UNIVERSITY OF
CAMBRIDGE

Cambridge Working Papers in Economics

Testing against Changing Correlation

Andrew Harvey and Stephen Thiele, Faculty of Economics, Cambridge University

CWPE 1439

Testing against Changing Correlation

Andrew Harvey and Stephen Thiele
Faculty of Economics, Cambridge University

October 31, 2014

Abstract

A test for time-varying correlation is developed within the framework of a dynamic conditional score (DCS) model for both Gaussian and Student t-distributions. The test may be interpreted as a Lagrange multiplier test and modified to allow for the estimation of models for time-varying volatility in the individual series. Unlike standard moment-based tests, the score-based test statistic includes information on the level of correlation under the null hypothesis and local power arguments indicate the benefits of doing so. A simulation study shows that the performance of the score-based test is strong relative to existing tests across a range of data generating processes. An application to the Hong Kong and South Korean equity markets shows that the new test reveals changes in correlation that are not detected by the standard moment-based test.

KEYWORDS: Dynamic conditional score, EGARCH, Lagrange multiplier test, Portmanteau test, Time-varying covariance matrices.

JEL classification: C14, C22, F36

1 Introduction

The possibility that the correlations between financial assets are changing over time is an important issue in many areas of finance, such as portfolio construction and risk management; see Lumsdaine (2009) for a recent discussion. The aim here is to provide a test for time-varying correlation that is powerful, yet simple to implement. The proposed approach is based on the dynamic conditional score (DCS) models recently developed by Creal *et al*

(2011, 2013) and Harvey (2013). It is shown that Lagrange multiplier (LM) tests can be constructed from the autocorrelations of the conditional scores, with a modified test taking account of estimated dynamic variances. Without this modification the test is based on a simple portmanteau statistic. The scores incorporate information on the level of correlation, and local power arguments indicate that the resulting test can be expected to be more powerful as the level of correlation under the null hypothesis moves away from zero. This is not the case with the standard moment-based portmanteau test, introduced by Bollerslev (1990), which simply uses the cross-product of standardised residuals.

The tests are developed for a bivariate Gaussian model, with a subsequent extension to the bivariate Student t-distribution. Monte Carlo experiments are used to compare the performance of these tests with existing tests, including those of Tse (2000, 2002) and Bera and Kim (2002). The results show that, on the whole, the proposed tests perform much better than existing tests across a range of data generating processes. Although the competing tests, which include portmanteau tests, residual regression tests and Lagrange multiplier tests, are based on a variety of approaches, they generally rely on the cross-product of standardised residuals to identify potential time variation and so share the same weakness relative to the scores. This point is highlighted by an application to the Hong Kong and South Korean equity markets, where it is found that the score-based tests can identify changing correlations that are undetectable by a moment-based test.

The paper is organised as follows. Section 2 reviews the bivariate DCS model for time-varying correlation and Section 3 shows how the new tests can be derived as LM tests within this framework. Section 4 presents the Monte Carlo results while Section 5 reports the application.

2 The DCS Model for Time-Varying Correlation

Consider a bivariate Gaussian model, with zero means and constant variances, in which the covariance matrix is

$$\mathbf{\Sigma}_{t|t-1} = \left[\begin{array}{ll} \sigma_1^2 & \rho_{t|t-1}\sigma_1\sigma_2 \\ \rho_{t|t-1}\sigma_1\sigma_2 & \sigma_2^2 \end{array} \right],$$

where $\rho_{t|t-1}$ denotes the (changing) correlation, based on information at time $t-1$. Rather than working directly with $\rho_{t|t-1}$, a transformation is applied so as to keep it in the range, $-1 < \rho_{t|t-1} < 1$. The link function

$$\rho_{t|t-1} = \frac{\exp(2\gamma_{t|t-1}) - 1}{\exp(2\gamma_{t|t-1}) + 1}, \quad t = 2, ..., T, \tag{1}$$

is eminently suitable in that it allows the new variable, $\gamma_{t|t-1}$, to be unconstrained. The inverse is the arctanh transformation.

The log-density of the $t-th$ observation, conditional on information at time $t-1$, is

$$\begin{aligned} \ln f(y_t; \boldsymbol{\psi}, \lambda_1, \lambda_2) &= -\ln 2\pi - \lambda_1 - \lambda_2 - \frac{1}{2}\ln(1 - \rho_{t|t-1}^2) \\ &\quad - \frac{1}{2(1 - \rho_{t|t-1}^2)} \left(\frac{{y_{1t}}^2}{\exp(2\lambda_1)} - \frac{2\rho y_{1t} y_{2t}}{\exp(\lambda_1 + \lambda_2)} + \frac{{y_{2t}}^2}{\exp(2\lambda_2)} \right), \end{aligned}$$

where $\boldsymbol{\psi}$ denotes the parameters upon which $\rho_{t|t-1}$, and hence $\gamma_{t|t-1}$, depend. The score with respect to $\gamma_{t|t-1}$, that is $\partial \ln f_t / \partial \gamma_{t|t-1}$, can be written in terms of $\rho_{t|t-1}$ as

$$u_t = \frac{1}{4}(x_{1t} + x_{2t})^2 \frac{1 - \rho_{t|t-1}}{1 + \rho_{t|t-1}} - \frac{1}{4}(x_{1t} - x_{2t})^2 \frac{1 + \rho_{t|t-1}}{1 - \rho_{t|t-1}} + \rho_{t|t-1}, \tag{2}$$

where $x_{it} = y_{it}\exp(-\lambda_i)$, $i = 1, 2$; see Harvey (2013, ch 7). We can also write

$$u_t = \frac{1}{1 - \rho_{t|t-1}^2} \left[(1 + \rho_{t|t-1}^2) x_{1t} x_{2t} - \rho_{t|t-1}(x_{1t}^2 + x_{2t}^2) \right] + \rho_{t|t-1}.$$

