%	0.142%	0.142%	0.142%	0.142%	0.142%	0.142%	0.142%	**0.142%**
	Utility	-2.112%	-3.068%	-5.811%	-3.067%	-7.715%	0.102%	**0.354%**	0.083%	0.140%	0.311%	0.198%	0.214%	0.001%
30 min	HE	93.744%	93.653%	93.705%	93.701%	93.652%	**94.134%**	94.122%	93.980%	93.595%	94.056%	94.072%	94.045%	94.092%
	S.R.	-0.041%	-0.112%	-0.322%	-0.113%	-0.462%	**0.476%**	0.179%	0.233%	0.373%	0.133%	0.282%	0.297%	0.341%
	VaR	0.147%	0.148%	0.147%	0.147%	0.149%	0.146%	**0.146%**	0.146%	0.151%	0.146%	0.146%	0.146%	0.146%
	ES	0.184%	0.185%	0.185%	0.185%	0.186%	0.183%	**0.183%**	0.183%	0.189%	0.183%	0.183%	0.183%	0.183%
	Utility	-5.683%	-6.690%	-9.398%	-6.661%	-11.304%	**1.218%**	-2.554%	-1.960%	-0.317%	-3.187%	-1.283%	-1.097%	-0.521%

Note: The Table reports the hedge effectiveness (HE), the Sharpe Ratio, the Value-at-Risk, the Expected Shortfall (ES), both at 95% confidence level, and the average utility levels (γ=3). Bold values indicate the best performing model for each asset and each criterion.

5.7 Transaction Costs

So far, the evaluation of the hedging performance of the competing models ignores the impact of transaction costs on the profitability of the hedging strategy. All models (except for the naïve and the constant OLS, for which only the cost of opening the position is relevant) require daily rebalancing of the hedged portfolio, thus entailing large transaction cost that could diminish any economic gains. The expected utility for each hedging model assuming an investor with a mean-variance utility is reevaluated taking into consideration transaction costs. Following Lee (2009), I assume transaction cost of 0.02%, which is typical for a round-trip transaction. The aggregate realized utility with transaction cost during the out-of-sample period for each competing model is presented in Table 5.11 assuming a degree of risk aversion equal to three.

Notably, for all four assets, the best performing model is one of the models that forecast directly the RMVHR. Moreover, when transaction costs are taken into consideration, the best performing model, in terms of the highest utility gains, remains unchanged. Models on realized hedge ratio do not outperform all of the conventional daily models only for a limited number of cases e.g. the regime switching model for S&P 500 index. Interestingly, in the case of EUR/USD, none of the dynamic conventional hedging models based on daily data outperforms the constant hedge ratio models, namely the naïve and the constant static OLS, and only the ARMA and ARMA-GARCH models on RMVHR outperform the constant hedge ratio.

Table 5.11: Out-Of-Sample Hedging Performance: Average Daily Utility with Transaction Costs

	Conventional Hedge Ratio					RMVHR							
	Naïve	**Static OLS**	**Rolling OLS**	**ECM**	**DCC**	**ARMA**	**ARMA-GARCH**	**ARFIMA**	**RS**	**HAR**	**Equal**	**Schwarz**	**Time-varying**
S&P 500	-0.204	-0.190	-0.194	-0.195	-0.137	0.026	0.003	0.095	-0.182	**0.129**	0.020	0.018	0.040
FTSE 100	-0.220	-0.124	-0.121	-0.140	-0.162	-0.038	-0.001	**0.086**	-0.007	0.023	0.022	0.021	0.020
EUR/USD	-0.045	-0.047	-0.055	-0.052	-0.076	-0.031	**-0.027**	-0.048	-0.102	-0.062	-0.051	-0.051	-0.069
GBP/USD	-0.045	-0.054	-0.084	-0.055	-0.121	-0.038	-0.028	-0.029	**-0.009**	-0.027	-0.024	-0.024	-0.029

Note. The Table presents the expected utility with transaction cost during the out-of-sample period for the hedge ratio forecasts. The transaction cost is assumed equal to 0.02% for a round-trip transaction while the degree of risk aversion is equal to 3. Bold values indicate the best performing model for each asset.

5.8 Conclusions

In this Chapter, I study the distributional properties of the realized variance, covariance and hedge ratio and find that the long-memory traits of the individual processes is eliminated in the case of the Realized Minimum Variance Hedge Ratio series, calculated from high-frequency data as the ratio of realized spot and futures covariance to futures variance. To this end, I propose the direct modelling of the evolution of the series RMVHR under alternative econometric specifications. The economic significance of the obtained forecasts is compared to traditional techniques of estimating and forecasting the optimal hedge ratio such as the OLS or GARCH-based hedge ratio. Data on two equity indices and two foreign exchange rates are used to study any differential pattern across different asset classes.

Several statistical measures suggest the presence of predictable pattern in the evolution of the realized hedge ratio series. The findings suggest that, in terms of correct magnitude prediction, and in specific RMSE, the time-varying combination forecast for the two stock indices and the HAR model for the two exchange rates outperform the competing models. Employed models correctly predict the directional change of the hedge ratio approximately 70% of times, throughout the out-of-sample period.

From an economic perspective, empirical results suggest that forecasting directly the evolution of the realized hedge ratio, calculated from informative intraday returns, results in substantial improvement of hedging performance, in terms of risk reduction, Sharpe ratio, expected utility and Value-at-Risk, while being magnified in the case of stock indices. When transaction costs from daily rebalancing are taken into account or the data frequency is changed, the hedging performance superiority of RMVHR forecasts is maintained for almost all models and assets. The results of the analysis suggest that forecasting the realized hedge ratio directly can provide substantial benefits to risk managers and hedgers.

