

NCER Working Paper Series

Forecasting Equicorrelation

Adam E Clements
Christopher A Coleman-Fenn
Daniel R Smith

Working Paper #72
April 2011

Forecasting Equicorrelation

Adam E. Clements, Christopher A. Coleman-Fenn, and Daniel R. Smith
School of Economics and Finance, Queensland University of Technology

March 31, 2011

Abstract

This article examines the out-of-sample forecast performance of several time-series models of equicorrelation, a mean of the off-diagonal elements of a covariance matrix. Building on the existing Dynamic Conditional Correlation and Linear Dynamic Equicorrelation models, we propose adapting the latter to include measures of equicorrelation based on high-frequency intraday data, as well as a forecast of equicorrelation implied by the options market. Using state-of-the-art statistical evaluation technology, we find that the use of both the realised measures and the implied equicorrelation outperform those models that use daily data alone. However, the out-of-sample forecasting benefits of implied equicorrelation disappear when used in conjunction with the realised measures.

Corresponding author

Adam Clements
School of Economics and Finance
Queensland University of Technology
GPO Box 2434, Brisbane, Qld, Australia 4001
email a.clements@qut.edu.au
Ph +61 7 3138 2525, Fax +61 7 3138 1500.

1 Introduction

Recently, there has been developing interest in the utilisation and modelling of the average correlation, or equicorrelation, of assets; which is defined as the mean of the off-diagonal elements of a correlation matrix. There are several reasons why this variable may be of importance in the financial economics literature.

To begin with, the concept of equicorrelation was developed nearly forty years ago by Elton and Gruber (1973) who showed that assuming all of the off-diagonal pairs of a correlation matrix were equal lead to superior portfolio allocation results and also reduced estimation noise. More recently, Pollet and Wilson (2010) develop a theoretical argument and provide empirical evidence for why the average correlation of a stock market index is strongly related to future market returns while stock market variance is not. Relatedly, Driessen, Maenhout, and Vilkov (2009) conduct an empirical exercise in which the entire S&P 100 Index variance risk premium is attributable to the correlation risk premium.

The concept of equicorrelation may also be of use for portfolio managers interested in assessing the level of diversification of their assets. The equicorrelation of a portfolio is the only scalar measure we are aware of that summarises the degree of interdependence within the portfolio and hence diversification benefits. Forecasts of equicorrelation may then provide portfolio managers a simple guide to the interrelationships of their portfolio constituents into the future that may be more readily interpretable than forecasting each of the potentially numerous individual correlation pairs.

Equicorrelation, or at least the equicorrelation implied by options markets, is also of use in the derivatives market. The return on a strategy known as dispersion trading, where one goes long an option on a basket of assets and shorts options on each of the constituents, is dependent only on correlations after each of the individual options are delta hedged. It is common to make the assumption that all of these correlations are equal, resulting in the value of the position depending upon the evolution of the implied equicorrelation (Engle and Kelly, 2008). Partially motivated by its use in dispersion trading, the Chicago Board of Exchange (CBOE) has published the Implied Correlation Index, the mean correlation of the S&P 500 Index for the proceeding 22-trading-days, since July 2009. Therefore, in addition to being used in forming expectations of market returns, equicorrelation is also of use in popular derivatives trading strategies.

From an econometric point of view, the assumption of equicorrelation imposes structure on problems that are otherwise intractable. In the description of candidate forecasting models that follows, it becomes apparent that for the majority of models discussed the length of the time-series available is required to be significantly larger than the number of assets in the portfolio for the results to be reliable; this is problematic for very large portfolios, such as the S&P 500 Index. In the vast majority of models, the time-span available for estimation purposes is limited to the shortest lived stock within the portfolio, which is conceivably quite short; for example, even the very large firm Kraft

Inc. has only been a publicly traded firm since mid-2007. The utilisation of equicorrelation can circumvent this restriction and allows for solutions to problems that would otherwise be intractable. Given the theoretical and empirical relevance of equicorrelation just described, we now provide a brief overview of previous literature that informs the equicorrelation forecasting problem.

While we focus on equicorrelation forecasting here, there has been significant interest in correlation forecasting more generally and a large range of competing candidate models exist that we may still utilise for our more narrow purposes. It is beyond the scope of this paper to provide a thorough review of all of these models and such theoretical surveys already exist in Bauwens, Laurent, and Rombouts (2006) and Silvennoinen and Teräsvirta (2009), while Laurent, Rombouts, and Violante (2010) conduct an extensive empirical comparison of the out-of-sample forecast performance of 125 conditional correlation models. Instead, for guidance on filtering the list of candidate models we refer to the review of Multivariate Generalised Autoregressive Conditional Heteroscedasticity (MGARCH) models by Silvennoinen and Teräsvirta (2009). They state that an ideal time-series model of conditional covariance or correlation matrices faces competing requirements; while the specification must be flexible enough to model the dynamic structure of variances and covariances, it is also desirable to remain parsimonious for the purposes of estimation.

The Dynamic Conditional Correlation (DCC) model of Engle (2002), adapted for consistent estimation by Aielli (2009) in his cDCC model, allows for the forecasting of conditional correlations with the optimisation of just two parameters while still retaining a reasonable degree of flexibility; it is on this model and variations thereof that we focus for generating our equicorrelation forecasts. In addition to meeting the criteria of flexibility and parsimony suggested by Silvennoinen and Teräsvirta (2009), the cDCC model has also become a benchmark in the correlation forecasting literature and provides a natural starting point to which competing forecasts will be compared. This is primarily due to the fact that the cDCC model was the basis of construction for the Dynamic Equicorrelation (DECO) model of Engle and Kelly (2008), and also the Linear Dynamic Equicorrelation (LDECO) from the same paper. As the majority of properties invoked in this article rely on results from the DECO and LDECO models, it is natural to use the unrestricted cDCC as the benchmark. The LDECO model of Engle and Kelly (2008) is an adaptation of the cDCC model which imposes the equicorrelation assumption, yet also introduces added flexibility by allowing the constituents of the portfolio of interest to enter and exit freely, and also the number of portfolio constituents to change. We propose two modifications borrowed from the univariate volatility literature to this LDECO model.

In the univariate volatility forecasting literature, it is now well known that Realised Volatility (RV), which is defined as a sum of squared high-frequency intraday returns, provides a superior measure of the latent volatility of an asset than the square of daily closing price returns. This RV measure also has the advantage that traditional time-series models can be constructed for it, a popular example of such a formulation is the Heterogenous Autoregressive (HAR) model of Corsi (2009). Further, when added as an exogenous regressor

to GARCH models, Blair, Poon, and Taylor (2001) find that the co-efficient acting on RV has statistically significant explanatory power. In the current context, it is possible to use the RV of an index and the RVs of each of its constituents to form a measure of equicorrelation that is based on this improved method of analysing volatility. We believe that similar benefits to those just mentioned for univariate volatility forecasting may also exist in using this measure of equicorrelation over the use of daily returns that is utilised in the cDCC and DECO specifications.

