

A Simplified Approach to Modelling the Comovement of Asset Returns

Paper Number 04/08

Richard D.F. Harris*
Evarist Stoja
Jon Tucker
Xfi Centre for Finance and Investment
University of Exeter

October 2004

Abstract

This paper proposes a simplified multivariate GARCH model that involves the estimation of only univariate GARCH models, both for the individual return series and for the sum and difference of each pair of series. The covariance between each pair of return series is then imputed from these variance estimates. The model that we propose is considerably easier to estimate than existing multivariate GARCH models and does not suffer from the convergence problems that characterize many of these models. Moreover, the model can be easily extended to include more complex dynamics or alternative forms of the GARCH specification. We use the simplified multivariate GARCH model to estimate the minimum-variance hedge ratio for the FTSE 100 index portfolio, hedged using index futures, and compare it to four of the most widely used multivariate GARCH models. The simplified multivariate GARCH model performs at least as well as the other models that we consider, and in some cases better than them.

Keywords: Multivariate GARCH; Hedging; Minimum-variance hedge ratio; FTSE 100 index.

Acknowledgements: We would like to thank Mark Freeman, Cherif Guermat and Jian Shen for useful comments and suggestions.

Corresponding author: Professor Richard D. F. Harris, Xfi Centre for Finance and Investment, University of Exeter, Exeter EX4 4ST, UK. Tel: +44 (0) 1392 263215. Fax: +44 (0) 1392 262525. Email: R.D.F.Harris@exeter.ac.uk.

1. Introduction

There are many applications in finance that rely on an estimate of the multivariate conditional covariance matrix of returns. Such applications include conditional asset pricing models, portfolio optimization, minimum-variance hedging, value at risk and the pricing of options that depend on more than one underlying asset. Perhaps the most widely used approach to modeling the conditional covariance matrix of returns is the multivariate GARCH class of models.¹ A number of different multivariate GARCH models have been proposed, each imposing a different set of restrictions on the dynamic process that governs the covariance matrix of returns. These models include the Vech and Diagonal Vech models of Bollerslev, Engle and Woolridge (1988), the BEKK model of Engle and Kroner (1995), the Constant Correlation model of Bollerslev (1990), the Factor ARCH model of Engle, Ng and Rothschild (1990) and the Dynamic Conditional Correlation model of Engle and Sheppard (2001).²

While commonly employed in the academic literature, multivariate GARCH models suffer from a number of problems in practice. First, they tend to be computationally burdensome, typically involving the simultaneous estimation of a large number of parameters. This is particularly true of the Vech and BEKK models, both of which impose relatively few restrictions on the dynamic process that governs the evolution of the covariance matrix. Despite recent advances in technology, there are many instances when computational cost is important such as when estimating out-of-sample forecasts of the conditional covariance matrix using a rolling window over a large sample, or where forecasts of the conditional covariance matrix of returns must be computed for a large number of assets in a short period of time (such as when estimating intra-day VaR for a derivatives trading desk). In these instances, multivariate GARCH models are often eschewed by practitioners in favour of simpler alternatives such as exponentially weighted estimators of the covariance matrix. Second, owing to the large number of parameters that must be estimated simultaneously, and the non-concavity of the

¹ Other approaches to estimating the conditional covariance matrix include rolling estimators of the sample covariance matrix, exponentially weighted estimators (JP Morgan, 1994) and multivariate stochastic volatility models (Harvey, Ruiz and Shephard, 1994).

² For a summary of multivariate GARCH models see Bollerslev, Chou and Kroner (1992) and Kroner and Ng (1998).

likelihood function, maximum likelihood estimation of multivariate GARCH models can be problematic. Establishing that estimation has properly converged (i.e. to parameter values that represent a global maximum of the likelihood function rather than a local maximum) involves a potentially computationally intensive grid-search over all of the parameters in the model. Third, compared with their univariate counterparts, it is relatively difficult to construct multi-period forecasts of the covariance matrix using multivariate GARCH models (see, for example, Kroner and Ng, 1998). Fourth, owing to their computational complexity, it is often difficult to extend multivariate GARCH models to include more complicated dynamics such as longer lag specifications, the asymmetric response of volatility to return shocks, and dummy variables to capture seasonality, outliers and structural breaks.