Chapter 6

Conclusions & Suggestions for Future Research

Estimates and forecasts of correlation are key inputs in a number of financial applications. This thesis contributes to the discussion of correlation modelling and especially to the information content of alternative correlation measures in optimizing portfolio allocation, risk management and hedging decisions. Overall, findings of the thesis show that understanding and modelling the dynamics that govern the evolution of alternative measures of correlation results in optimized portfolio allocation, risk management decisions as well as to substantially enhanced hedging and arbitrage trading strategies.

The purpose of Chapter 2 has been twofold. First, an extensive literature review of advances in correlation modelling and relevant applications is presented. The main reason behind the tremendous evolution of research focusing on the second moments of distribution is that volatility and correlation cannot be directly observed and should thus, be modelled within a specific econometric specification. On an attempt to overcome the ambiguities arising from inherent assumptions of every model, the measures of implied correlation, inferred from option prices, and realized correlation, computed from high-frequency intraday returns, have stimulated the research interest. Several studies have focused on the information content of above-mentioned measures in forecasting the latent process, optimizing asset allocation and forming profitable strategies, thus providing a direct test of the market efficiency hypothesis.

Chapter 3 discusses the statistical properties and the information content of macroeconomic and other market-specific variables in predicting the correlation risk premium, considered as an insurance premium that investors require to hedge against correlation increase. Correlation risk premium is defined as the difference of risk neutral correlation, the realized correlation, and correlation under the physical measure, implied correlation. A model-free measure of implied correlation, inferred from currently observed option prices and adjusted to account for the early exercise premium inherent in the American options, is proposed. Previous studies have provided evidence of a negatively priced correlation risk premium, in the sense that assets that capitalize their gains in high correlation market conditions earn negative excess returns. Based on the decomposition of volatility risk premium of the index into individual stocks' volatility risk premium and inter-asset correlation structure, the significance of the volatility risk premium is also assessed. The statistical analysis is performed on an extensive dataset of fifteen years, while the robustness of the results is also assessed for different sample sizes, thus allowing the examination of the time series behavior during several recent periods of financial distress with scrutiny.

Consistent with previous literature, both implied correlation and realized correlation increase substantially during periods of financial turmoil. Both correlation measures culminated during the Asian (July 1997) and Russian (August 1998) financial crises, the devaluation of local currencies flowing the great depression in Brazil (early 1999) and Argentine (1998 - 2002), the dot-com bubble in 2001, and the global financial crisis of 2007 – 2009. Interestingly, during the subprime mortgage crisis the implied correlation reached a maximum value of 0.97, indicating almost perfect correlation and zero diversification benefit.

The empirical analysis provides evidence in favor of a negatively priced volatility risk premium, both for the index and the constituent individual stocks as well as a negative price of correlation risk premium throughout the sample period. However, during turbulent periods, implied and realized correlations do not differ substantially and the null hypothesis of zero correlation risk premium cannot be rejected. A plausible explanation might stem either from investors' irrationality or from increased transaction costs and margin requirements that eliminated any arbitrage opportunities between the index and the individual options. The statistical analysis suggests that the distribution of the correlation risk premium is leptokurtic and negatively skewed while exhibits strong persistence. Furthermore, index returns, variance risk premium of index options and correlation risk premium are positively correlated suggesting that a decrease in returns will decrease the negative price of correlation risk premium, i.e. CRP will be forced to less negative values, coinciding with the empirical finding of zero CRP during the subprime mortgage crisis. Finally, macroeconomic variables, e.g. the term structure, the default spread structure and the USD Libor, fail to provide accurate forecasts of the CRP. In contrast, the volatility risk premium of the index and the index returns embed significant information content of future CRP.

Chapter 4 of the thesis is concerned with the existence of predictable dynamics in the evolution of the model-free implied correlation measure derived in Chapter 3. To this end, several time-series econometric specifications that gauge alternative dynamics of the series, including combination forecasts, are employed to model and forecast the evolution of the series. The out-of-sample performance of the models is initially assessed with statistical measures that estimate both the magnitude of forecast error and the directional efficiency of obtained forecasted values. Secondly, a trading strategy that exploits daily changes of the implied correlation is employed to assess the economic significance of predictable patterns. Finally, the robustness of obtained results with regard to transaction costs and across different sampling periods is examined.

The empirical results suggest that combination forecasts, which essentially accumulate the obtained forecasts of several econometric models, provide superior forecasting performance, both in terms of correct magnitude forecast and directional prediction. The comparison of forecasting performance with the benchmark model of random walk constitutes a direct test of the efficient market hypothesis. Both on individual and on aggregate level, all models produce superior forecasting accuracy, thus providing strong evidence in favour of predictable patterns and against the market efficiency of the S&P 100 option market. In economic terms, results from a trading strategy designed to exploit daily changes of correlation, suggest that modelling the implied correlation process under specific econometric specifications is able to generate abnormal profits, which however disappear when transaction costs are taken into account. To conclude, although statistical measures favoured the rejection of market efficiency hypothesis, results from the trading strategy suggest that, by virtue of transaction costs, investors cannot attain significant economic profits and thus, the efficiency of the S&P

100 options market cannot be rejected. Obtained results remain robust across different in-sample windows and forecast periods.