The score reduces to $x_{1t}x_{2t}$ when $\rho_{t|t-1} = \gamma_{t|t-1} = 0$, but more generally the term involving squared observations makes important modifications capturing information on the level of correlation.

The first-order dynamic equation for correlation is

$$\gamma_{t+1|t} = (1 - \phi)\omega + \phi\gamma_{t|t-1} + \kappa u_t, \qquad t = 1, ..., T, \tag{3}$$

with $\gamma_{1|0} = \omega$, where u_t is the score. It is convenient to specify the standard deviations using an exponential link function. The covariance matrix is therefore $\mathbf{\Sigma}_{t|t-1} = \mathbf{D}\mathbf{R}_{t|t-1}\mathbf{D}$, where the diagonal matrix $\mathbf{D}$ has elements $\exp(\lambda_1)$ and $\exp(\lambda_2)$ and

$$\mathbf{R}_{t|t-1} = \begin{bmatrix} 1 & \rho_{t|t-1} \\ \rho_{t|t-1} & 1 \end{bmatrix}. \tag{4}$$

When scale (standard deviation in a Gaussian model) is time varying, the dynamic equations will be assumed to take a similar form to that for $\gamma_{t|t-1}$, namely

$$\lambda_{i,t+1|t} = \omega_i(1-\phi_i) + \phi_i\lambda_{i,t|t-1} + \kappa_i u_{it}, \quad i = 1, 2, \tag{5}$$

with $\lambda_{i,1|0} = 0$, $i = 1, 2$.The exponential link function ensures that the variances remain positive.

The information matrix for λ_1, λ_2 and γ in the static model depends only on γ. In terms of ρ it is

$$\mathbf{I}\begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \gamma \end{pmatrix} = \begin{bmatrix} \frac{2-\rho^2}{1-\rho^2} & \frac{-\rho^2}{1-\rho^2} & -\rho \\ \frac{-\rho^2}{1-\rho^2} & \frac{2-\rho^2}{1-\rho^2} & -\rho \\ -\rho & -\rho & 1+\rho^2 \end{bmatrix}. \tag{6}$$

Remark 1 *If the score vector for λ_1, λ_2 and γ is pre-multiplied by the inverse of the information matrix, as is often the practice in formulating DCS models, the modified score for $\gamma_{t|t-1}$ becomes*

$$\begin{aligned} u_t &= \frac{1}{8}(x_{1t}+x_{2t})^2\frac{1-\rho_{t|t-1}}{1+\rho_{t|t-1}} - \frac{1}{8}(x_{1t}-x_{2t})^2\frac{1+\rho_{t|t-1}}{1-\rho_{t|t-1}} + \frac{1}{2}x_{1t}x_{2t} \qquad (7) \\ &= \frac{1}{1-\rho_{t|t-1}^2}\left[x_{1t}x_{2t} - \frac{\rho_{t|t-1}}{2}(x_{1t}^2+x_{2t}^2)\right]. \end{aligned}$$

In this case, the variance of u_t is unity in all time periods and so the condition $|\phi| < 1$ ensures that $\gamma_{t+1|t}$ is covariance stationary. As regards the volatility equations, (5), the $u_{it}'s$ are the same as they would be in a univariate model (apart from a factor of 1/2). In other words the score-driven approach suggests that the volatility for each series is driven solely by its own movements.

The model may be generalized so that the joint distribution is multivariate Student's t, as in Creal et al (2011).

3 Testing

The model of the previous section provides a framework for testing for time varying correlation. Under the null hypothesis of constant correlation in a Gaussian model with constant variances, the score for γ is

$$u_t = \frac{1}{4}(x_{1t} + x_{2t})^2 \frac{1-r}{1+r} - \frac{1}{4}(x_{1t} - x_{2t})^2 \frac{1+r}{1-r} + r, \tag{8}$$

where r is the sample correlation and the $x_{it}'s$ are standardized observations, that is $x_{it} = y_{it}/s_i$, $i = 1, 2$, where s_i^2 is the sample variance. The portmanteau statistic is

$$Q_u(P) = T \sum_{j=1}^{P} r_u^2(j), \tag{9}$$

where $r_u(j)$ is the $j-th$ sample autocorrelation of u_t. The Ljung-Box statistic

$$Q_u^*(P) = T(T+2) \sum_{j=1}^{P} (T-j)^{-1} r_u^2(j),$$

may also be used; the asymptotic distribution of both statistics under the null hypothesis is χ_P^2. When $r = 0$ the $Q_u(P)$ statistic reduces to the *moment-based* portmanteau test of Bollerslev (1990), because $u_t = x_{1t}x_{2t}$.