In an attempt to overcome these computational issues, a number of simpler specifications of the multivariate GARCH model have been proposed. However, the simplifications that these models entail generally come at the cost of imposing severe, and often implausible, cross-equation restrictions on the elements of the covariance matrix. For example, in the Diagonal Vech model, each element of the covariance matrix is assumed to evolve independently, meaning that shocks to the variances of individual assets have no impact on the future covariance between them. In the Constant Correlation model, the covariance between individual assets is determined solely by their individual variances. The Factor ARCH model assumes that the covariance between any two assets derives solely from a common covariance with one or more underlying factors. None of these models is able to capture, for example, the well documented feature that correlation among financial assets tends to increase as volatility increases (see, for example, Longin and Solnik, 1995).

In this paper, we propose an alternative, simplified multivariate GARCH model. The model that we propose involves the estimation of only univariate GARCH models, both for the individual return series and for the sum and difference of each pair of series. The covariance between each pair of return series is then imputed from these variance estimates. The model that we propose is considerably more straightforward to estimate than the Vech and BEKK models, and because the estimation involves only univariate GARCH models – and hence only a small number of parameters in any single estimation – it does not suffer from the convergence problems that typically characterize

these models. Moreover, it is easily extended to include the more complex dynamics that are commonly found in the univariate GARCH literature, or to use alternative forms of the GARCH specification. The model that we propose is less restrictive than the Diagonal Vech, Constant Correlation and Factor ARCH models, allowing the covariance between two assets to depend on the history of both their covariance and their individual variances, without imposing the restriction that the correlation coefficient between them is constant over time or that their covariance derives solely from a common covariance with an underlying factor.

We illustrate the simplified multivariate GARCH model by estimating the minimum-variance hedge ratio for the FTSE 100 index portfolio, hedged using index futures. We compare the simplified model to four of the most widely used multivariate GARCH models, namely the Diagonal Vech, Constant Correlation, BEKK and Dynamic Conditional Correlation models. We evaluate the performance of each model both statistically, using a regression of each element of the realized covariance matrix on the corresponding element of the estimated covariance matrix, and economically, by considering the performance of the hedged portfolio. We find that the simplified multivariate GARCH model performs at least as well as the other models that we consider, and in some cases better than them. Moreover, the computation time for the simplified model is considerably lower than for the other models.

The rest of the paper is organised as follows. The following section introduces the simplified multivariate GARCH model. Section 3 presents the empirical application. Section 4 concludes.

2. The Simplified Multivariate GARCH Model

Consider two assets, i and j , whose per-period abnormal returns are given by $\varepsilon_{i,t} = r_{i,t} - \mu_{i,t}$ and $\varepsilon_{j,t} = r_{j,t} - \mu_{j,t}$, where $r_{i,t}$ and $r_{j,t}$ are actual returns and $\mu_{i,t}$ and $\mu_{j,t}$ are conditional mean returns. Given the time $t-1$ information set, Ω_{t-1} , an estimate of the conditional covariance matrix requires estimates of the conditional variances $\sigma_{i,t}^2 = \text{var}(\varepsilon_{i,t} | \Omega_{t-1})$ and $\sigma_{j,t}^2 = \text{var}(\varepsilon_{j,t} | \Omega_{t-1})$, and an estimate of the conditional covariance $\sigma_{ij,t} = \text{cov}(\varepsilon_{i,t}, \varepsilon_{j,t} | \Omega_{t-1})$.

The multivariate GARCH model that we propose involves firstly estimating the conditional variances, $\sigma_{i,t}^2$ and $\sigma_{j,t}^2$, using a univariate GARCH model. We then construct the new series $\varepsilon_{+,t} = \varepsilon_{i,t} + \varepsilon_{j,t}$ and $\varepsilon_{-,t} = \varepsilon_{i,t} - \varepsilon_{j,t}$ and use a univariate GARCH model to estimate $\sigma_{+,t}^2 = \text{var}(\varepsilon_{+,t} | \Omega_{t-1})$ and $\sigma_{-,t}^2 = \text{var}(\varepsilon_{-,t} | \Omega_{t-1})$. An estimate of the conditional covariance, $\sigma_{ij,t}$, can then be obtained using the following identities.

$$\sigma_{+,t}^2 \equiv \sigma_{i,t}^2 + \sigma_{j,t}^2 + 2\sigma_{ij,t} \quad (1)$$

$$\sigma_{-,t}^2 \equiv \sigma_{i,t}^2 + \sigma_{j,t}^2 - 2\sigma_{ij,t} \quad (2)$$