Chapter 5 is concerned with the dynamic evolution of the Realized Minimum Variance Hedge Ratio (RMVHR). In this study, four issues are addressed. First, the RMVHR is constructed from intraday returns for two asset classes, namely stock indices and exchange rates, and the distributional properties are compared to the time-series characteristics of the realized volatility and covariance processes. Second, instead of forecasting the variance-covariance matrix of the spot and future returns, a number of time-series models are employed to capture the dynamic evolution of the RMVHR series per se. Third, the predictability of the series is assessed based on statistical measures. Finally, obtained forecasts from a number of econometric models are compared to traditional methods of deriving hedge ratios that use daily data.

The comparative statistical analysis of realized volatility, covariance and hedge ratio suggest that long memory and persistence is less pronounced in the case of the hedge ratio. Moreover, equity markets are far more volatile than the FX markets; a finding that is also present in the realized hedge ratio series. With regard to the predictability of the realized hedge ratio series, empirical evidence suggests the superior forecasting performance of combination forecasts. Interestingly, the forecasting performance of the random walk model is inferior to any other econometric specification employed to forecast the realized hedge ratio, thus providing evidence of predictable dynamics in the evolution of the series and challenging the hypothesis of market efficiency. The out-of-sample hedging effectiveness of the realized hedge ratio is compared to the hedge ratio estimated from daily returns. Specifically, the hedging performance is evaluated with a view to risk management and portfolio optimization applications. To this end, the percentage of portfolio risk reduction, the Sharpe ratio and the expected utility, the Value-at-Risk (VaR) and the Expected Shortfall (ES) measures are employed. Main findings suggest that, the proposed methodology of directly forecasting the RMVHR results in considerable improvement in terms of portfolio risk and expected utility, risk-return tradeoff and Value-at-Risk. The results hold across the different asset classes, although the benefits are lower in the case of exchange rates, which are considerably more liquid. The superior forecasting performance of the realized hedge ratio series remains robust to transactions costs and sampling frequency.

Several research questions arise from this thesis. In Chapter 3, the information content of the most widely spread macroeconomic variables in predicting the correlation risk premium is found to be rather limited. The information content and the anticipated economic significance of alternative macroeconomic and market specific variables in forecasting the correlation risk premium can be further explored and discussed. Moreover, it would be interesting to assess the information content of macroeconomic variables in forecasting long run component of the correlation risk premium. As an alternative course of future research, econometric specification can be applied to examine the presence of predictable dynamics in the evolution of the correlation risk premium.

In Chapter 4, the statistical measures provided evidence of market inefficiency whereas the inclusion of transaction costs in the employed trading strategy eliminates any profitability. Future work could further investigate the profitability of correlation trading strategies using different datasets, alternative model specifications and longer forecast horizons. Alternatively, an interesting exercise would be to assess the profitability of the trading strategy for different levels of transaction costs. Taking into consideration that the implemented trading strategy requires transactions on the index and the 100 constituent stocks, an alternative approach would include the construction a mimicking portfolio that tracks closely the basket of individual options, e.g. including representative liquid stocks of the index (e.g. those representing the 75% of the index capitalization). This strategy could significantly reduce transaction costs and enhance profitability. Additionally, institutional investors typically face smaller transaction costs. It would be thus interesting to examine the profitability for different types of investors.

Finally, in Chapter 5, the hedging effectiveness of the realized minimum variance hedge ratio, is assessed. An interesting extension would be to apply the methodology to alternative hedge ratios that seek to maximize the risk-return function, e.g. M-MEG and M-GSV. Additionally, improving the realized hedge ratio estimates and forecasts by accounting for microstructure noise and non-synchronous trading is an interesting topic for future research.

Bibliography

Ahoniemi, K. (2008). Modeling and Forecasting the VIX Index. Working paper, Helsinki School of Economics and HECER.

Aït-Sahalia, Y., Mykland, P. A., & Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial studies*, *18*, 351-416.

Aksu, C., & Gunter, S. I. (1992). An empirical analysis of the accuracy of SA, OLS, ERLS and NRLS combination forecasts. *International Journal of Forecasting*, *8*, 27-43.

Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts.*Review of Financial studies*, *15*, 1137-1187.

Ang, A., & Chen, J. (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, *63*, 443-494.

Amira, K., Taamouti, A., & Tsafack, G. (2011). What drives international equity correlations? Volatility or market direction?. *Journal of International Money and Finance*, *30*, 1234-1263.

Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International economic review*, *39*, 885-905.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics*, *61*, 43-76.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American statistical association*, *96*, 42-55.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Wu, G. (2006). Realized beta: Persistence and predictability. *Advances in econometrics*, *20*, 1-39.

Andersen, T.G., Bollerslev, T., Diebold, F.X. (2007). Roughing it up: Including Jump components in the measurement, modelling and forecasting of return volatility. *The Review of Economics and Statistics*, 89, 701-720.

Baillie, R. T. (1989). Econometric tests of rationality and market efficiency. *Econometric Reviews*, *8*, 151-186.

Bakshi, G., & Kapadia, N. (2003a). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, *16*, 527-566.

Bakshi, G., & Kapadia, N. (2003b). Volatility risk premiums embedded in individual equity options: Some new insights. *The Journal of Derivatives*, *11*, 45-54.

Bakshi, G., & Madan, D. (2006). A theory of volatility spreads. *Management Science*, *52*, 1945-1956.