Alternatively, the use of realised measures of latent variables has been extended into the multivariate setting by, for example, Barndorff-Nielson, Hansen, Lunde, and Shephard (2010) and Corsi and Audrino (2007). In these papers, it is shown that utilising high-frequency intraday data provides superior estimates of the level of latent covariance between assets, although one must be wary of market microstructure effects. A time-series model for correlation based on intraday data has been put forth by Corsi and Audrino (2007), who extend the univariate HAR model for RV to its multivariate analogue, and demonstrate promising results in the bi-variate setting. We use one of these multivariate measures to estimate the realised equicorrelation by taking a mean of the off-diagonal elements of the covariance matrix, this may then be utilised in the analysis of equicorrelation forecasting.

In the univariate volatility forecasting literature, it has also been shown that forecasts generated from the options market, implied volatility (IV), can contain information incremental to that in the physical market. The standard argument for the inclusion of such implied measures is that as options are priced with reference to a future-dated payoff, an efficient options market should incorporate both historical information as well as a forecast of information relevant to the pricing of the options. The use of IV has been reported by Poon and Granger (2003) to outperform time-series based forecasts in the majority of research that they reviewed. Further, in a recent study of IV forecasting ability which incorporated an estimate of the volatility risk-premium, Becker, Clements and Coleman-Fenn (2009) show that models incorporating IV cannot be statistically separated from the best performing time-series models unconditionally, and dominate in low-volatility periods.

Motivated by prior results in the univariate volatility forecasting literature, we examine two potential improvements to the LDECO model, one that relies on high-frequency intraday measures of equicorrelation, in a similar vein to RV in the univariate volatility literature, and a forecast of equicorrelation from the options market, in a similar vein to IV in the univariate volatility literature. We propose replacing the equicorrelation measure suggested by Engle and Kelly (2008), which uses daily closing prices in its calculation, with two realised equicorrelation measures which incorporate higher frequency data. As part of our robustness checks, we also propose utilising an alternative non-parametric statistic of equicorrelation from intraday data, the Spearman rank correlation.

The other amendment that we propose is to extend the LDECO model to include a forecast of equicorrelation generated from the options market, implied equicorrelation (IC). The inclusion of IC in a time-series model of equicorrelation is analogous to the inclusion of implied volatility (IV) as an exogenous

regressor in GARCH models adopted in, for example, Blair, Poon, and Taylor (2001). As Silvennoinen and Teräsvirta (2009) have also found that the VIX is important in forecasting correlation matrices, we see promise in the inclusion of IC in equicorrelation forecasting similar to the success observed in the univariate volatility framework. We are also able to combine the two amendments that we propose, this allows us to analyse whether any improved forecasting ability over LDECO gained from incorporating IC disappears when one includes the realised equicorrelation measures.

We generate 22-trading-day ahead forecasts of equicorrelation using ten models that include existing time-series specifications and the amendments that we propose, which will all be specified in detail below. To evaluate the forecast performance of these models, we employ the Model Confidence Set (MCS) methodology of Hansen, Lunde, and Nason (2003, 2010). The MCS has been utilised previously in the univariate volatility context by, among others, Becker and Clements (2008) and in the multivariate setting by Laurent, Rombouts, and Violante (2010). An interesting result of the later paper in the context of this article is that in turbulent times the DECO model, which is closely linked to the LDECO model employed here, dominates among DCC models, including those which relax the equicorrelation assumption and even include asymmetry terms.

We find that both of the proposed variations to the LDECO model individually lead to superior in-sample fit and out-of-sample forecast performance. While the specification that includes both of the discussed extensions perhaps unsurprisingly provides the best in-sample fit, for the purposes of out-of-sample forecasting, the use of realised equicorrelation alone is optimal for the data considered here.

The paper proceeds as follows. Section 2 provides an overview of the nesting framework and the models considered in this paper, Section 3 describes how forecast performance is statistically evaluated, Section 4 details the data utilised, Section 5 presents and analyses the results and Section 6 concludes.

2 General Framework and Models Considered

We begin this Section with a brief description of the general framework that nests the problem considered and a review of previously proposed MGARCH models. As defined in Bollerslev (1990) and Engle and Kelly (2008), the multivariate distribution of asset returns, when assumed to be Gaussian, can be written as

$$r_{t|t-1} \sim N(0, H_t), \quad H_t = D_t R_t D_t \quad (1)$$

where D_t is the diagonal matrix of conditional standard deviations and R_t is a conditional correlation matrix. The multivariate Gaussian log-likelihood

function of this distribution is given by

$$\begin{aligned}
L &= -\frac{1}{T} \sum_{t=1}^T (n \log(2\pi) + \log|H_t| + r_t' H_t^{-1} r_t) \\
&= -\frac{1}{T} \sum_{t=1}^T (n \log(2\pi) + 2 \log|D_t| + r_t' D_t^{-2} r_t - \tilde{r}_t' \tilde{r}_t) \\
&\quad -\frac{1}{T} \sum_{t=1}^T (\log|R_t| + \tilde{r}_t' R_t^{-1} \tilde{r}_t) \\
L &= L_{Vol}(\theta) + L_{Corr}(\theta, \Phi)
\end{aligned} \tag{2}$$

where \tilde{r}_t are the volatility standardised returns given by $\tilde{r}_t = D_t^{-1} r_t$.

The above definition demonstrates that the optimisation problem can be decomposed into an optimisation over the volatility specific parameters and a secondary optimisation over the correlation parameters, which depend on the volatility specific parameters through the standardised returns. As the focus here is on forecasting equicorrelation and not univariate volatility forecasting, we follow Engle and Kelly (2008) and assume that conditional volatilities follow a GARCH(1,1) process; this allows us to concentrate solely on the modelling of the conditional correlation matrix, R_t , through time. Further, in the log-likelihood comparisons presented in the forthcoming results, we only compare the $L_{Corr}(\theta, \Phi)$ component of L .

While numerous time-series models have been proposed for the forecasting of R_t , we begin with the Dynamic Conditional Correlation (DCC) model of Engle (2002), adapted for consistent estimation (cDCC) by Aielli (2009), which allows for time-varying pairwise correlations to be optimised across only two parameters. Under the cDCC model, conditional correlation is given by

$$R_t^{cDCC} = \tilde{Q}_t^{\frac{1}{2}} Q_t \tilde{Q}_t^{\frac{1}{2}} \tag{3}$$

where Q_t has the following dynamics

$$Q_t = \bar{Q}(1 - \alpha - \beta) + \alpha \tilde{Q}_{t-1}^{\frac{1}{2}} \tilde{r}_{t-1} \tilde{r}_{t-1}' \tilde{Q}_{t-1}^{\frac{1}{2}} + \beta Q_{t-1} \tag{4}$$

where \bar{Q} is the unconditional correlation matrix, \tilde{Q}_t replaces the off-diagonal elements of Q_t with zeros but maintains its main diagonal, and the following conditions must hold to ensure stationarity, $\alpha > 0$, $\beta > 0$, $\alpha + \beta < 1$. Similar in structure to the univariate GARCH model, the cDCC model allows for an unconditional correlation matrix, or correlation targeting, as well as an innovation term on the lagged residuals, and a persistence term for lagged values of Q_t . The cDCC model is attractive given its analytical tractability, flexibility, and low number of parameters; however, for the practical applications for which portfolio managers require solutions, the cDCC model begins to falter. As can be seen in (2), the optimisation process requires finding the inverse and determinant of potentially very large matrices. While alternatives may be found

for taking the inverse, such as Gauss-Jordan elimination, we are not aware of any such alternative pathways for the determinant. As this procedure must be repeated at each time step for each iteration of the optimisation algorithm, the estimation procedure can quickly become cumbersome.