In particular, combining (1) and (2), we have

$$\sigma_{ij,t} \equiv (1/4)(\sigma_{+,t}^2 - \sigma_{-,t}^2) \quad (3)$$

The identity given by (3) is commonly used in the statistics literature in order to derive (unconditional) covariance estimators from (unconditional) variance estimators when no obvious multivariate extension of the variance estimator exists such as in the case of robust estimation of the covariance matrix (see, for example, Huber, 1981). In the context of conditional volatility, Harris and Shen (2003) employ this identity to generalise a univariate robust EWMA estimator to the multivariate case. In this paper, we extend the application of this identity to the multivariate GARCH model.

The simplified multivariate GARCH model involves the estimation of only univariate GARCH models and is therefore considerably easier to implement than the Vech and BEKK models. In particular, because only a few parameters are estimated in each model, it is more likely that maximum likelihood estimation will converge properly and hence much less experimentation is required with different starting values for the model parameters. In the empirical example below, we use the simplest GARCH(1,1) model in order to estimate $\sigma_{i,t}^2$, $\sigma_{j,t}^2$, $\sigma_{+,t}^2$ and $\sigma_{-,t}^2$. However, it would be straightforward to extend the model to allow for more complicated dynamics in both the mean and

volatility of returns using, for example, a more general ARMA(p,q)-GARCH(r,s) specification, or one of the many alternative specifications of the univariate GARCH model, such as the EGARCH model of Nelson (1990). In particular, it would be straightforward to include terms that capture the asymmetric response of volatility to return shocks due to changes in financial leverage, or dummy variables that capture seasonality, outliers and structural breaks, in either the mean or the volatility of returns.

In the simplified model, $\sigma_{ij,t}$ is determined by $\sigma_{+,t}^2$ and $\sigma_{-,t}^2$, which are functions of $\varepsilon_{+,t-1}^2$ and $\varepsilon_{-,t-1}^2$. From the definitions of $\varepsilon_{+,t-1}^2$ and $\varepsilon_{-,t-1}^2$ it can be seen that $\sigma_{ij,t}$ is a function of both $\varepsilon_{i,t-1}^2$ and $\varepsilon_{j,t-1}^2$, and $\varepsilon_{i,t-1}\varepsilon_{j,t-1}$. Therefore the specification of the simplified model allows shocks to the variances of $\varepsilon_{i,t}$ and $\varepsilon_{j,t}$ to affect their future covariance (as well as their respective variances), without imposing the restriction that their correlation is constant over time. In this respect, the simplified model is considerably more flexible than the Diagonal Vech model (which assumes that the covariance between $\varepsilon_{i,t}$ and $\varepsilon_{j,t}$ is determined solely by their lagged covariance) and the Constant Correlation model (which assumes that the covariance between $\varepsilon_{i,t}$ and $\varepsilon_{j,t}$ is determined solely by their lagged variances).

3. Empirical Illustration: Estimation of the Minimum Variance Hedge Ratio

In this section, we illustrate the simplified multivariate GARCH model by estimating the minimum-variance hedge ratio for the FTSE 100 index portfolio, hedged using the FTSE100 index futures contract. When the conditional covariance matrix of spot and futures returns is time-varying, the minimum-variance hedge ratio at time t is equal to

$$h_t = \frac{\sigma_{sf,t}}{\sigma_{f,t}^2} \quad (4)$$

where $\sigma_{sf,t}$ is the conditional covariance of spot and futures returns and $\sigma_{f,t}^2$ is the conditional variance of futures returns (see, for example, Kroner and Sultan, 1993).

Data

We obtained daily closing prices for the FTSE 100 index from Datastream, and for the FTSE 100 index futures contracts from LIFFE, for the period 04 May 1984 to 03 May 2002, which is the longest common sample available. At any one time, there are four futures contracts outstanding. On each day, we use the nearest contract to delivery but rollover to the next nearest contract on the first day of the delivery month in order to avoid thin trading and expiration effects. Using the daily closing spot and futures prices, we computed continuously compounded returns. We removed the returns on the rollover dates from the sample in order to avoid spurious jumps in the futures price that arise from suddenly increasing the maturity of the futures contract. Table 1 gives summary statistics for the spot and futures returns. The mean returns are almost identical for the two series, and very close to zero. The volatility of futures returns is somewhat higher than the volatility of spot returns, which is a common empirical finding (see, for example, Kroner and Sultan, 1993). The ARCH(4) portmanteau test for up to fourth order serial correlation in squared returns shows that both spot and futures returns display significant volatility clustering. Both spot and futures returns are highly leptokurtic (which is consistent with the existence of time-varying volatility) and negatively skewed. The Jarque-Bera statistic very strongly rejects the null hypothesis of normality.