For a bivariate t-distribution with degrees of freedom ν, the scores in (8) can be amended by modifying the observations so that they enter as

$$x_{it}^* = x_{it} \sqrt{\frac{\nu+2}{\nu} \frac{1}{w_t}}, \qquad i = 1, 2, \qquad t = 1, ..., T, \tag{10}$$

where

$$w_t = 1 + \frac{1}{\nu(1-r^2)} \left({x_{1t}}^2 - 2r x_{1t} x_{2t} + {x_{2t}}^2 \right)$$

and σ_1, σ_2 and ρ are jointly estimated by maximum likelihood (ML).

When changing volatility is estimated, the residuals are redefined as $x_{it} = y_{it}/\widetilde{\sigma}_{it}$, $i-1, 2$, where the $\widetilde{\sigma}_{it}'s$ is obtained from an EGARCH volatility model. An extra term must be then added to $Q_u(P)$ to give the LM test statistic.

The first sub-section below derives the LM test. The second sub-section uses a local power argument to demonstrate the value of using the scores to capture information on the level of correlation. This is then followed by a discussion of the choice of P and the use of an information criterion to determine a suitable value. The test of Nyblom (1989), which is also based on the scores of (8), is given in the last sub-section.

3.1 Lagrange multiplier tests

The portmanteau test may be derived as an LM test of the null hypothesis that $\kappa_0 = \kappa_1 = ... = \kappa_{P-1} = 0$, against the alternative $\kappa_i \neq 0, i = 0, ..., P-1$, in the dynamic model

$$\gamma_{t|t-1} = \omega + \kappa_0 u_{t-1} + ... + \kappa_{P-1} u_{t-P}, \quad t = 1, ..., T. \tag{11}$$

Let $\boldsymbol{\theta} = (\omega, \sigma_1, \sigma_2)'$ denote fixed parameters other than those in $\boldsymbol{\kappa} = (\kappa_0, .., \kappa_{P-1})'$. The LM test statistic is

$$LM_u(P) = \frac{1}{T} \begin{bmatrix} \partial \ln L/\partial \boldsymbol{\kappa}' & \mathbf{0}' \end{bmatrix} \begin{bmatrix} \mathbf{I}_{\kappa\kappa} & \mathbf{I}_{\kappa\theta} \\ \mathbf{I}_{\theta\kappa} & \mathbf{I}_{\theta\theta} \end{bmatrix}^{-1} \begin{bmatrix} \partial \ln L/\partial \boldsymbol{\kappa} \\ \mathbf{0} \end{bmatrix}, \tag{12}$$

where $\mathbf{I}_{\boldsymbol{\kappa\kappa}}$ denotes the information matrix for $\boldsymbol{\kappa}$ for a single observation and so on. For the t-th observation

$$\frac{\partial \ln f_t}{\partial \boldsymbol{\kappa}} = \frac{\partial \ln f_t}{\partial \gamma_{t|t-1}} \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}} = u_t \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}}$$

and so $\mathbf{I}_{\boldsymbol{\kappa\kappa}}$ is

$$\begin{aligned} E\left[\frac{\partial \ln f_t}{\partial \boldsymbol{\kappa}} \frac{\partial \ln f_t}{\partial \boldsymbol{\kappa}'}\right]_{\kappa=0} &= EE_{t-1}\left[\frac{\partial \ln f_t}{\partial \gamma_{t|t-1}} \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}} \frac{\partial \ln f_t}{\partial \gamma_{t|t-1}} \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right] \\ &= E\left[E_{t-1}\left[\left(\frac{\partial \ln f_t}{\partial \gamma_{t|t-1}}\right)^2\right] \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}} \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right] \\ &= E\left[\left(\frac{\partial \ln f_t}{\partial \gamma}\right)^2\right] E\left[\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}} \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right] \\ &= \sigma_u^2 E\left[\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}} \frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right]. \end{aligned}$$

Under the null hypothesis, the conditional expectation of the squared score is fixed and hence equal to the information quantity in the static model, this is σ_u^2, the variance of the scores.

We have

$$\frac{\partial \gamma_{t|t-1}}{\partial \kappa_j} = \sum_{i=1}^{P} \kappa_{i-1} \frac{\partial u_{t-i}}{\partial \kappa_j} + u_{t-j-1}, \quad j = 0, ..., P-1,$$

but under the null hypothesis $\boldsymbol{\kappa} = \mathbf{0}$, so $\partial\gamma_{t|t-1}/\partial\boldsymbol{\kappa} = \mathbf{u}_{t-1}$,where $\mathbf{u}_{t-1} = (u_{t-1},\ u_{t-2,}..., u_{t-P})'$. Hence

$$E\left(\frac{\partial\gamma_{t|t-1}}{\partial\boldsymbol{\kappa}}\frac{\partial\gamma_{t|t-1}}{\partial\boldsymbol{\kappa}'}\right) = \sigma_u^2\mathbf{I}_P,$$