2.1 The Dynamic Equicorrelation Model

To circumvent the computational issue just described, Engle and Kelly (2008) make use of the simplifying assumption of equicorrelation. For each point in time, all off-diagonal pairs of the conditional correlation matrix are assumed to be equal to the equicorrelation scalar, ρ_t . It is this equicorrelation which varies through time and they propose the Dynamic Equicorrelation model (DECO) which is similar in specification to the cDCC model:

$$R_t^{DECO} = (1 - \rho_t^{DECO})I_n + \rho_t^{DECO} J_n \quad (5)$$

where ρ_t^{DECO} is the average of the off-diagonal elements of Q_t as specified in (4), and J_n is the $n \times n$ matrix of ones; similarly,

$$\rho_t^{DECO} = \frac{2}{n(n-1)} \sum_{i>j} \frac{q_{i,j,t}}{\sqrt{q_{i,i,t}q_{j,j,t}}} \quad (6)$$

where $q_{i,j,t}$ is the i, j th element of the matrix Q_t .

The restriction of equicorrelation employed by the DECO model significantly decreases estimation time by allowing the inverse and determinant of the conditional correlation matrix to be respectively solved analytically by

$$R_t^{-1} = \frac{1}{1 - \rho_t} \left(I_n - \frac{\rho_t}{(1 + [n-1]\rho_t)} \right) J_n \quad (7)$$

and

$$|R_t| = (1 - \rho_t)^{n-1} (1 + [n-1]\rho_t) \quad (8)$$

where ρ_t is the equicorrelation, and I_n is the n -dimensional identity matrix where n is the number of assets in the portfolio; Engle and Kelly (2008) show that R_t^{-1} exists *iff* $\rho_t \neq 1$ and $\rho_t \neq \frac{-1}{n-1}$, and R_t is positive definite *iff* $\rho_t \in \left(\frac{-1}{n-1}, 1 \right)$.

Again if we consider the practical perspective of the portfolio manager, a limitation of the cDCC, DECO, and MGARCH models in general is that they are unable to handle changes in the portfolio constituents or the number of assets within the portfolio. An example of where this may raise a problem is for the S&P 500 Index, where the portfolio constituents may change frequently. Hence, Engle and Kelly (2008) propose a variation of the DECO model which allows for such changes; the Linear DECO (LDECO) model.

In their original paper, Engle and Kelly (2008) note of the LDECO model that “key in this approach is extracting a measurement of the equicorrelation in each time period using a statistic that is insensitive to the indexing of assets in the return vector”, with their statistic of choice being

$$u_t = \frac{[(\sum_i \tilde{r}_{i,t})^2 - \sum_i (\tilde{r}_{i,t}^2)]/n(n-1)}{\sum_i (\tilde{r}_{i,t}^2)/n} \quad (9)$$

The equicorrelation innovation term, u_t , can be decomposed into an estimate of the covariance of returns given by the numerator, while the denominator is an estimate of the variance for all assets. As $\tilde{r}_{i,t}$ are volatility standardised returns they should have unit variance and, therefore, the numerator should be a correlation estimate. Technically, the numerator is not restricted to lie in the range that ensures positive definiteness of R_t , and it lacks robustness to deviations from unity for the conditional variance estimate Engle and Kelly (2008). However, Engle and Kelly (2008) demonstrate that the denominator of u_t standardizes this covariance estimate by an estimate of the common variance; this ensures that u_t lies within the range necessary for positive definiteness of the correlation matrix. The time-series model of correlation based on the assumptions of LDECO is given by

$$\rho_t = \omega + \alpha u_{t-1} + \beta \rho_{t-1} \quad (10)$$

It is here that we propose our first modification of the LDECO model. We contend that a more robust measure of the equicorrelation on a given day is available from the realised variance technology. By utilising high-frequency intraday data, one is able to extract a greater understanding of the interrelationships of the relevant assets during the trading day, rather than relying on the closing prices alone¹.

Our argument for using realised correlation within a time-series model is not new, and has been applied by Corsi and Audrino (2007) in their multivariate HAR model for conditional correlation matrices, although it is the first occasion we are aware of where the argument has been applied in the context of equicorrelation. In order to implement our modification, we first require a method of calculating the realised equicorrelation. We utilise one method that relies on univariate RV estimates alone, and a method that directly estimates the entire covariance matrix through multivariate realised covariance; we begin, with the univariate case.

The RV of an asset is well known as a superior measure of latent volatility relative to alternatives such as daily squared returns and is calculated by the sum of squared intraday returns, the RV of an asset on day t is given by

$$RV_t^{(m)} \equiv \sum_{\tau=1}^m r_{\tau,m,t}^2, \quad \tau = 1, \dots, m \quad (11)$$

where $r_{\tau,m,t}^2$ is the squared intraday log-return from period $\tau - 1$ to τ for each of the m fixed-length² periods within day t . Noting that one may calculate the RV of an index as well as the RVs for each of the index's constituents, a realised measure of equicorrelation may be constructed by using the portfolio variance identity

$$\sigma^2 = \sum_{i=1}^N w_i^2 s_i^2 + 2 \sum_{i=1}^{N-1} \sum_{j>i}^N w_i w_j s_i s_j \rho_{i,j} \quad (12)$$

¹See Andersen, Bollerslev, Christoffersen, and Diebold (2006) for a review.

²Based on the research of Hansen and Lunde (2006) and related articles, we calculate the RV based on 5-minute intervals; 1-, 15-, and 30-minute intervals are also used for robustness with no qualitative effect on our results.

and making the assumption of equicorrelation, we can then re-arrange to yield

$$DREC = \frac{\sigma^2 - \sum_{i=1}^N w_i^2 s_i^2}{2 \sum_{i=1}^{N-1} \sum_{j>i}^N w_i w_j s_i s_j} \quad (13)$$

where σ^2 is the RV of the index, w_i the weight in that index placed on asset i , and s_i^2 is the RV of asset i . This realised equicorrelation measure is defined here as DREC, where the D denotes that it only uses the individual RVs, which are the diagonal elements of a covariance matrix. By using just the individual RVs to estimate the equicorrelation, a potential benefit of the DREC measure is that we do not encounter the Epps effect (Epps, 1979), whereby estimates of covariance may be biased downwards due to asynchronous trading. We then propose our first amendment to the LDECO model by replacing the u_t variable in (10) with the DREC measure,

$$\rho_t = \omega + \alpha DREC_{t-1} + \beta \rho_{t-1} \quad (14)$$

Alternatively, the equicorrelation may be estimated by using the entire covariance matrix constructed from multivariate realised measures. Suppressing the day t subscripts for notational convenience, the realised covariance (RCOV) on a given day may be calculated as follows³. Within a trading day we obtain a $N \times 1$ vector of stock returns for the assets $i = 1, \dots, N$ over each trading period, \mathbf{r}_q , where $q = 1, \dots, Q$ in a day containing Q trading periods, which is distinct and may differ from the m periods in the univariate case. The length of each of these Q periods is allowed to vary such that all of the asset returns for a given period are non-zero, with the restriction that the minimum window length is 15-minutes to avoid any market microstructure effects⁴. The realised covariance matrix relating to the trading portion on a given day is the sum of the products of these vectors,

$$RCOV = \sum_{q=1}^Q \mathbf{r}_q \mathbf{r}_q' \quad (15)$$

$$q = \min_q r_{i,q} \neq 0, \forall i$$

where q is the minimum length of time such that all returns are non-zero. The realised equicorrelation (REC) on a given day is the mean of the off-diagonal elements of (15)

$$REC = \frac{2}{N(N-1)} \sum_{i>j} \frac{RCOV_{i,j}}{\sqrt{RCOV_{i,i} RCOV_{j,j}}} \quad (16)$$

³This is the methodology is similar to that used in Laurent, Rombouts, and Violante (2010) and is based on results from Andersen, Bollerslev, Diebold, and Labys (2003) and Barndorff-Nielson and Shephard (2004).