[Table 1]

Methodology

In order to implement the multivariate GARCH model, we first estimate the conditional variances of $r_{s,t}$ and $r_{f,t}$. In principal, any conditional volatility model could be used, but to illustrate our approach, we use the simplest GARCH(1,1) specification. Since our interest is in modeling the conditional covariance matrix of spot and futures returns, and not expected returns, we include only a constant in the mean equation for both spot and futures returns. The model for $r_{s,t}$ is therefore given by

$$r_{s,t} = \mu_s + \varepsilon_{s,t} \tag{5}$$

$$\sigma_{s,t}^2 = \beta_{s,0} + \beta_{s,1}\sigma_{s,t-1}^2 + \beta_{s,2}\varepsilon_{s,t-1}^2 \quad (6)$$

The model for $r_{f,t}$ is given by

$$r_{f,t} = \mu_f + \varepsilon_{f,t} \quad (7)$$

$$\sigma_{f,t}^2 = \beta_{f,0} + \beta_{f,1}\sigma_{f,t-1}^2 + \beta_{f,2}\varepsilon_{f,t-1}^2 \quad (8)$$

After estimating these two models, we compute the residuals $\hat{\varepsilon}_{s,t}$ and $\hat{\varepsilon}_{f,t}$, and using these, compute the new series $\hat{\varepsilon}_{+,t} = \hat{\varepsilon}_{s,t} + \hat{\varepsilon}_{f,t}$ and $\hat{\varepsilon}_{-,t} = \hat{\varepsilon}_{s,t} - \hat{\varepsilon}_{f,t}$. We then estimate the conditional variances of $\hat{\varepsilon}_{+,t}$ and $\hat{\varepsilon}_{-,t}$ using a GARCH(1,1) model with no mean equation specified (since $\hat{\varepsilon}_{s,t}$ and $\hat{\varepsilon}_{f,t}$, and hence $\hat{\varepsilon}_{+,t}$ and $\hat{\varepsilon}_{-,t}$, have zero mean by construction).³ The models for $\hat{\varepsilon}_{+,t}$ and $\hat{\varepsilon}_{-,t}$ are therefore given by

$$\sigma_{+,t}^2 = \beta_{+,0} + \beta_{+,1}\sigma_{+,t-1}^2 + \beta_{+,2}\hat{\varepsilon}_{+,t-1}^2 \quad (9)$$

$$\sigma_{-,t}^2 = \beta_{-,0} + \beta_{-,1}\sigma_{-,t-1}^2 + \beta_{-,2}\hat{\varepsilon}_{-,t-1}^2 \quad (10)$$

The estimated conditional variances of $\hat{\varepsilon}_{+,t}$ and $\hat{\varepsilon}_{-,t}$ are then used to compute the conditional covariance of $r_{s,t}$ and $r_{f,t}$ using equation (3):

$$\sigma_{sf,t} = (1/4)(\sigma_{+,t}^2 - \sigma_{-,t}^2) \quad (11)$$

The estimated conditional variance of futures returns, $\sigma_{f,t}^2$, and conditional covariance of spot and futures returns, $\sigma_{sf,t}$, are then used to compute the minimum-variance hedge

³ The series $\hat{\varepsilon}_{+,t}$ and $\hat{\varepsilon}_{-,t}$ are subject to measurement error but since the information matrix is block diagonal, maximum likelihood estimation is consistent, and when the conditional distribution is normal, fully efficient (see, for example, Kroner and Ng, 1998).

ratio using equation (4). We estimate each of the univariate GARCH models above by maximum likelihood with a conditional normal distribution, using the BHHH algorithm with a convergence criterion of 0.00001 applied to the coefficient values. Note that if $r_{f,t}$ and $r_{s,t}$ are both conditionally normally distributed, then $\varepsilon_{+,t}$ and $\varepsilon_{-,t}$ will also be conditionally normally distributed by construction. If $r_{f,t}$ and $r_{s,t}$ are conditionally non-normal then we can rely on the consistency results of Quasi-Maximum Likelihood (see Bollerslev and Woolridge, 1992).⁴