Bandi, F. M., & Russell, J. R. (2008). Microstructure noise, realized variance, and optimal sampling. *The Review of Economic Studies*, *75*, 339-369.

Barndorff-Nielsen, O. E., & Shephard, N. (2002). Estimating quadratic variation using realized variance. *Journal of Applied Econometrics*, *17*, 457-477.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., & Shephard, N. (2011). Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics*, *162*, 149-169.

Barone-Adesi, G. & Whaley, R. E. (1987). Efficient analytic approximation of American option values. *The Journal of Finance*, *42*, 301-320.

Bates, D. S. (1991). The crash of'87: Was it expected? The evidence from options markets. *The journal of finance*, *46*, 1009-1044.

Bauwens, L., Laurent, S., & Rombouts, J. V. (2006). Multivariate GARCH models: a survey. *Journal of applied econometrics*, *21*, 79-109.

Becker, R., & Clements, A. E. (2008). Are combination forecasts of S&P 500 volatility statistically superior. *International Journal of Forecasting*, *24*, 122–133.

Beckers, S. (1981). Standard deviations implied in option prices as predictors of future stock price variability. *Journal of Banking & Finance*, *5*, 363-381.

Bekaert, G., Ehrmann, M., Fratzscher, M., & Mehl, A. J. (2011). *Global crises and equity market contagion* (No. w17121). National Bureau of Economic Research.

Bennett, P., & Kelleher, J. (1988). The international transmission of stock prices disruption in October 1987. *Quarterly Review*, (Sum), 17-33.

Bernales, A. and Guidolin, M. (2010), "Can we forecast the implied volatility surface dynamics for CBOE equity options? Predictability and economic value tests", *Working Paper*, University of Manchester.

Bessembinder, H. & Chan, K. (1992). Time varying risk premia and forecastable returns in futures markets. *Journal of Financial Economics*, *32*, 169-193.

Blair, B. J., Poon, S. H., & Taylor, S. J. (2010). Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. In *Handbook of Quantitative Finance and Risk Management* (pp. 1333-1344). Springer US.

Bollen, N. P., & Whaley, R. E. (2004). Does Net Buying Pressure Affect the Shape of Implied Volatility Functions? *The Journal of Finance*, *59*, 711-753.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity.*Journal of econometrics*, *31*, 307-327.

Bollerslev, T., Engle, R. F., & Wooldridge, J. M. (1988). A capital asset pricing model with time-varying covariances. *The Journal of Political Economy*, 116-131.

Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics*, 498-505.

Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, *22*, 4463-4492.

Brenner, R. J., & Kroner, K. F. (1995). Arbitrage, cointegration, and testing the unbiasedness hypothesis in financial markets. *Journal of Financial and Quantitative Analysis*, *30*, 23-42.

Britten-Jones, M., & Neuberger, A. (2000). Option Prices, Implied Price Processes, and Stochastic Volatility. *The Journal of Finance, 55*, 839-866.

Brooks, C., & Oozeer, M.C. (2002). Modelling the implied volatility of options on long gilt futures. *Journal of Business, Finance & Accounting, 29*, 111-137.

Buraschi, A., Porchia, P., & Trojani, F. (2010). Correlation risk and optimal portfolio choice. *The Journal of Finance*, *65*, 393-420.

Buraschi, A., Trojani, F., & Vedolin, A. (2013). Economic uncertainty, disagreement, and credit markets. Disagreement, and Credit Markets, Working paper.

Buss, A., & Vilkov, G. (2012). Measuring equity risk with option-implied correlations. *Review of Financial Studies*, *25*, 3113-3140.

Campa, J. M. and Chang, P. (1998). The forecasting ability of correlations implied in foreign exchange options, *Journal of International Money and Finance*, 17, 855 – 880.

Campbell, R., Koedijk, K., & Kofman, P. (2002). Increased correlation in bear markets. *Financial Analysts Journal*, 87-94. Working Paper Series 0447, European Central Bank.

Castrén, O., & Mazzotta, S. (2005). Foreign exchange option and returns based correlation forecasts: evaluation and two applications.

Chalamandaris, G. and Tsekrekos, A. E. (2010), "Predictable dynamics in implied volatility surfaces from OTC currency options", *Journal of Banking and Finance*, *34*, 1175-1188.

Chang, B. Y., Christoffersen, P., Jacobs, K., & Vainberg, G. (2012), Option-implied measures of equity risk. *Review of Finance*, *16*, 385-428.

Chen, S. S., Lee, C. F., & Shrestha, K. (2001). On a Mean—Generalized Semivariance Approach to Determining the Hedge Ratio. *Journal of Futures Markets*, *21*, 581-598.

Chen, S. S., Lee, C. F., & Shrestha, K. (2004). An empirical analysis of the relationship between the hedge ratio and hedging horizon: A simultaneous estimation of the short- and long-run hedge ratios. *Journal of Futures Markets*, *24*, 359-386.

Chen, Z., & Petkova, R. (2012). Does idiosyncratic volatility proxy for risk exposure?. *Review of Financial Studies*, *25*, 2745-2787.

Cheung, C. S., Kwan, C. C., & Yip, P. C. (1990). The hedging effectiveness of options and futures: A mean-gini approach. *Journal of Futures Markets*, *10*, 61-73.

Cheung, W., Fung, S., & Tsai, S. C. (2010). Global capital market interdependence and spillover effect of credit risk: evidence from the 2007–2009 global financial crisis. *Applied Financial Economics*, *20*, 85-103.