⁴The choice of 15 minutes is partially motivated by the findings of Sheppard (2006) who finds that a minimum length of 10 minutes is sufficient to get unbiased estimates of the correlation between assets.

The alternative LDECO model we propose, LDECO-REC, is specified as

$$\rho_t = \omega + \alpha REC_{t-1} + \beta \rho_{t-1} \quad (17)$$

We also make use of another non-parametric measure of the equicorrelation between portfolio constituents on a given day that similarly utilises intraday data, the Spearman rank correlation⁵. The Spearman rank correlation between two assets i and j , $SR_{i,j}$, is calculated from 15-minute log-returns

$$SR_{i,j} = 1 - \frac{6 \sum_n^{Nt} d_n^2}{Nt(Nt - 1)} \quad (18)$$

where d_n is the difference in rankings of returns for period n for each of the Nt 15-minute periods in a day. For use in our equicorrelation problem, we again take the mean of the off-diagonal elements of this matrix of Spearman rank correlations,

$$SREC = \frac{2}{N(N-1)} \sum_{i>j} \frac{SR_{i,j}}{\sqrt{SR_{i,i}SR_{j,j}}} \quad (19)$$

Which may be substituted into the LDECO specification to yield the LDECO-SR model

$$\rho_t = \omega + \alpha SREC_{t-1} + \beta \rho_{t-1} \quad (20)$$

2.2 Incorporating Implied Equicorrelation

As well as proposing the use of alternative measures of equicorrelation, we investigate whether the information implicit in options markets is able to offer forecasting power above that contained in historical returns alone. Our motivation for this extension to LDECO again comes from the univariate volatility literature, where numerous papers demonstrate the advantages of implied volatility in generating forecasts; see Poon and Granger (2003) for a review and Becker, Clements, and Coleman-Fenn (2009) for more recent developments. In their paper introducing the LDECO model, Engle and Kelly (2008) calculate the Dow Jones Industrial Average (DJIA) IC and show that it closely matches the fitted equicorrelation from both the DCC-DECO and LDECO models, although they do not use the IC directly in model estimation or forecasting. Further, the use of IC for correlation modelling has been used previously by Castren and Mazzotta (2005) in the bivariate setting of exchange rates and they find that a combination forecast of IC and a MGARCH model is preferred, these conclusions are based on the in-sample adjusted R^2 values and they do not conduct a forecasting exercise. Encouraged by these results and the publishing of IC for the S&P 500 Index on the CBOE, we aim to adapt the methodology employed in the univariate context to the equicorrelation forecasting problem.

For a market index on which options trade, the DJIA for example, it is well known that a model-free estimate of the implied volatility of the index as

⁵We thank Andrew Harvey for suggesting this measure.

a whole can be constructed, i.e. the VXD ⁶. For each one of the constituent stocks of that index on which options trade, we can similarly determine its model-free implied volatility. Again relying on the portfolio variance identity in (12) and making the assumption of equicorrelation, we can re-arrange the portfolio variance identity to yield

$$IC = \frac{\sigma^2 - \sum_{j=1}^N w_j^2 s_j^2}{2 \sum_{i=1}^{N-1} \sum_{j>i}^N w_i w_j s_i s_j} \quad (21)$$

where σ is now the annualised *implied* 22-day-ahead standard deviation of the index, w_i is the portfolio weight given to asset i , and s_i is the annualised *implied* 22-day-ahead standard deviation of asset i .

Following Blair, Poon and Taylor (2001) who add the IV to a univariate GARCH model of volatility, we propose extending the LDECO specification to include the IC as an exogenous variable in forecasting equicorrelation. This combines the historical information available from the returns series of the portfolio constituents with the information implied by the options market,

$$\rho_t = \omega + \alpha u_{t-1} + \beta \rho_{t-1} + \gamma IC_{t-1} \quad (22)$$

Our model retains the attractive property of analytical solutions to the inverse and determinant of R_t as given by (7) and (8), respectively, as equicorrelation is assumed in calculating IC and we simply take a linear combination of two equicorrelation measures.

In a fully efficient options market, the co-efficient on u_t is expected to be statistically indistinguishable from zero as the historical time-series information should be incorporated by options market participants in generating their IC forecast. It is worth noting that even if we restrict our model to utilise data from the options market only, we would not expect a co-efficient of unity for γ as, similar to the volatility literature, one might expect a correlation risk premium. Hence, even if the options market generates perfect forecasts through the IC measure, we would not expect the results typical of a Mincer-Zarnowitz (1969) regression of zero as a constant and unity as a co-efficient.

Finally, we can combine both alternative measures of daily equicorrelation with the IC from options market into the following LDECO-DREC-IC, LDECO-REC-IC and LDECO-SR-IC models respectively

$$\rho_t = \omega + \alpha DREC_{t-1} + \beta \rho_{t-1} + \gamma IC_{t-1} \quad (23)$$

$$\rho_t = \omega + \alpha REC_{t-1} + \beta \rho_{t-1} + \gamma IC_{t-1} \quad (24)$$

$$\rho_t = \omega + \alpha SREC_{t-1} + \beta \rho_{t-1} + \gamma IC_{t-1} \quad (25)$$

⁶The VXD is the DJIA equivalent of the potentially more well known VIX for the S&P 500 Index; a model-free, risk-neutral, option implied forecast of the mean annualised volatility of the index over a fixed 22 trading day horizon.

3 Forecast Evaluation

In this Section, we detail the procedure by which point forecasts of correlation are generated and also describe the Model Confidence Set methodology for comparing the statistical performance of the respective forecasts.

3.1 Generating Forecasts

As well as an in-sample comparison of the log-likelihood values and parameter estimates of the models considered, we also generate multi-step-ahead point forecasts of equicorrelation up to the 22-day horizon over which the IC is defined. Unlike variance and covariance, however, one cannot aggregate correlation through time and each point forecast must be evaluated individually, rather than the total 22-day correlation. That is, we evaluate the forecasting performance of each of the models for each k -day ahead forecast, up to $k = 22$ days.