Evaluation

We compare our results with four of the most commonly used multivariate GARCH models, namely the Diagonal Vech (DVech) model of Bollerslev, Engle and Woolridge (1988), the BEKK model of Engle and Kroner (1995), the Constant Correlation (CC) model of Bollerslev (1990) and the Dynamic Conditional Correlation (DCC) model of Engle and Sheppard (2001). As with the simplified model, we estimate each of these multivariate GARCH models by maximum likelihood with a conditional normal distribution, using the BHHH algorithm with a convergence criterion of 0.00001 applied to the coefficient values. In each case, the mean equation for both spot and futures return is specified to include only a constant.

In order to evaluate the performance of the five multivariate GARCH models, we employ two approaches. The first is a statistical evaluation. If a multivariate GARCH model is correctly specified then it should generate estimates of the realized covariance matrix that are conditionally unbiased. For each element of the covariance matrix, we test this using a regression of the realized variance (or covariance) on the estimated variance (or covariance). As a measure of the realized covariance matrix, we use the squares and cross-product of $\hat{\varepsilon}_{s,t}$ and $\hat{\varepsilon}_{f,t}$. We therefore estimate the following three equations.

⁴ Although Quasi-Maximum Likelihood estimation is consistent if $r_{f,t}$ and $r_{s,t}$ are conditionally non-normal, more efficient estimators of the conditional covariance matrix can be obtained by using, for example, the APARCH model of Ding, Granger and Engle (1993) (see Nelson and Foster, 1996). This is easily accommodated in the simplified multivariate GARCH model.

$$\hat{\varepsilon}_{s,t}^2 = \delta_{s,0} + \delta_{s,1}\sigma_{s,t}^2 + v_{s,t} \quad (12)$$

$$\hat{\varepsilon}_{f,t}^2 = \delta_{f,0} + \delta_{f,1}\sigma_{f,t}^2 + v_{f,t} \quad (13)$$

$$\hat{\varepsilon}_{s,t}\hat{\varepsilon}_{f,t} = \delta_{sf,0} + \delta_{sf,1}\sigma_{sf,t} + v_{sf,t} \quad (14)$$

If the estimated covariance matrix is conditionally unbiased, the intercept in each regression should be zero and the slope coefficient should be unity (see, for example, Andersen and Bollerslev, 1998). We estimate these regressions using OLS. The null hypothesis of conditional unbiasedness is tested for each regression using an F-statistic. The second approach is an economic evaluation. In particular, we report the standard deviation of the hedged portfolio for each model. The more accurate the estimated conditional covariance matrix, the lower should be the standard deviation of the hedged portfolio. We also report the standard deviation of the hedge ratio itself, which gives some indication of the likely transaction costs associated with a dynamic hedging strategy based on each model.

Results

Table 2 reports the correlation matrices for $\sigma_{s,t}^2$, $\sigma_{f,t}^2$ and $\sigma_{sf,t}$ across the five models. As expected, owing to the different restrictions that they impose, the five multivariate GARCH models yield quite different estimates of the covariance matrix of spot and futures returns. For all three elements of the covariance matrix, the lowest correlation is between the DCC and CC models, while the highest correlation is between the simplified model and the CC model, except for $\sigma_{f,t}^2$, where the correlation between the DCC and DVech models is marginally higher than between the simplified model and the CC model. Overall, the lowest correlation is 0.81 (between the DCC model and the CC model for $\sigma_{f,t}^2$), while the highest correlation is 0.99 (between the simplified model and the CC model for $\sigma_{s,t}^2$).

[Table 2]

Table 3 reports the results of the regressions to test the conditional unbiasedness of $\sigma_{s,t}^2$, $\sigma_{f,t}^2$ and $\sigma_{sf,t}$ for each of the five multivariate GARCH models. For $\sigma_{s,t}^2$, the null hypothesis of conditional unbiasedness is rejected at the five percent significance level for all of the models except the simplified model, although for the CC and DCC models, the rejection is marginal. The rejection is particularly strong for the BEKK model. In contrast, for $\sigma_{f,t}^2$, the null hypothesis of conditional unbiasedness is rejected for all five models, although the rejection is weakest for the CC model, followed by the DVEch model, the simplified model and the DCC model. Again, the null hypothesis of conditional unbiasedness is very strongly rejected for the BEKK model. For $\sigma_{sf,t}$, the null hypothesis of conditional unbiasedness is rejected at the five percent level for all of the models except the CC and BEKK models, although for the remaining models, the rejection is very marginal. Overall, it would appear that in terms of conditional unbiasedness, there is little to choose between the simplified model and the DVEch, CC and DCC models, but that the BEKK model performs significantly worse than these models.