Chicago Board of Exchange, CBOE S&P 500 Implied Correlation Index, 2009. http://www.cboe.com/micro/impliedcorrelation/ImpliedCorrelationIndicator.pdf

Chiras, D. P., & Manaster, S. (1978). The information content of option prices and a test of market efficiency. *Journal of Financial Economics*, *6*, 213-234.

Chordia, T., & Shivakumar, L. (2002). Momentum, business cycle, and time-varying expected returns. *The Journal of Finance*, *57*, 985-1019.

Chowdhury, A. R. (1991). Futures market efficiency: evidence from cointegration tests. *Journal of Futures Markets*, *11*, 577-589.

Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, *50*, 125-150.

Christoffersen, P., Jacobs, K., & Vainberg, G. (2008, April). Forward-looking betas. In *EFA 2007 Ljubljana Meetings*.

Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, *7*, 174-196.

Cosemans, M. (2011). The Pricing of Long and Short Run Variance and Correlation Risk in Stock Returns, Working paper, Erasmus University - Rotterdam School of Management

Cotter, J., & Hanly, J. (2012). Hedging effectiveness under conditions of asymmetry. *The European Journal of Finance*, *18*, 135-147.

Cox, J. C., Ross, S. A. and Rubinstein, M. (1979). Option pricing: A simplified approach, *Journal of Financial Economics*, *7*, 229–263.

Dacorogna, M. and Lundin, M.: 1999. Correlation of high-frequency financial time series. in P. Lequeux (ed.), Financial Markets Tick by Tick. Wiley & Sons.

Darbar, S. M., & Deb, P. (1997). Co-movements in international equity markets. *Journal of Financial Research*, *20*, 305-22.

Das, S. R., & Uppal, R. (2004). Systemic risk and international portfolio choice. *The Journal of Finance*, *59*, 2809-2834.

David, A., & Veronesi, P. (2002). Option prices with uncertain fundamentals. Olin School of Business Working Paper, 07-001.

De Jong, A., De Roon, F., & Veld, C. (1997). Out-of-sample hedging effectiveness of currency futures for alternative models and hedging strategies. *Journal of Futures Markets*, *17*, 817-837.

De Jong, F., Driessen, J., & Pelsser, A. (2004). On the information in the interest rate term structure and option prices. *Review of Derivatives Research*, *7*, 99-127.

DeMiguel, V., Plyakha, Y., Uppal, R., & Vilkov, G. (2012). Improving portfolio selection using option-implied volatility and skewness. *Journal of Financial and Quantitative Analysis*, *48*, 1-57.

Deng, Q. (2008). Volatility dispersion trading. *Available at SSRN 1156620*.

Driessen, J., Maenhout, P.J. & Vilkov, G. (2009). The price of correlation risk: Evidence from equity options, *Journal of Finance*, *9*, 1377–1406.

Driessen, J., Maenhout, P.J., & Vilkov, G. (2012). Option-implied correlations and the price of correlation risk. *Advanced Risk & Portfolio Management Paper*.

Duarte, J., & Jones, C. S. (2007). The price of market volatility risk. *Unpublished working paper*, *University of Southern California*.

Dungey, M., Fry, R., Tang, C., & Martin, V. (2010). *Are financial crises alike?*. International Monetary Fund.

Ederington, L. H. (1979). The hedging performance of the new futures markets. *The Journal of Finance*, *34*, 157-170.

Elton, E. J., & Gruber, M. J. (1973). Estimating the dependence structure of share prices—implications for portfolio selection. *The Journal of Finance*, *28*, 1203-1232.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, 987-1007.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, *20*, 339-350.

Engle, R., & Kelly, B. (2012). Dynamic equicorrelation. *Journal of Business & Economic Statistics*, *30*, 212-228.

Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical Association*, *74*, 291-298.

Erb, C.B., Harvey, C.R. & Viskanta, T.E. (1994). Forecasting international equity correlations. *Financial Analysis Journal*, *50*, 32-45.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work*. *The journal of Finance*, *25*, 383-417.

Fernandes, M., Medeiros, M. C., & Scharth, M. (2014). Modeling and predicting the CBOE market volatility index. *Journal of Banking & Finance*, *40*, 1-10.

Feunou, B., Fontaine, J-B., Taamouti, A., and Tedognap, R., (2014). Risk premium, variance premium, and the maturity structure of uncertainty. *Review of Finance*, *18*, 219-269.

Fleming, J., Kirby, C., & Ostdiek, B. (2003). The economic value of volatility timing using “realized” volatility. *Journal of Financial Economics*, *67*, 473-509.

Galai, D. (1978). Empirical tests of boundary conditions for CBOE options. *Journal of Financial Economics*, *6*, 187-211.

Garcia, R., & Tsafack, G. (2011). Dependence structure and extreme comovements in international equity and bond markets. *Journal of Banking & Finance*, *35*, 1954-1970.

Garleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-based option pricing. *Review of Financial Studies*, *22*, 4259-4299.

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*, *48*, 1779-1801.

Goncalves, S., & Guidolin, M. (2006) Predictable dynamics in the S&P 500 index options implied volatility surface. *Journal of Business, 79*, 1591-1635.

Goyal, A., & Saretto, A. (2009). Cross-section of option returns and volatility. *Journal of Financial Economics*, *94*, 310-326.

Granger, C. W., & Joyeux, R. (1980). An introduction to long-memory time series models and fractional differencing. *Journal of time series analysis*, *1*, 15-29.