To generate a multi-period forecast, we need to assume that $\mathbb{E}_t[u_{t+k}] \approx \mathbb{E}_t[\rho_{t+k}]$, and similarly for alternative equicorrelation measures, which can then be utilised to generate recursive forecasts,

$$\mathbb{E}_t[\rho_{t+k}] = \omega + (\alpha + \beta)\mathbb{E}_t[\rho_{t+k-1}] + \gamma\mathbb{E}_t[IC_{t+k-1}] \quad (26)$$

It is apparent that we fall into difficulty with forecasting the implied equicorrelation forward through time and we have no *a priori* guidance as to its dynamics. A simple solution is to estimate a standard time-series model, namely an Auto-Regressive model of order one, AR(1). Under such an assumption, the multi-period forecast of IC is given trivially by

$$\mathbb{E}_t[IC_{t+K}] = \theta_1^K IC_t + \mu(1 - \theta_1^K) \quad (27)$$

Using recursive substitution of (26) and substituting in (27) leads to an equivalent multi-step-ahead forecast,

$$\begin{aligned} \rho_{t+K} &= \omega \left[\frac{1 - (\alpha + \beta)^{(K-1)}}{1 - \alpha - \beta} \right] + (\alpha + \beta)^{K-1} \rho_{t+1} \\ &+ \gamma \sum_{k=0}^{K-2} (\alpha + \beta)^{K-2-k} \left[\mu \left[\frac{1 - \theta^{(k+1)}}{1 - \theta} \right] + \theta^{(k+1)} IC_t \right] \end{aligned}$$

where μ is the drift term in the AR(1) process and θ_1 is the co-efficient acting on the lagged value of IC_t .

We also generate equicorrelation forecasts from the cDCC model of Aielli (2009). We believe this provides a natural alternative to those models specifically designed for equicorrelation forecasting. One is able to forecast the conditional correlation matrix forward without the assumption of equicorrelation already imposed, and then simply take the mean of the off-diagonal elements of this conditional correlation matrix and utilise that scalar value as the equicorrelation forecast. This forecast is then compared with those models that impose the equicorrelation assumption in their estimation process.

3.2 Statistical Evaluation of Forecasts

In order to statistically evaluate the relative forecast performance of the models considered, we require a measure of the “true” equicorrelation on each of the days for which point forecasts are generated in order to gauge their accuracy. We have already argued in favour of realised equicorrelation as a superior measure of the daily relationship between assets of interest when constructing the REC measure for use within the LDECO model and the majority of our results will be discussed with this measure in mind. However, as a robustness check, we also utilise the Engle and Kelly (2008) measure of equicorrelation defined in (9), u_t , the diagonal realised equicorrelation (DREC) defined in 13, and the Spearman rank equicorrelation defined in (18), $SREC$, as the “true” equicorrelation values.

We employ the Model Confidence Set (MCS) approach to examine the forecast performance of each of the models considered. At the heart of the methodology (Hansen, Lunde and Nason, 2003, 2010) as it is applied here, is a forecast loss measure. While there are many options available, the loss functions utilized here are mean-square-error (MSE) and QLIKE,

$$MSE_k^{\mathcal{M}} = (\rho_{t+k} - f_{t,k}^i)^2, \quad (28)$$

$$QLIKE_k^{\mathcal{M}} = \log(f_{t,k}^i) + \frac{\rho_{t+k}}{f_{t,k}^i}, \quad (29)$$

where $f_{t,k}^{\mathcal{M}}$ are individual forecasts (formed at time t for k -days ahead) obtained from the individual models, \mathcal{M} , and ρ_{t+k} is the measure of true equicorrelation.

While these loss functions allow forecasts to be ranked, they give no indication of whether the top performing model is statistically superior to any of the lower-ranked models; the MCS approach allows for such conclusions to be drawn. The construction of a MCS is an iterative procedure that requires sequential testing of equal predictive accuracy (EPA); the set of candidate models is trimmed by deleting models that are found to be statistically inferior. The interpretation attached to a MCS is that it contains the best forecasting models which are of EPA with a given level of confidence. We defer to Hansen, Lunde and Nason (2003, 2010) for any further detail required regarding the implementation of the MCS methodology.

4 Data

Our results are based on the DJIA over the period starting on the 1st of November 2001 through to the 30th of October 2009, leaving us with 1964 observations⁷. After allowing for a 1000-trading-day initial estimation window, we are able generate 963 out-of-sample forecasts. We utilise three data sources,

⁷We choose the DJIA as we are able to obtain the implied volatilities of each of its constituent stocks for each day of the sample and are therefore able to calculate the implied equicorrelation with certainty. This is not true of the S&P 500 Index, for which the CBOE publishes its IC based on an approximation from the largest 50 stocks within the index, as not all of its constituent stocks have listed options traded.

the OptionsMetrics IvyDB US database in calculating implied volatilities for individual stocks, the CBOE for the daily closing values of the VXD index, and ThomsonReuters Tick History for minute-by-minute intraday prices used in calculating the realised equicorrelation measures.

Similar to the more commonly known VIX for the S&P 500 Index, the VXD is a model-free 22-day-ahead at-the-money implied volatility forecast for the DJIA. To fix ideas, the day t implied equicorrelation is given by

$$IC_t = \frac{VXD_t^2 - \sum_{j=1}^n w_{j,t}^2 s_{j,t}^2}{2 \sum_{i=1}^{N-1} \sum_{j>i}^N w_{i,t} w_{j,t} s_{i,t} s_{j,t}} \quad (30)$$

where the weights and standard deviations now have a t subscript to denote that the constituents of the index vary through time ⁸.

For comparative purposes we now plot each of the five equicorrelation measures utilised in this paper, the measure proposed by Engle and Kelly (2008), the implied equicorrelation, the realised diagonal equicorrelation, the realised equicorrelation and the Spearman rank equicorrelation for the sample period investigated here in Figure 1.

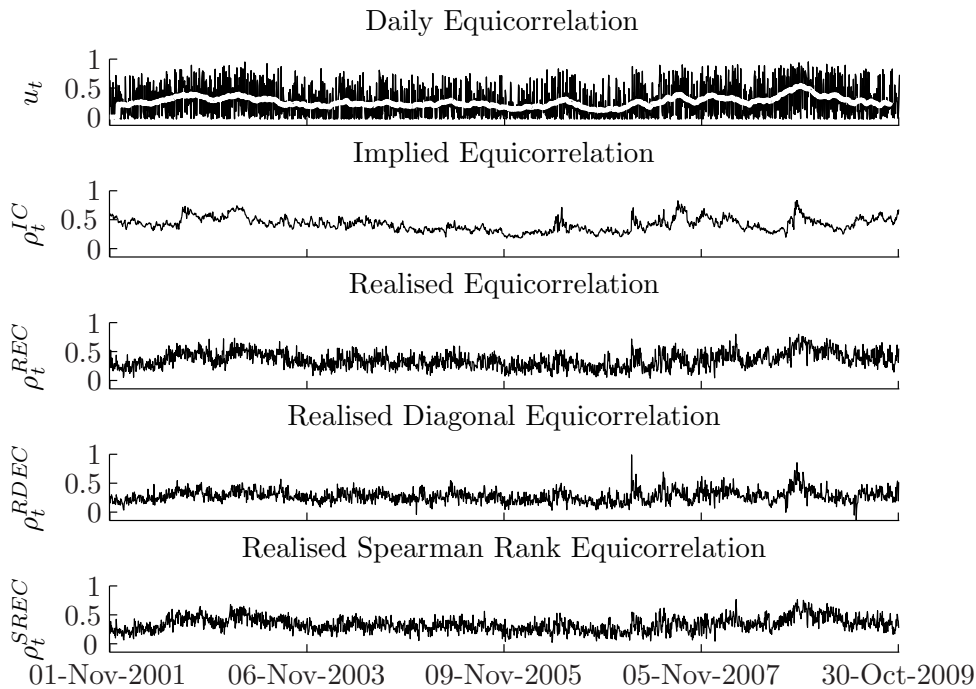


Figure 1: Plots of each of the five equicorrelation measures utilised in this paper for the full sample period of 1st November 2001 through to 30th of October 2009.