[Table 3]

The first row of Table 4 reports the standard deviation of the estimated hedge ratio for the five multivariate GARCH models. The CC model yields the lowest hedge ratio standard deviation, reflecting the fact that it imposes the restriction that the correlation coefficient between spot and futures returns is constant, while the other models allow the correlation coefficient to vary over time. The BEKK model yields the highest hedge ratio standard deviation. The simplified model yields a hedge ratio standard deviation that is marginally higher than the DCC model. The second row of Table 4 reports the standard deviation of the hedge portfolio daily return for the five multivariate GARCH models. The DVEch and DCC models yield the lowest hedge portfolio standard deviation, although the simplified model yields a standard deviation that is only marginally higher than these. The standard deviation for the CC and BEKK models is significantly higher than for the other models. Thus, in terms of hedge portfolio

standard deviation, the simplified model performs almost as well as the DVech and DCC models, and substantially better than the CC and BEKK models.

[Table 4]

Finally, Table 5 reports the estimation time for each of the five multivariate GARCH models. Even in this simple bivariate case, it is clear that estimation of multivariate GARCH models can be relatively time-consuming, with the BEKK model taking almost six minutes to estimate.⁵ However, the estimation time for the simplified model is less than half that of the other models. Moreover, these figures significantly understate the true computational advantage of the simplified model since they are based on the final estimation for each model after experimenting with different starting values for the model parameters. The simplified model converged directly, irrespective of the starting values for the parameters, while, in contrast, estimation of the other models required considerable experimentation with different starting values in order to achieve convergence. The DVech and BEKK models proved to be particularly problematic in this respect.

[Table 5]

4. Conclusion

While commonly employed in the academic literature, multivariate GARCH models suffer from a number of problems in practice owing to the complexity of their specification. In particular, because of the large number of parameters that must be estimated simultaneously, multivariate GARCH models tend to be computationally burdensome. Moreover, because the likelihood function of these models is not globally concave, there is no guarantee that maximum likelihood estimation will converge to the correct parameter values, particularly when the models are supplemented by more complicated dynamics, asymmetric terms or dummy variables.

⁵ The models were estimated using the RATS 5.01 package on a Pentium IV 2.8 GHz PC.

In this paper, we propose a simple but effective multivariate GARCH model that overcomes these problems. The model that we propose involves the estimation of only univariate GARCH models, both for the individual return series and for the sum and difference of each pair of series. The covariance terms in the covariance matrix are then imputed from these variance estimates. Since the estimation involves only univariate GARCH models, it is considerably more straightforward to estimate than existing multivariate GARCH models and does not suffer from the convergence problems that typically characterise many of these models.

We illustrate the simplified multivariate GARCH model, and compare it to four of the most widely used multivariate GARCH models – the Diagonal Vech model, the Constant Correlation model, the BEKK model and the Dynamic Conditional Correlation model – by estimating the minimum-variance hedge ratio for the FTSE 100 index portfolio, hedged using index futures. We evaluate the performance of each model both statistically, using a regression of each element of the realized covariance matrix on the corresponding element of the estimated covariance matrix, and economically, by considering the performance of the hedged portfolio. We find that by both measures, the simplified multivariate GARCH model performs at least as well as the other models that we consider, and in some cases better than them.