Guidolin, M., & Timmermann, A. (2003). Option prices under Bayesian learning: implied volatility dynamics and predictive densities. *Journal of Economic Dynamics and Control*, *27*, 717-769.

Guo, D. (2000). Dynamic volatility trading strategies in the currency option market, *Review of Derivatives Research, 4*, 133-154.

Han, B. (2007). Stochastic volatilities and correlations of bond yields. *Journal of Finance, 62*, 1491-1524.

Härdle, W. K., & Silyakova, E. (2012). *Implied basket correlation dynamics*(No. 2012-066). SFB 649 Discussion Paper.

Harris, R. D., Shen, J., & Stoja, E. (2010). The Limits to Minimum-Variance Hedging. *Journal of Business Finance & Accounting*, *37*, 737-761.

Harvey, C. R., & Whaley, R. E. (1992). Market volatility prediction and the efficiency of the S & P 100 index option market. *Journal of Financial Economics*, *31*, 43-73.

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of forecasting*, *13*, 281-291.

Hong, H., & Yogo, M. (2012). What does futures market interest tell us about the macroeconomy and asset prices?. *Journal of Financial Economics*, *105*, 473-490.

Hong, Y., Tu, J., & Zhou, G. (2007). Asymmetries in stock returns: Statistical tests and economic evaluation. *Review of Financial Studies*, *20*, 1547-1581.

Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distributions from option prices. *The Journal of Finance*, *51*, 1611-1631.

Jacod, J., Li, Y., Mykland, P. A., Podolskij, M., & Vetter, M. (2009). Microstructure noise in the continuous case: the pre-averaging approach. *Stochastic Processes and their Applications*, *119*, 2249-2276.

Jacquier, A., S. Slaoui (2010): Variance dispersion and correlation swaps. Working paper of Imperial College London

Jennrich, R. I. (1970). An asymptotic χ2 test for the equality of two correlation matrices. *Journal of the American Statistical Association*, *65*, 904-912.

Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. *Journal of financial economics*, *6*, 95-101.

Jiang, G. T., & Tian, Y. S. (2005). The Model-Free Implied Volatility and Its Information Content. *The Review of Financial Studies*, *18*, 1305-1342.

Jobson, J. D., & Korkie, R. M. (1981). Putting Markowitz theory to work. *The Journal of Portfolio Management*, *7*, 70-74.

Johnson, L. L. (1960). The theory of hedging and speculation in commodity futures. *The Review of Economic Studies*, 139-151.

Juhl, T., Kawaller, I. G., & Koch, P. D. (2012). The effect of the hedge horizon on optimal hedge size and effectiveness when prices are cointegrated. *Journal of Futures Markets*, *32*, 837-876.

Karolyi, G. A., & Stulz, R. M. (1996). Why do markets move together? An investigation of US-Japan stock return comovements. *The Journal of Finance*, *51*, 951-986.

Karolyi, G. A., & Stulz, R. M. (2003). Are financial assets priced locally or globally?. *Handbook of the Economics of Finance*, *1*, 975-1020.

Kelly, B. T., Lustig, H., & Van Nieuwerburgh, S. (2011). *Too-systemic-to-fail: What option markets imply about sector-wide government guarantees* (No. w17149). National Bureau of Economic Research.

Kho, B. C. (1996). Time-varying risk premia, volatility, and technical trading rule profits: Evidence from foreign currency futures markets. *Journal of Financial Economics*, *41*, 249-290.

Klitgaard, T., & Weir, L. (2004). Exchange rate changes and net positions of speculators in the futures market. *FRBNY Economic Policy Review*, *10*, 17-28.

Koch, P. D., & Koch, T. W. (1991). Evolution in dynamic linkages across daily national stock indexes. *Journal of International Money and Finance*, *10*, 231-251.

Kolassa, S. (2011). Combining exponential smoothing forecasts using Akaike weights. *International Journal of Forecasting*, *27*, 238-251.

Kolb, R. W., & Okunev, J. (1993). Utility maximizing hedge ratios in the extended mean Gini framework. *Journal of Futures Markets*, *13*, 597-609.

Konstantinidi, E., Skiadopoulos, G., & Tzagkaraki, E. (2008). Can the evolution of implied volatility be forecasted? Evidence from European and US implied volatility indices. *Journal of Banking & Finance*, *32*, 2401-2411.

Konstantinidi, E., & Skiadopoulos, G. (2011). Are VIX futures prices predictable? An empirical investigation. *International Journal of Forecasting*, *27*, 543-560.

Konstantinidi, E., & Skiadopoulos, G. (2013). How does the market variance risk premium vary over time? Evidence from S&P 500 variance swap investment returns.

Koopman, S. J., Jungbacker, B., & Hol, E. (2005). Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measurements. *Journal of Empirical Finance*, *12*, 445-475.

Kostakis, A., Panigirtzoglou, N., & Skiadopoulos, G. (2011). Market timing with option-implied distributions: A forward-looking approach. *Management Science*, *57*, 1231-1249.

Krishnan, C. N. V., Petkova, R., & Ritchken, P. (2009). Correlation risk. *Journal of Empirical Finance*, *16*, 353-367.

Lai, K. S., & Lai, M. (1991). A cointegration test for market efficiency. *Journal of Futures Markets*, *11*, 567-575.

Lai, Y. S., & Sheu, H. J. (2008). An Application of Realized Regression to the Futures Hedging Problem. *Asia Pacific Management Review*, *13*, 655-666.