From the plot of u_t in Panel A of Figure 1, it may be observed that the measure proposed by Engle and Kelly (2008) is a noisy measure, perhaps more noisy

⁸Although the DJIA is relatively more stable than, say, the S&P 100, only 17 of the original 30 constituents remain in the index consistently for our sample period.

than one would expect of the mean correlation of thirty of the largest US firms. However, we also plot a centered 44-day-moving-average, the mean equicorrelation of data one month either side of a given day, in white to demonstrate that the measure is still quite persistent.

As can be seen in Panel B of Figure 1, the IC of options markets appears significantly less noisy than the other measures. It also appears to track the realised measures quite closely, which augurs well for the out-of-sample forecasting exercise given these measures are used as alternative “true” equicorrelation proxies. Further anecdotal support for the use of IC in correlation forecasting comes from the fact that the IC_t tends to peak in times of market turmoil; when large indexes fall, the majority of assets suffer losses and this is reflected in a high level of correlation across assets.

In Panels B, C and D of Figure 1 we plot the realised equicorrelation measures and the Spearman rank equicorrelation over the sample period, it may be observed that the measures follow similar dynamics although the diagonal realised equicorrelation is the least noisy of the three, this may suggest that it will possess more power in separating the out-of-sample forecast performance of the competing models.

To reinforce our point regarding the noisiness of the respective equicorrelation measures, we provide some descriptive statistics for each of the series in Table 1. It can be seen that the u_t measure is extremely noisy, with its standard deviation of 0.2681 larger than its mean of 0.2631; the u_t measure is also weakly correlated with all of the alternative measures. The standard deviation of the other four measures are significantly smaller, with all falling between 0.105 and 0.135, and they are more highly correlated with each other. The two measures of realised equicorrelation, DREC and REC, are somewhat different in their means at 0.3556 and 0.2782 respectively, with the standard deviation of the DREC measure slightly smaller, they are unsurprisingly highly correlated with each other at 0.7237. The mean of the IC measure, at 0.4218, is higher than all of the other measures of equicorrelation and probably reflects a correlation risk premium being priced in the derivatives market. Finally, it should be noted that our measures of equicorrelation from the physical market are not that highly correlated with the IC from the options market, so there should not be any adverse effects from multicollinearity by including multiple measures of equicorrelation.

5 Results

This Section begins with a summary of the in-sample performance of the models considered before analysing the out-of-sample forecast performance.

5.1 In-Sample Estimation Results

We begin with a comparison of the in-sample fit of our modified LDECO models to the standard specification and present the results in Table 2; all proposed

Descriptive Statistics of Equicorrelation Measures

Measure	Mean	Std	ρ_{u_t}	ρ_{IC_t}	ρ_{REC_t}	ρ_{DREC_t}	ρ_{SREC_t}
u_t	0.2631	0.2681	1	0.2075	0.2112	0.1624	0.1680
IC_t	0.4218	0.1152	0.2075	1	0.4649	0.5047	0.4429
REC_t	0.3556	0.1347	0.2112	0.4649	1	0.7237	0.6932
$DREC_t$	0.2782	0.1054	0.1624	0.5047	0.7237	1	0.5523
$SREC_t$	0.3363	0.1221	0.1680	0.4429	0.6932	0.5523	1

Table 1: Descriptive statistics over the full-sample period of the four equicorrelation measures utilised in this paper. The correlation statistic is of the measure for that row with the measure in the column header.

amendments result in a superior log-likelihood function value to the standard LDECO model, even a simple univariate model of IC without a persistence term, $\rho_t = \omega + \gamma IC_t$, outperforms LDECO. As all three of our realised equicorrelation measures provide superior in-sample fits relative to the u_t measure, our results reinforce earlier theoretical and empirical research that utilising high-frequency intraday data more accurately measures a latent variable than daily closing returns.

Further, the IC from options markets appears to subsume any information regarding daily equicorrelation provided by the Engle and Kelly (2008) measure as the co-efficient acting on the innovation term, α , is statistically insignificant. Again, this is similar to the majority of research surveyed by Poon and Granger (2003) where models utilising IV typically outperformed daily-returns based models. Regarding the models which incorporate both intraday data and options market implied information, measures from both markets have statistically significant power even when utilised in combination for three of the four measures; the co-efficient on IC is not statistically significant in the LDECO-DREC-IC model.

For a visual comparison, we plot the fitted equicorrelation, ρ_t , for the standard LDECO specification and some of our variations⁹ in Figure 2. It can be seen that the extensions of LDECO generate relatively similar in-sample fits of ρ_t , even LDECO-IC and LDECO-REC, which derive their measures of equicorrelation quite differently, follow each other closely. It is apparent that, given

⁹We exclude the univariate model of IC as it is obvious from the results in Table 2 that a persistence term is highly significant. Further, we exclude the models incorporating the SREC measure as they are qualitatively similar to the REC models.

the higher log-likelihood of all of its competitors reported in Table 2, the fitted ρ_t of LDECO appears to miss variations in the object of interest. While obviously not a constant conditional correlation model, this apparent inability to follow prevailing market conditions does not bode well for its out-of-sample forecast performance, though it seems that any of its extensions may generate fairly similar forecasts to each of the other alternatives.

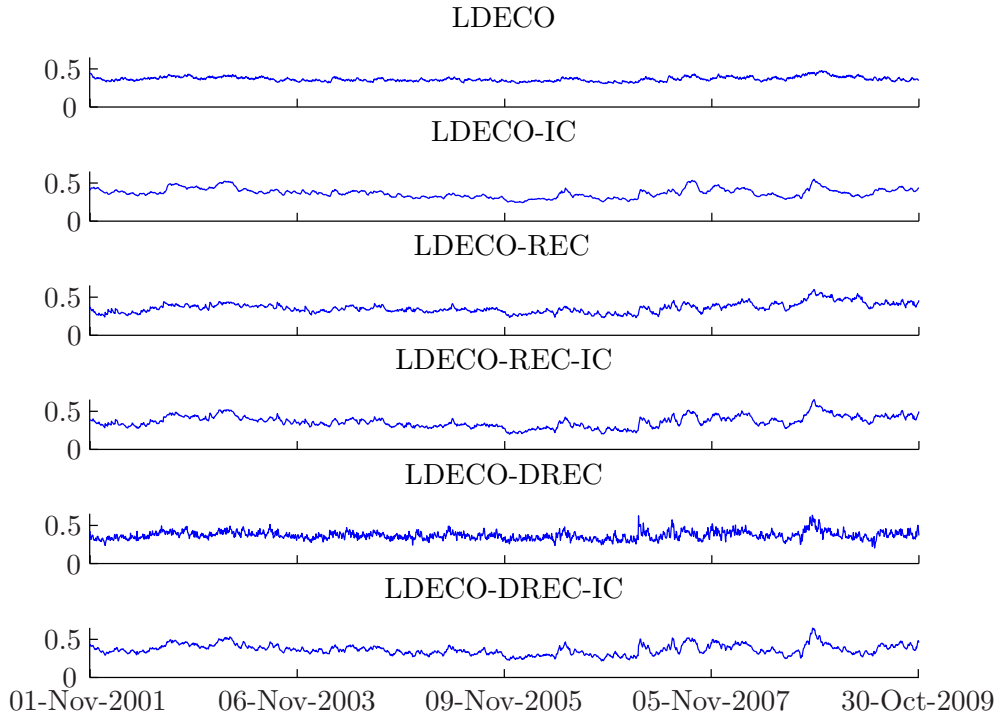


Figure 2: The in-sample fitted equicorrelations, ρ_t , of four of the competing models.