References

- Andersen, T., and T. Bollerslev, 1998, “Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts”, *International Economic Review* 39, 885-905.
- Bollerslev, T., 1990, “Modelling the Coherence in Short Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model”, *Review of Economics and Statistics* 72, 498–505.
- Bollerslev, T., R. Chou and K. Kroner, 1992, “ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence”, *Journal of Econometrics*.
- Bollerslev, T, R. Engle and J. Wooldridge, 1988, A Capital Asset Pricing Model with Time-Varying Covariances, *Journal of Political Economy* 96, 116–131.
- Bollerslev, T. and J. Wooldridge, 1992, “Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances”, *Econometric Reviews* 11, 143–172.
- Ding, Z., C. Granger and R. Engle, 1993, “A Long Memory Property of Stock Market Returns and a New Model”, *Journal of Empirical Finance* 1, 83-106.
- Engle, R., and K. Kroner, 1995, “Multivariate Simultaneous Generalised ARCH”, *Econometric Theory* 11, 122–50.
- Engle, F., and K. Sheppard, K., 2001, “Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH”, NBER Working Paper 8554.
- Engle, F., V. Ng and M. Rothschild, 1990, “Asset Pricing with a Factor ARCH Covariance Structure: Empirical Estimates for Treasury Bills”, *Journal of Econometrics* 45, 213–237.
- Glosten, L., R. Jagannathan and D. Runkle, 1993, “On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks”, *Journal of Finance* 48, 1779–1801.
- Harris, R. D. F., and J. Shen, 2003, “Robust Estimation of the Optimal Hedge Ratio”, *Journal of Futures Markets* 23, 799-816.
- Harvey, A., E. Ruiz and N. Shephard, 1994, “Multivariate Stochastic Variance Models”, *Review of Economic Studies* 61, 247-264.
- Huber, P., 1981, *Robust Statistics*, Wiley.
- JP Morgan, 1996, *RiskMetrics* Technical Document, fourth edition, New York.
- Kroner, K., and V. Ng, 1998, “Modelling Asymmetric Co-Movements of Asset Returns”, *Review of Financial Studies* 11, 817–44.

Kroner, K., and J. Sultan, 1993, "Time Varying Distribution and Dynamic Hedging with Foreign Currency Futures", *Journal of Finance and Quantitative Analysis* 28, 535-551.

Longin, F., and B. Solnik, 1995, "Is the Correlation in International Equity Returns Constant: 1960-1990?", *Journal of International Money and Finance* 14, 3-26.

Nelson, D., 1990, "Conditional Heteroskedasticity in Asset Returns: A New Approach", *Econometrica* 59, 347-370.

Nelson, D., and D. Foster, 1996, "Asymptotic Filtering Theory for Univariate ARCH Models", *Econometrica* 62, 1-41.

Table 1 Summary Statistics of FTSE 100 Spot and Futures Returns

| | $r_{s,t}$ | $r_{f,t}$ |
|---------------------------|-----------|-----------|
| Mean | 0.034% | 0.034% |
| Standard Deviation | 1.002% | 1.144% |
| Skewness | -0.706 | -1.101 |
| Excess kurtosis | 9.507 | 16.602 |
| Jarque-Bera | 17414.126 | 52870.870 |
| ARCH(4) | 734.221 | 340.620 |

Notes: The table reports summary statistics for continuously compounded spot and futures returns for the FTSE 100 for the period 04 May 1984 to 03 May 2002. The Jarque-Bera statistic tests the null hypothesis of zero skewness and excess kurtosis, and has a $\chi^2(2)$ distribution with a critical value of 5.99 at the 5% significance level. The ARCH(4) statistic tests the null hypothesis that the first four partial autocorrelations of squared returns are zero, and has a $\chi^2(4)$ distribution with a critical value of 9.49 at the 5% significance level.

Table 2 Correlation Matrices of $\sigma_{s,t}^2$, $\sigma_{f,t}^2$ and $\sigma_{sf,t}^2$

Panel A: Correlation Matrix of $\sigma_{s,t}^2$

| | DVech | CCR | BEKK | DCC | Simplified Model |
|-------------------------|--------------|------------|-------------|------------|-------------------------|
| DVech | 1.000 | 0.913 | 0.938 | 0.975 | 0.933 |
| CCR | | 1.000 | 0.905 | 0.855 | 0.994 |
| BEKK | | | 1.000 | 0.935 | 0.936 |
| DCC | | | | 1.000 | 0.895 |
| Simplified Model | | | | | 1.000 |

Panel B: Correlation Matrix of $\sigma_{f,t}^2$

| | DVech | CCR | BEKK | DCC | Simplified Model |
|-------------------------|--------------|------------|-------------|------------|-------------------------|
| DVech | 1.000 | 0.879 | 0.882 | 0.967 | 0.959 |
| CCR | | 1.000 | 0.881 | 0.809 | 0.966 |
| BEKK | | | 1.000 | 0.919 | 0.934 |
| DCC | | | | 1.000 | 0.929 |
| Simplified Model | | | | | 1.000 |