Lai, Y. S., & Sheu, H. J. (2010). The incremental value of a futures hedge using realized volatility. *Journal of Futures Markets*, *30*, 874-896.

Latane, H. A., & Rendleman, R. J. (1976). Standard deviations of stock price ratios implied in option prices. *The Journal of Finance*, *31*, 369-381.

Lee, H. T. (2009). Optimal futures hedging under jump switching dynamics. *Journal of Empirical Finance*, *16*, 446-456.

Lee, H. T., & Yoder, J. (2007). Optimal hedging with a regime-switching time-varying correlation GARCH model. *Journal of Futures Markets*, *27*, 495-516.

Lien, D., & Luo, X. (1993). Estimating the extended mean-gini coefficient for futures hedging. *Journal of Futures Markets*, *13*, 665-676.

Lien, D., & Tse, Y. K. (1998). Hedging time-varying downside risk. *Journal of Futures Markets*, *18*, 705–722.

Lien, D., & Tse, Y. K. (2000). Hedging downside risk with futures contracts. *Applied Financial Economics*, *10*, 163–170.

Liu, Q. (2009). On portfolio optimization: How and when do we benefit from high-frequency data?. *Journal of Applied Econometrics*, *24*, 560-582.

Longin, F., & Solnik, B. (1995). Is the correlation in international equity returns constant: 1960–1990?. *Journal of international money and finance*, *14*, 3-26.

Longin, F., & Solnik, B. (2001). Extreme correlation of international equity markets. *The Journal of Finance*, *56*, 649-676.

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial studies*, *14*, 113-147.

Longstaff, F. A. (2010). The subprime credit crisis and contagion in financial markets. *Journal of Financial Economics*, *97*, 436-450.

Lopez, J. and Walter, C. (2000). Is implied correlation worth calculating? Evidence from foreign exchange options and historical data, *Journal of Derivatives,* 7, 65 – 81.

MacDonald, R., & Taylor, M. (1988). Metals prices, efficiency and cointegration: some evidence from the London Metal Exchange. *Bulletin of Economic Research*, *40*, 235-240.

Makridakis, S. G., & Wheelwright, S. C. (1974). An analysis of the interrelationships among the major world stock exchanges. *Journal of Business Finance & Accounting*, *1*(2), 195-215.

Markowitz, H. (1952). Portfolio selection*. *The Journal of Finance*, *7*, 77-91.

Martens, M. (2002). Measuring and forecasting S&P 500 index-futures volatility using high-frequency data. *Journal of Futures Markets*, *22*, 497-518.

Martens, M., & Zein, J. (2004). Predicting financial volatility: High-frequency time-series forecasts vis-à-vis implied volatility. *Journal of Futures Markets*, *24*, 1005-1028.

McCracken, M. (2007). Asymptotics for out of sample tests of Granger causality. *Journal of Econometrics, 140,* 719-752.

McMillan, D. G., & Quiroga Garcia, R. (2010). Realized hedge ratio properties, performance and implications for risk management: Evidence from the Spanish IBEX 35 spot and futures markets. *Journal of Risk*, *12*, 33.

Miffre, J. (2001). Efficiency in the pricing of the FTSE 100 futures contracts. *European Financial Management, 7*, 9-22.

Miffre, J. (2004). Conditional OLS minimum variance hedge ratios. *Journal of Futures Markets*, *24*, 945-964.

Morana, C. (2007). Estimating, filtering and forecasting realized betas. *Journal of Financial Forecasting*, *1*, 83-111.

Moskowitz, T. (2003). An Analysis of Covariance Risk and Pricing Anomalies. *Review of Financial Studies 16*, 417-457.

Mueller, P., Stathopoulos, A., & Vedolin, A. (2013). *International correlation risk*. Financial Markets Group, The London School of Economics and Political Science.

Myers, R. J. (1991). Estimating time-varying optimal hedge ratios on futures markets. *Journal of Futures Markets*, *11*, 39-53.

Mykland, P. A., & Zhang, L. (2006). ANOVA for diffusions and Ito processes. *The Annals of Statistics*, *34*, 1931-1963.

Nelken, I. (2006). Variance swap volatility dispersion. *Derivatives Use*, *Trading Regulation*, *11*, 334-344.

Neumann, M., & Skiadopoulos, G. (2013). Predictable Dynamics in Higher-Order Risk-Neutral Moments: Evidence from the S&P 500 Options. *Journal of Financial and Quantitative Analysis*, *48*, 947-977.

Ni, S., Pan, J., & Poteshman, A. (2008). Volatility information trading in the option market. *Journal of Finance*, *63*, 1059-1091.

Noh, J., Engle, R. F., & Kane, A. (1994). Forecasting volatility and option prices of the S&P 500 index. *The Journal of Derivatives*, *2*, 17-30.

Oomen, R. C. A. (2006). Properties of realized variance under alternative sampling schemes. *Journal of Business & Economic Statistics*, *24*, 219-237.

Palandri, A. (2009). The Effects of Interest Rate Movements on Assets' Conditional Second Moments. *School of Economics and Management Aarhus University, Creates Research Paper*, *32*.

Park, T. H., & Switzer, L. N. (1995). Bivariate GARCH estimation of the optimal hedge ratios for stock index futures: A note. *Journal of Futures Markets*, *15*, 61-67.

Patton, A. J. (2004). On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *Journal of Financial Econometrics*, *2*, 130-168.

Patton, A. J., & Verardo, M. (2012). Does beta move with news? Firm-specific information flows and learning about profitability. *Review of Financial Studies*, *25*, 2789-2839.