To summarise our in-sample results, all of the adaptations of LDECO proposed in this paper generate superior in-sample performance, indicating that intraday measures of equicorrelation are superior to the u_t measure of Engle and Kelly (2008) and that IC provides information incremental to that contained in the physical market.

5.2 Multi-Step-Ahead Forecast Results

We now turn our attention to the out-of-sample forecasts of equicorrelation. In Table 3 we summarise the MCS results under the Mean Square Error loss function for the range statistic using the REC as our measure of “true” equicorrelation¹⁰; two main observations may be made. Firstly, the LDECO-REC model

¹⁰Results are qualitatively similar for both the QLIKE and MSE loss functions under both the range and semi-quadratic test statistics as well as when the DREC, SREC and u_t measures are used as the “true” equicorrelation. Hence, for brevity’s sake we only present the results for one of the sub-set of results with the remainder available upon request.

is almost universally the best out-of-sample forecasting method, it outperforms all its competitors at every forecast horizon considered under both loss functions and test statistics; however, there is one exception to this case that is discussed shortly. Secondly, as the forecast horizon increases, it becomes increasingly difficult to statistically distinguish between the competing models. Beyond the 5 day forecast horizon, the only models to be excluded under any loss function, test statistic of measure of equicorrelation are the DCC, LDECO, LDECO-DREC, and LDECO-REC-IC models. That is, those models that use either daily data or the diagonal realised equicorrelation

We note that in our robustness checks of using four measures of “true” equicorrelation, there is one exception to the LDECO-REC providing the best forecast. At the one-day horizon under both loss functions and test statistics, the LDECO-DREC model provides the best forecast under the DREC measure of equicorrelation. That is, the LDECO-DREC model is superior at producing a one-step-ahead forecast of DREC, a perhaps unsurprising result. We believe the fact that for all other loss functions, tests statistics, time horizons, and measures of equicorrelation the LDECO-REC result provides the superior out-of-sample forecast is a highly robust result.

It is important to note that all of the models that rely on daily closing price returns for their equicorrelation measure are typically the worst performing, even taking the mean of the forecast correlation matrix from the DCC model results in poor forecasting performance; a simple linear regression on IC typically yields superior forecasts to those models utilising daily returns measures. Further, the inclusion of the IC to form the LDECO-IC model also generally outperforms the simple LDECO model.

However, when examining those models that incorporate realised measures of equicorrelation (the DREC, REC or SREC measures), the inclusion of the IC measure typically worsens the out-of-sample forecasting performance. This is perhaps due to in-sample over-fitting, or that the chosen AR(1) specification for IC does not adequately match the true dynamics of IC. The fact that LDECO-REC dominates at all time horizons suggests that even an improved forecasting method for IC would not reverse the rankings of the models. Perhaps if LDECO-REC-IC dominated for shorter horizons before LDECO-REC became the superior forecasting tool, a more thorough investigation of the dynamics of IC would be warranted, but we do not believe that is the case here. In either scenario, the superior in-sample fit by including IC is not replicated in the out-of-sample forecasting exercise, where information from the physical market alone generates the best forecasts.

6 Conclusions

The concept of equicorrelation has been of increasing interest in the financial economics and econometrics literature. We analyse the in-sample fit and out-of-sample forecasting performance of ten candidate models of equicorrelation after adapting the LDECO model to utilise improved measures of equicorrelation provided by the realised covariance literature and extend LDECO to

include a measure of equicorrelation implied by the options market as an exogenous regressor. We find for our in-sample results that the information in implied equicorrelation subsumes the equicorrelation measure proposed by Engle and Kelly (2008) and also provides information incremental to the intraday equicorrelation measures.

In our out-of-sample multi-period forecast results, the standard DCC and LDECO models are typically the worst performing with any of the specifications that replace a daily measure of equicorrelation with a realised measure almost universally generating superior out-of-sample forecasts, even a simple linear regression on implied correlation outperforms the DCC and LDECO models. However, while the implied equicorrelation is generally significant in-sample, it does not yield superior forecasting power when included with specifications that also include a realised measure of equicorrelation. Further, the realised equicorrelation statistic typically generates superior forecasts to the other non-parametric measures, the diagonal realised equicorrelation and the Spearman rank equicorrelation. These results generally hold across both loss functions, test statistics, forecast horizons, and the alternative “true” measures of equicorrelation; the exception is under the diagonal realised equicorrelation measure at the one day horizon.

These results reinforce empirical work in the univariate volatility literature that realised measures of a latent variable are superior to measures reliant upon daily closing prices alone. Further, the information implicit in options markets also generates forecasts superior to the standard LDECO model which utilises daily returns as its equicorrelation measure. However, the implied equicorrelation does not provide information incremental to the realised measure of equicorrelation for the purposes of forecasting.

In-Sample Results

Model	ω	α	β	γ	L_{Corr}
LDECO	0.0287 (0.0067)	0.0339 (0.0134)	0.8975 (0.0229)	- -	-18.9682
LDECO-IC	0.0250 (0.0098)	0.0141 (0.0200)	0.8090 (0.0386)	0.0996 (0.0307)	-18.8847
REC	0.0161 (0.0199)	0.1164 (0.0348)	0.8404 (0.0628)	- -	-18.8694
REC-IC	0.0052 (0.0106)	0.1020 (0.0286)	0.8008 (0.0494)	0.0743 (0.0282)	-18.8235
DREC	0.1375 (0.0657)	0.3876 (0.1378)	0.3354 (0.2661)	- -	-18.8827
DREC-IC	0.0261 (0.0340)	0.1278 (0.1396)	0.7069 (0.2551)	0.1089 (0.0625)	-18.8480
SREC	0.0188 (0.0060)	0.0581 (0.0136)	0.8946 (0.0206)	- -	-18.9211
SREC-IC	0.0140 (0.0085)	0.0505 (0.0155)	0.8313 (0.0351)	0.0735 (0.0280)	-18.8653
$\rho_t = \omega + \gamma IC_t$	0.1917 (0.0407)	- -	- -	0.4234 (0.0943)	-18.9229

Table 2: In-sample parameter estimates, robust standard errors in parentheses, and log-likelihood values for the five competing models.