Panel C: Correlation Matrix of $\sigma_{sf,t}^2$

| | DVech | CCR | BEKK | DCC | Simplified Model |
|-------------------------|--------------|------------|-------------|------------|-------------------------|
| DVech | 1.000 | 0.882 | 0.886 | 0.976 | 0.916 |
| CCR | | 1.000 | 0.9140 | 0.825 | 0.986 |
| BEKK | | | 1.000 | 0.892 | 0.940 |
| DCC | | | | 1.000 | 0.885 |
| Simplified Model | | | | | 1.000 |

Notes: The table reports the correlation matrices of $\sigma_{s,t}^2$, $\sigma_{f,t}^2$ and $\sigma_{sf,t}^2$ across the five multivariate GARCH models.

Table 3 Results for Conditional Bias Regressions

| Panel A $\hat{\varepsilon}_{s,t}^2 = \delta_{s,0} + \delta_{s,1}\sigma_{s,t}^2 + v_{s,t}$ | | | | | |
|--|--------------|------------|-------------|------------|-------------------------|
| | DVech | CCR | BEKK | DCC | Simplified Model |
| $\hat{\delta}_{s,0}$ | 6.4E-06 | -7.0E-06 | -3.1E-05 | 6.6E-06 | 9.7E-06 |
| $\hat{\delta}_{s,1}$ | 0.881 | 1.115 | 1.263 | 0.881 | 0.918 |
| F-statistic | 4.436 | 3.771 | 27.654 | 4.793 | 2.190 |

| Panel B $\hat{\varepsilon}_{f,t}^2 = \delta_{f,0} + \delta_{f,1}\sigma_{f,t}^2 + v_{f,t}$ | | | | | |
|--|--------------|------------|-------------|------------|-------------------------|
| | DVech | CCR | BEKK | DCC | Simplified Model |
| $\hat{\delta}_{f,0}$ | 1.9E-05 | 1.7E-05 | 3.1E-05 | 2.0E-05 | 2.0E-05 |
| $\hat{\delta}_{f,1}$ | 0.804 | 0.849 | 0.722 | 0.797 | 0.824 |
| F-statistic | 7.664 | 6.277 | 32.771 | 9.181 | 8.213 |

| Panel C $\hat{\varepsilon}_{sf,t}\hat{\varepsilon}_{f,t} = \delta_{sf,0} + \delta_{sf,1}\sigma_{sf,t} + v_{sf,t}$ | | | | | |
|--|--------------|------------|-------------|------------|-------------------------|
| | DVech | CCR | BEKK | DCC | Simplified Model |
| $\hat{\delta}_{sf,0}$ | 3.5E-06 | -4.0E-06 | -1.1E-05 | 2.9E-06 | 1.0E-05 |
| $\hat{\delta}_{sf,1}$ | 0.908 | 1.066 | 1.052 | 0.913 | 0.909 |
| F-statistic | 3.440 | 1.848 | 2.372 | 3.392 | 3.150 |

Notes: The table reports the results of estimating the conditional bias regressions for $\sigma_{s,t}^2$, $\sigma_{f,t}^2$ and $\sigma_{sf,t}^2$. The F-statistic tests the null hypothesis that $\hat{\delta}_{i,0} = 0$ and $\hat{\delta}_{i,1} = 1$ ($i = s, f, sf$), and has an $F(2, 4518)$ distribution with a critical value of 3.00 at the 5% significance level.

Table 4 Results for Minimum Variance Hedge Portfolio

| | DVech | CC | BEKK | DCC | Simplified Model |
|--------------------|--------------|-----------|-------------|------------|-------------------------|
| $\sigma_{\hat{h}}$ | 0.083 | 0.072 | 0.105 | 0.097 | 0.100 |
| σ_p | 0.400% | 0.481% | 0.512% | 0.400% | 0.403% |

Notes: The table reports the standard deviation of the estimated hedge ratio, $\sigma_{\hat{h}}$, and the standard deviation of the hedge portfolio return, σ_p , for each of the five multivariate GARCH models.

Table 5 Estimation Time

| DVech | CC | BEKK | DCC | Simplified Model |
|--------------|--------------|--------------|--------------|-------------------------|
| 5mins 23secs | 4mins 31secs | 5mins 49secs | 4mins 36secs | 2mins 15secs |

Notes: The table reports the estimation time for each of the five multivariate GARCH models. The models were estimated using RATS 5.01 on a Pentium IV 2.8 GHz PC.