Paye, B. S. (2008). Do Macroeconomic Variables Predict Aggregate Stock Market Volatility?. *Unpublished Working Paper, Rice University*.

Peng, L. (2005). Learning with information capacity constraints. *Journal of Financial and Quantitative Analysis*, *40*, 307-329.

Peng, L., & Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, *80*, 563-602.

Peng, L., Xiong, W., & Bollerslev, T. (2007). Investor Attention and Time-varying Comovements. *European Financial Management*, *13*, 394-422.

Perez-Quiros, G., & Timmermann, A. (2001). Business cycle asymmetries in stock returns: Evidence from higher order moments and conditional densities. *Journal of Econometrics*, *103*, 259-306.

Politis, D., & Romano, J. (1994). The Stationary Bootstrap. Journal *of the American Statistical Association, 89*, 1303-1313.

Pollet, J.M., & Wilson, M. (2010). Average correlation and stock market returns. *Journal of Financial Economics*, *96*, 364-380.

Pong, S., Shackleton, M. B., Taylor, S. J., & Xu, X. (2004). Forecasting currency volatility: A comparison of implied volatilities and AR (FI) MA models. *Journal of Banking & Finance*, *28*, 2541-2563.

Poon, S. H., & Granger, C. W. (2003). Forecasting volatility in financial markets: A review. *Journal of economic literature*, *41*, 478-539.

Pooter, M. D., Martens, M., & Dijk, D. V. (2008). Predicting the daily covariance matrix for s&p 100 stocks using intraday data—but which frequency to use?. *Econometric Reviews*, *27*, 199-229.

Ratner, M. (1992). Portfolio diversification and the inter-temporal stability of international stock indices. *Global Finance Journal*, *3*, 67-77.

Rubinstein, M. (1976). The valuation of uncertainty income streams and the pricing of options. *Bell Journal of Economics*, *7*, 407-425.

Schmalensee, R., & Trippi, R. R. (1978). Common stock volatility expectations implied by option premia. *The Journal of Finance*, *33*, 129-147.

Schürhoff, N., & Ziegler, A. (2010). *The pricing of systematic and idiosyncratic variance risk*. Working Paper, University of Zurich.

Schwert, G. W. (1989). Why does stock market volatility change over time?. *The Journal of Finance*, *44*, 1115-1153.

Shalit, H. (1995). Mean-Gini hedging in futures markets. *Journal of Futures Markets*, *15*, 617-635.

Sheppard, K. (2008). Economic factors and the covariance of equity returns. Working Paper, University of California, San Diego.

Sheu, H. J., & Lai, Y. S. (2013). Incremental Value of a Futures Hedge Using Realized Ranges. *Journal of Futures Markets*, *34*, 676-689.

Siegel, A. F. (1997). International currency relationship information revealed by cross-option prices. *Journal of Futures Markets*, *17*, 369-384.

Skintzi, V. and Refenes, A. (2005). Implied correlation index: A new measure of diversification, *Journal of Futures Markets*, *25*, 171–197.

Syllignakis, M. N., & Kouretas, G. P. (2011). Dynamic correlation analysis of financial contagion: Evidence from the Central and Eastern European markets. *International Review of Economics & Finance*, *20*, 717-732.

Szymanowska, M., Roon, F., Nijman, T., & Goorbergh, R. (2014). An Anatomy of Commodity Futures Risk Premia. *The Journal of Finance*, *69*, 453-482.

Tacchella, A., Cristelli, M., Zaccaria, A., & Pietronero, L. (2012). Time Evolution of Financial Cross-Correlation Coefficients across Market Crisis. In *International*

Journal of Modern Physics: Conference Series (Vol. 16, pp. 82-92). World Scientific Publishing Company.

Tang, G. Y. (1995). Intertemporal stability in international stock market relationships: A revisit. *The Quarterly Review of Economics and Finance*, *35*, 579-593.

Thomakos, D. D., & Wang, T. (2003). Realized volatility in the futures markets. *Journal of Empirical Finance, 10*, 321-353.

Tian, Y. S. (2011), Extracting Risk-Neutral Density and Its Moments from American Option Prices. *Journal of Derivatives, 18*, 17-34.

Timmerman, A., & Pesaran, M. (1992). A simple nonparametric test of predictive performance. *Journal of Business and Economic Statistics, 10*, 461-465.

Wang, C. (2004). Futures trading activity and predictable foreign exchange market movements. *Journal of Banking & Finance, 28*, 1023-1041.

Wang, T., Wu, J., & Yang, J. (2008). Realized volatility and correlation in energy futures markets. *Journal of Futures Markets*, *28*, 993-1011.

White, H. (2000). A reality check for data snooping. *Econometrica, 68*, 1097-1126.

Yeh, J. H., Huang, C. W., & Hsu, C. C. (2008, December). Hedging Ideally with Realized Covariance. Paper presented at 2008 NTU International Conference on Finance, Taipei, Taiwan

Yoo, J. & Maddala, G.S. (1991). Risk premia and price volatility in futures markets. *Journal of Futures Markets*, *11*, 165-177.

Zhang, L., Mykland, P. A., & Aït-Sahalia, Y. (2005). A tale of two time scales. *Journal of the American Statistical Association*, *100*, 1394-1411.

Zhou, B. (1996). High-frequency data and volatility in foreign-exchange rates. *Journal of Business & Economic Statistics, 14*, 45-52.