where $\mathbf{I}_P$ is a $P \times P$ identity matrix, and so $\mathbf{I}_{\boldsymbol{\kappa\kappa}} = \sigma_u^4\mathbf{I}_P$. Furthermore

$$\begin{aligned} E\left[\frac{\partial \ln f_t}{\partial\boldsymbol{\theta}}\frac{\partial \ln f_t}{\partial\boldsymbol{\kappa}'}\right]_{\kappa=0} &= EE_{t-1}\left[\frac{\partial \ln f_t}{\partial\boldsymbol{\theta}}\frac{\partial \ln f_t}{\partial\gamma_{t|t-1}}\frac{\partial\gamma_{t|t-1}}{\partial\boldsymbol{\kappa}'}\right] \\ &= E\left[\frac{\partial \ln f_t}{\partial\boldsymbol{\theta}}\frac{\partial \ln f_t}{\partial\gamma}\right]E\left[\frac{\partial\gamma_{t|t-1}}{\partial\boldsymbol{\kappa}'}\right] = \mathbf{0}. \end{aligned}$$

Note that because ω appears directly in the dynamic equation,

$$\frac{\partial \ln f_t}{\partial\omega} = \frac{\partial \ln f_t}{\partial\gamma_{t|t-1}}\frac{\partial\gamma_{t|t-1}}{\partial\omega} = u_t.1$$

under the null hypothesis. Thus $\mathbf{I}_{\boldsymbol{\theta\kappa}} = \mathbf{0}$ and so

$$LM_u(P) = \frac{1}{T}\frac{\partial \ln L}{\partial\boldsymbol{\kappa}'}\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}\frac{\partial \ln L}{\partial\boldsymbol{\kappa}}. \tag{13}$$

On substituting for $\mathbf{I}_{\kappa\kappa}$ and noting that

$$\frac{\partial \ln L}{\partial\kappa_j} = \sum\frac{\partial \ln f_t}{\partial\gamma_{t|t-1}}\frac{\partial\gamma_{t|t-1}}{\partial\kappa_j} = \sum u_t u_{t-1-j}, \quad j = 0, 1, ..., P-1,$$

the $Q_u(P)$ statistic, (9), is obtained.

Remark 2 *Although the form of the link function is important for estimation, it does not affect the LM statistic in (9).*

The above derivation is as Harvey (2013, sub-section 2.5.1), but stated more generally, and it applies to any time-varying parameter in a DCS model when the other parameters are fixed[1]. Now suppose some of the other parameters, denoted $\boldsymbol{\lambda}$, are time-varying, with dynamics depending on a set of parameters $\boldsymbol{\psi}$, but not depending on $\gamma_{t|t-1}$. In the present context this means that each volatility comes from a univariate model; see the Remark at the end

[1] Calvori et al (2014) also propose tests based on conditional scores but develop the methods in a different direction.

of Section 2. Suppose, for simplicity, that the only other constant parameter is ω. Then $\boldsymbol{\theta} = (\boldsymbol{\psi}', \omega)'$. Assuming identifiability under the null hypothesis, the formula for a partitioned inverse means that the LM statistic, (12), can be written

$$LM_u(P) = \frac{1}{T}\frac{\partial \ln L}{\partial \boldsymbol{\kappa}'}\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}\frac{\partial \ln L}{\partial \boldsymbol{\kappa}} + \frac{1}{T}\frac{\partial \ln L}{\partial \boldsymbol{\kappa}'}\left[\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}\mathbf{I}_{\boldsymbol{\kappa\theta}}\left(\mathbf{I}_{\boldsymbol{\theta\theta}} - \mathbf{I}_{\boldsymbol{\kappa\theta}}\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}\mathbf{I}_{\boldsymbol{\kappa\theta}}'\right)^{-1}\mathbf{I}_{\boldsymbol{\kappa\theta}}'\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}\right]\frac{\partial \ln L}{\partial \boldsymbol{\kappa}},$$

where the second term on the right hand side is positive semi-definite[2] resulting in a modified LM statistic that cannot be less than the LM statistic with fixed $\boldsymbol{\lambda}$, which is the portmanteau statistic of (13). Hence the $Q_u(P)$ test is more conservative than the LM test because $Q_u(P) \leq LM_u(P)$.

The second term in the LM statistic acts as a correction for the estimation of $\boldsymbol{\psi}$ and it can be shown to be equivalent to the result by Pierce (1982), which has been used in the GARCH literature to correct specification tests based on estimated residuals; see, for example, Bera and Zuo (1996) and Tse (2002). We have $\mathbf{I}_{\boldsymbol{\kappa\theta}} = [\mathbf{I}_{\kappa\psi}', \mathbf{I}_{\kappa\omega}']' = [\mathbf{I}_{\kappa\psi}', \mathbf{0}']'$ because $\mathbf{I}_{\omega\boldsymbol{\kappa}} = \mathbf{0}$; see above (13). Following on from Pierce (1982),