Horizon	Mean Square Error Tr p -values of Model Confidence Set									
	DCC	LDECO	LDECO-IC	REC	REC-IC	SREC	SREC-IC	RC-DIAG	RC-DIAG-IC	IC
ρ_{t+1}				✓✓✓				✓✓✓	✓✓✓	
ρ_{t+2}				✓✓✓						
ρ_{t+3}				✓✓✓						
ρ_{t+4}				✓✓✓	✓					
ρ_{t+5}	✓✓	✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓	✓✓	✓✓✓
ρ_{t+6}		✓	✓✓	✓✓✓	✓✓	✓✓	✓		✓	✓✓
ρ_{t+7}	✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓	✓✓✓
ρ_{t+8}	✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓	✓✓	✓✓✓
ρ_{t+9}	✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓	✓✓✓
ρ_{t+10}	✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓	✓✓	✓✓✓
ρ_{t+11}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+12}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+13}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+14}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+15}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+16}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+17}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+18}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+19}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+20}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+21}	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
ρ_{t+22}	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓	✓✓

Table 3: Summary of p -values of the Model Confidence Set using the Mean-Square-Error Loss function under the range statistic when the *realised equicorrelation* is taken as the measure of “true” equicorrelation. ✓ indicates a p -value between 0.05 and 0.10, ✓✓ between 0.10 and 0.20, and ✓✓✓ greater than 0.20. The null hypothesis is that models are of equal predictive accuracy. Hence, a p -value of 0.05 implies a 5% probability of committing a Type I error if we exclude the model from the MCS.

References

- Aielli, G.P., Dynamic Conditional Correlations: On Properties and Estimation, 2009. SSRN Working Paper, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1507743
- Andersen, T. G., Bollerslev, T., Christoffersen, P. F., and Diebold, F. X., Volatility and Correlation Forecasting, 2006. In G. Elliot, C. Granger, and A. Timmermann, *Handbook of Economic Forecasting*, North Holland Press.
- Andersen, T., Bollerslev, T., Diebold, F. X., and Labys, P., (2003): "Modeling and Forecasting Realized Volatility", *Econometrica*, **71**, 579625.
- Barndorff-Nielson, O. E., Hansen, P. R., Lunde, A., Shephard, N., Designing Realized Kernels to Measure the ex post Variation of Equity Prices in the Presence of Noise, 2008. *Econometrica* 76-(6), 1481-1536.
- Barndorff-Nielson, O. E., Hansen, P. R., Lunde, A., Shephard, N., Multivariate Realised Kernels: Consistent Positive Semi-Definite Estimators of the Covariation of Equity Prices with Noise and Non-Synchronous Trading, Working Paper, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1154144
- Barndorff-Nielson, O. E., Shephard, N., Econometric analysis of realized volatility and its use in estimating stochastic volatility models, 2002. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64-(2), 253-280.
- Barndorff-Nielson, O. E., Shephard, N., Estimating quadratic variation using realized variance, 2002. *Journal of Applied Econometrics* 17-(5), 457-477.
- Barndorff-Nielson, O. E., Shephard, N., Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics, 2004. *Econometrica* 72-(3), 885-925.
- Bauwens, L., Laurent, S., Rombouts, J. V. K., Multivariate GARCH Models: A Survey, 2006. *Journal of Applied Econometrics* 21, 79-109.
- Becker, R., Clements, A., Are combination forecasts of S&P 500 volatility statistically superior?, 2008. *International Journal of Forecasting* 24, 122-133.
- Becker, R., Clements A. E., Coleman-Fenn, C. A., Forecast performance of implied volatility and the impact of the volatility risk premium, 2009. NCER Working Paper Series, 45. <http://www.ncer.edu.au/papers/documents/WPNo45.pdf>
- Blair, B. J., Poon, S-H., Taylor, S. J., Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns, 2001. *Journal of Econometrics* 105, 5-26.
- Bollerslev, T., Modelling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH model, 1990. *Review of Economics and Statistics* 72, 498-505.

- Bollerslev, T., Engle, R. F., Wooldridge, J. M., A Capital Asset Pricing Model with Time-Varying Covariances, 1988. *Journal of Political Economy* 96-(1), 116-131.
- Castren, O., Mazzotta, S., Foreign Exchange Option and Returns Based Correlation Forecasts: Evaluation and Two Applications, 2005. European Central Bank Working Paper Series, 447. <http://www.ecb.int/pub/pdf/scpwps/ecbwp447.pdf>
- Chicago Board of Exchange, CBOE S&P 500 Implied Correlation Index, 2009. <http://www.cboe.com/micro/IMPLIEDCORRELATION/IMPLIEDCORRELATIONINDICATOR.pdf>
- Corsi, F., A Simple Approximate Long-Memory Model of Realized Volatility, 2009. *Journal of Financial Econometrics* 7-(2), 174-196.
- Corsi, F., Audrino, F., Realized Correlation Tick-by-Tick, 2007. Universität St. Gallen Discussion Paper 2007-02, <http://www.vwa.unisg.ch/RePEc/usg/dp2007/DP02-Au.pdf>
- Driessen, J., Maenhout, P., and Vilkov, G., Option Implied Correlations and Correlation Risk, 2009. *Journal of Finance* 64-(3), 1061-1519.
- Elton, E. J., Gruber, M. J., Estimating the Dependence Structure of Share Prices-Implications for Portfolio Selection, 1973. *Journal of Finance* 28, 1203-1232.
- Engle, R. F., Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models, 2002. *Journal of Business and Economic Statistics* 20-(3), 339-350
- Engle, R. F., Kelly, B. T., Dynamic Equicorrelation, 2008. NYU Working Paper, <http://pages.stern.nyu.edu/~rengle/Dynamic%20Equicorrelation.pdf>
- Engle, R. F., Kelly, B. T., Dynamic Equicorrelation, 2009. SSRN Working Paper, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1354525
- Epps, T. W., (1979) "Comovements in Stock Prices in the Very Short Run", *Journal of the American Statistical Association*, **74**, 291-298.
- Hansen, P. R., Lunde, A., Realized Variance and Market Microstructure Noise, 2006. *Journal of Business and Economic Statistics* 24, 127-218
- Hansen, P. R., Lunde, A., Nason, J. M., (2003). "Choosing the best volatility models: The model confidence set approach", *Oxford Bulletin of Economics and Statistics* 65, 839-861.
- Hansen, P. R., Lunde, A., Nason, J. M., The Model Confidence Set, 2010. SSRN Working Paper, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=522382
- Jorion, P., Predicting Volatility in the Foreign Exchange Market, 1995. *Journal of Finance* 50, 507-528.

Laurent, S., Rombouts, J. V. K., and Violante, F., On the Forecasting Accuracy of Multivariate GARCH Models, 2010. CIRPEE Working paper, <http://www.cirpee.org/fileadmin/documents/Cahiers.2010/CIRPEE10-21.pdf>

Mincer, J., Zarnowitz, V., The evaluation of economic forecasts, 1969. In *Economic Forecasts and Expectations*, Columbia University Press.

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

Poon S-H., Granger C. W. J., Forecasting volatility in financial markets: a review, 2003. *Journal of Economic Literature* 41, 478-539.

Sheppard, K., (2006), "Realized Covariance and Scrambling", Working Paper, <http://www.kevinsheppard.com/images/7/7e/Scrambling.pdf>

Silvennoinen, A., Teräsvirta, T., Multivariate GARCH models, 2009. In *Handbook of Financial Time Series*, Springer.