$$\begin{aligned}
E\left[\frac{\partial \ln f_t}{\partial \boldsymbol{\psi}}\frac{\partial \ln f_t}{\partial \boldsymbol{\kappa}'}\right]_{\kappa=0} &= E\left[\frac{\partial^2 \ln f_t}{\partial \boldsymbol{\psi}\partial \boldsymbol{\kappa}'}\right]_{\kappa=0} = E\left[\frac{\partial}{\partial \boldsymbol{\psi}}\left(\frac{\partial \ln f_t}{\partial \gamma}\frac{\partial \gamma}{\partial \boldsymbol{\kappa}'}\right)\right] = E\left[\frac{\partial (u_t \mathbf{u}_{t-1}')}{\partial \boldsymbol{\psi}}\right] \\
&= E\left[\frac{\partial u_t}{\partial \boldsymbol{\psi}}\mathbf{u}_{t-1} + u_t\frac{\partial \mathbf{u}_{t-1}'}{\partial \boldsymbol{\psi}}\right] = E\left[\frac{\partial u_t}{\partial \boldsymbol{\psi}}\mathbf{u}_{t-1}\right] + E\left[E_{t-1}\left(u_t\frac{\partial \mathbf{u}_{t-1}'}{\partial \boldsymbol{\psi}}\right)\right] \\
&= E\left[\frac{\partial u_t}{\partial \boldsymbol{\psi}}\mathbf{u}_{t-1}\right] = E\left[\frac{\partial^2 \ln f_t}{\partial \boldsymbol{\psi}\partial \gamma}\mathbf{u}_{t-1}\right].
\end{aligned}$$

Once the model has been estimated under the null hypothesis, the above expression can be approximated numerically.

Suppose that $\mathbf{I}_{\boldsymbol{\lambda}\gamma}$ does not depend on $\boldsymbol{\lambda}$. This is the situation here when EGARCH models are used; see (6). Consider one of the elements, λ_i, in $\boldsymbol{\lambda}$.

[2] This follows from that fact that under identifiability, the full information matrix in (12) will be positive definite. It then follows that the sub matrix $\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}$ and its Schur complement $\left(\mathbf{I}_{\boldsymbol{\theta\theta}} - \mathbf{I}_{\boldsymbol{\kappa\theta}}\mathbf{I}_{\boldsymbol{\kappa\kappa}}^{-1}\mathbf{I}_{\boldsymbol{\kappa\theta}}'\right)^{-1}$ will also be positive definite; see Abadir and Magnus(2005, p 228).

Dropping the subscript on λ_i, we have

$$\begin{aligned}
\mathbf{I}_{\psi\kappa} &= E\left[\frac{\partial \ln f_t}{\partial \boldsymbol{\psi}}\frac{\partial \ln f_t}{\partial \boldsymbol{\kappa}'}\right]_{\kappa=0} = EE_{t-1}\left[\frac{\partial \ln f_t}{\partial \lambda_{t|t-1}}\frac{\partial \lambda_{t|t-1}}{\partial \boldsymbol{\psi}}\frac{\partial \ln f_t}{\partial \gamma_{t|t-1}}\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right] \\
&= E\left[E_{t-1}\left[\frac{\partial \ln f_t}{\partial \lambda_{t|t-1}}\frac{\partial \ln f_t}{\partial \gamma_{t|t-1}}\right]\frac{\partial \lambda_{t|t-1}}{\partial \boldsymbol{\psi}}\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right] \\
&= E\left[\frac{\partial \ln f_t}{\partial \lambda}\frac{\partial \ln f_t}{\partial \gamma}\right]E\left[\frac{\partial \lambda_{t|t-1}}{\partial \boldsymbol{\psi}}\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right] = -\rho E\left[\frac{\partial \lambda_{t|t-1}}{\partial \boldsymbol{\psi}}\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right].
\end{aligned}$$

The elements in $E\left[\frac{\partial \lambda_{t|t-1}}{\partial \boldsymbol{\psi}}\frac{\partial \gamma_{t|t-1}}{\partial \boldsymbol{\kappa}'}\right]$ will also depend on ρ because of the correlation between the (contemporaneous) scores. Thus $\mathbf{I}_{\psi\kappa} \neq \mathbf{0}$, unless $\rho = 0$.

When $\rho = 0$ the LM statistic reverts to the original portmanteau statistic, $Q_u(P)$, and this, in turn, is the same as the moment-based portmanteau statistic, $Q_x(P)$.

3.2 Local Power for P=1

Consider the Gaussian DCS model $\gamma_{t+1|t} = \omega + \kappa u_t$. We are interested in the power of the proposed score test for the null hypothesis $H_0 : \kappa = \kappa_0 = 0$, against local alternatives of the form $\kappa = \delta/\sqrt{T}$. The asymptotic distribution of the test statistic, $Q_u(1)$, is then $\chi_1^2(I(\kappa_0)\delta^2)$, a non-central χ^2 with noncentrality parameter $I(\kappa_0)\delta^2$; see Godfrey (1988, p 18). Because $I(\kappa_0)$ is the element of the information matrix for $\kappa_0 = 0$, we have $I(\kappa_0) = (1+\rho^2)^2$. (Estimation of the variances, σ_1^2 and σ_2^2, makes no difference; see sub-section 3.1). Thus for a given value of δ, the local power increases as $|\rho| \to 1$. This property will be apparent in the Monte Carlo results. By contrast, the power of the moment-based test does not increase with $|\rho|$.

3.3 Choice of P

Although the portmanteau test is derived against a moving average alternative, a stationary first-order model of the form, (3), is a more likely candidate for a dynamic model. In this case, it can be shown that the LM test is the portmanteau test with $P = 1$; see, for instance, Lee (1991). However, when the process driving $\gamma_{t+1|t}$ is very persistent, that is ϕ is close to one, the power may be increased by setting P to a relatively high value, perhaps selected by a criterion such as $P = \sqrt{T}$. An alternative way forward is to select P using a consistent information criterion, as in Escanciano and Lobato (2009); see appendix. Under the alternative, such a model selection procedure should select an increasing number of lags as ϕ goes to unity. Under the null hypothesis, only the first lag is selected in large samples with probability one. As a result, the asymptotic distribution under the null hypothesis is χ_1^2. Simulation results (not reported here) indicated that this last approach was the best option and so it was adopted for all tests based on portmanteau statistics. Such test statistics will be denoted simply as Q_u rather than $Q_u(P)$. The LM statistics are similarly denoted as $LM_u(P)$ and LM_u and the moment-based test statistics as $Q_x(P)$ and Q_x.

3.4 Nyblom test

Nyblom (1989) gives a general test for parameter constancy against a random walk alternative based on the LM principle. In the present context, the statistic ends up being based on the same scores as in the portmanteau test. It can be written

$$N = \frac{1}{T^2\sigma_u^2}\sum_{j=1}^{T}\left(\sum_{k=j}^{T}u_k\right)^2.$$

Under the null hypothesis of parameter constancy, the statistic follows a Cramer-von Mises distribution with a 5% critical value of 0.462. The same critical value can be used when the scores are constructed from dynamic volatility estimates. Although the Nyblom test is usually regarded as a test against a random walk alternative, it can also be interpreted a test against a very persistent, but stationary, alternative, as in Harvey and Streibel (1998).

4 Monte Carlo experiments

To evaluate the performance of the proposed testing procedure, a simulation study was conducted on a number of models. The results are confined to versions of the tests in which the number of lags is determined by an information criterion, as in sub-section 3.3. The Ljung-Box form of the portmanteau statistic was used and volatilities were estimated from univariate GARCH or EGARCH models[3].

[3]EGARCH models were always used for the DCS test, whereas GARCH models were used for the other tests when the true model was not the DCS; the exception is the Tse

Several tests from the existing literature were also considered. These are as follows.

i) The moment-based portmanteau test, as in Bollerslev (1990), based on autocorrelations constructed from the cross-product of standardized (volatility corrected) residuals. As with the score-based tests, the value of P is selected by an information criterion and so the test statistic is denoted as Q_x. (Since we are using the Box-Ljung form throughout this should actually be Q_x^* to be consistent with the original notation. However, it is neater to drop the star). The results for a version of the test that corrects for volatility estimation are omitted as they are very close to those of the Q_x test.

ii) A residual regression test, RR, proposed by Tse (2002), in which $x_1x_2-\rho$ is regressed on P lags. He also provides a correction based on Pierce (1982) to allow for the estimation of volatility. The third test considered is the LM test of Tse (2000) based on an alternative model $\rho_t = c + by_{it-1}y_{jt-1}$, with the score vector calculated using a set of recursive equations. Estimation of the volatility models was based on MLEs for the bivariate time series and all corrected statistics used numerical derivatives. The results for the residual regression test are based on a lag length of two in accordance with Tse (2002)[4].

iii) The test of Bera and Kim (2002), denoted BK, gets around the need to assume a functional form for the time-varying correlations by focussing on behaviour local to the constant parameter case. They use Taylor approximations based on the variance of the errors driving the time varying parameters being small. The test statistic is again constructed from standardized residuals, x_{it}, $i = 1, 2$, and is given by

$$BK = \frac{\left[\sum_{t=1}^{T}(\xi_{1t}^2\xi_{2t}^2 - 1 - 2\hat{\rho}^2)\right]^2}{4T(1+4\hat{\rho}^2+\hat{\rho}^4)},$$

where $\xi_{1t} = (x_{1t} - \hat{\rho}x_{2t})/(\sqrt{1-\hat{\rho}^2})$ and $\xi_{2t} = (x_{2t} - \hat{\rho}x_{1t})/(\sqrt{1-\hat{\rho}^2})$.

The simulation study consists of three models with a bivariate normal conditional distribution, and one with a t-distribution. The sample sizes

test where GARCH was used in all cases.

[4]All test statistics requiring a choice of lag length were also considered with fixed lag lenths of 2, 10 and 20 in a series of preliminary simulations. The relative performance of the various tests was similar for all lag lengths. Hence, only the preferred lag length is presented.

were $T = 500$ and 1000 with $5,000$ replications used in power comparisons and $10,000$ in size comparisons.

4.1 DCS model

The DCS model has dynamic equations for the correlation and volatility as in (3) and (5). The three parameters in the equation for correlation were varied across the sets $\omega = [0, 0.9]$, $\phi = [0.6, 0.99]$ and $\kappa = [0.01, 0.1]$, whereas the parameters governing the EGARCH volatility dynamics were fixed at $\omega_i = 0$, $\phi_i = 0.95$ and $\kappa_i = 0.2$, $i = 1, 2$.[5] Only one parameter was changed at a time, with the base set of parameters given by $\omega = 0.4$, $\phi = 0.9$ and $\kappa = 0.05$. Note that $\omega = 0.4$ and 0.8 correspond to $\rho = 0.38$ and 0.66 respectively.

4.1.1 Size of tests

From the results in Table 1, the LM_u test appears to be slightly oversized in finite samples as does the Q_u test, though to lesser extent (because Q_u cannot be greater than LM_u). This size distortion, which is due to the use of the information criterion to choose P, declines as the sample size increases and becomes negligible for $T = 1000$. The estimated rejection probabilities of the N test increase as the correlation increases, whereas those of the moment-based portmanteau test decrease.

Table 1: Estimated size ($\times 100$) of tests for a DCS model with EGARCH volatility.

		Test						
	$\omega(\rho)$	LM_u	Q_u	N	Q_x	BK	Tse	cRR
$T = 500$	0	7.08	6.40	4.60	6.50	5.52	6.03	5.18
	0.4 (0.38)	7.04	6.35	5.48	6.27	5.68	6.21	5.18
	0.8 (0.66)	7.35	6.64	7.33	5.95	6.02	6.73	5.86
$T = 1000$	0	5.94	5.67	4.58	5.9	5.22	5.45	4.70
	0.4	6.16	5.90	5.24	5.86	5.26	5.54	5.21
	0.8	5.70	5.29	5.95	5.27	6.05	6.66	5.38

Note: LM_u is score-based LM test, Q_u is score-based portmanteau test, N is Nyblom test, Q_x is moment-based (Bollerslev) test, BK is Bera and Kim test, Tse is Tse test, cRR is (corrected) residual regression test.

[5] The values for κ are relatively large but lower values give similar results.

4.1.2 Power comparisons

Table 2 shows powers, or, more precisely, estimated probabilities of rejection. The salient feature is the increasing extent to which the score-based tests dominate the moment-based tests as ω increases. A clearer impression of the relative performance of the tests comes from Figure 1 which shows the estimated powers for the Q_u, LM_u, N and Q_x tests for $T = 500$ as the parameter ω (governing the unconditional level of correlation) increases from zero to 0.8. We find that the new score-based tests and the Nyblom test outperform the competition across virtually the entire range of ω. The power of the score-based tests increases as the unconditional level of the correlation rises, as indicated by the local power results of sub-section 3.2, whereas the power of the moment-based test does not; in fact it shows a slight fall. When $T = 500$, the Q_u and LM_u tests outperform the Nyblom test for ω above 0.5, but when $T = 1000$ the break-even value falls to 0.3, as shown in Figure 2. The rejection probabilities with the conservative Q_u test are only slightly smaller than those for the LM_u test when $T = 1000$.

Table 2: Powers of Tests for DCS Model with Different Levels of Correlation

	Test									
	ω	ϕ	κ	LM_u	Q_u	N	Q_x	BK	Tse	cRR
$T = 500$	0	0.9	0.05	12.4	11.4	18.8	11.6	6.6	7.6	10.5
	0.4	0.9	0.05	21.0	18.5	21.6	10.0	7.0	9.2	9.8
	0.8	0.9	0.05	37.8	34.5	30.4	7.5	11.6	11.7	8.4
$T = 1000$	0	0.9	0.05	16.5	15.4	19.7	15.3	6.9	7.7	18.3
	0.4	0.9	0.05	28.5	26.9	22.7	10.7	9.1	8.8	13.6
	0.8	0.9	0.05	55.3	53.2	31.4	6.6	17.3	13.0	8.6

Figures 3 and 4 show the power of the tests as the parameter κ is varied across the set (0.01, 0.1) with sample sizes T = 500 and T = 1000 respectively. Once again the score-based tests, including the Nyblom test, outperform the others across almost the entire range examined. It seems that the Nyblom test has greater power against smaller deviations from the null. However as before, the range of values over which the score-based portmanteau test matches or improves upon the Nyblom test increases as the sample size reaches T = 1000. Once again the difference between the Q_u and LM_u tests is small throughout.

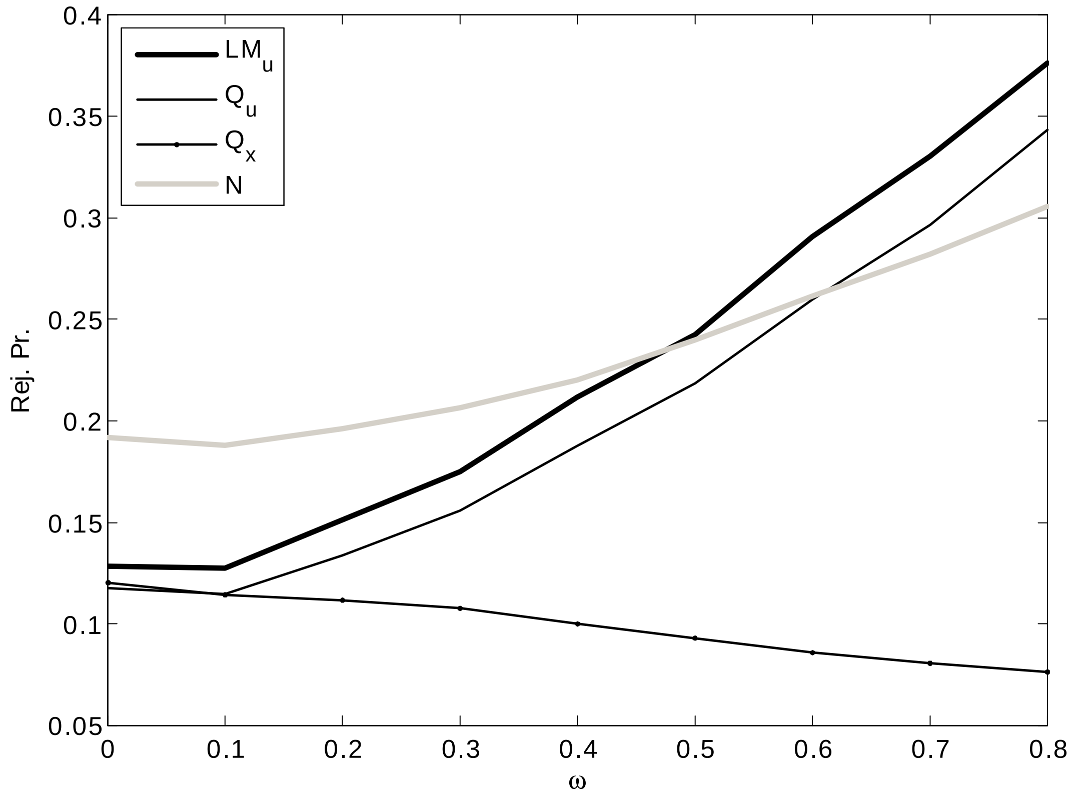

Figure 1: Power Comparison across ω with $\phi = 0.9$, $\kappa = 0.05$, $T = 500$. LM_u is DCS-LM test, Q_u is score-based portmanteau test, Q_x is moment-based portmanteau test and N is Nyblom test.

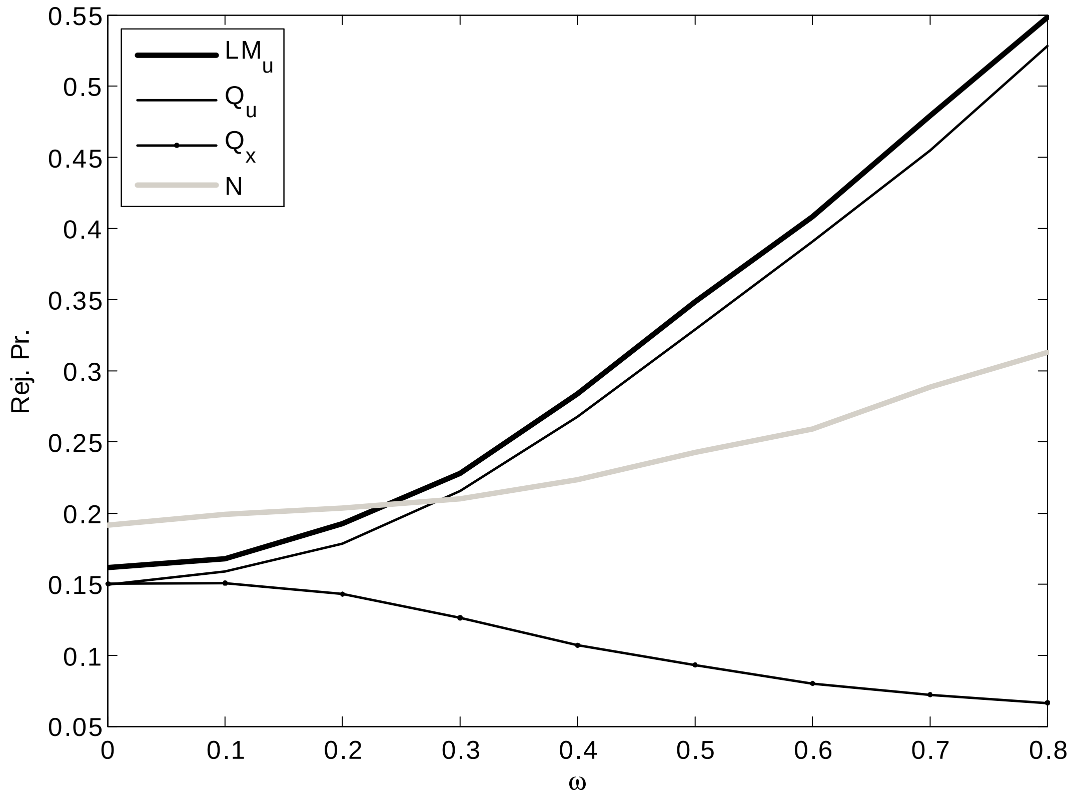

Figure 2: Power Comparison across ω with $\phi = 0.9$, $\kappa = 0.05$, $T = 1000$.

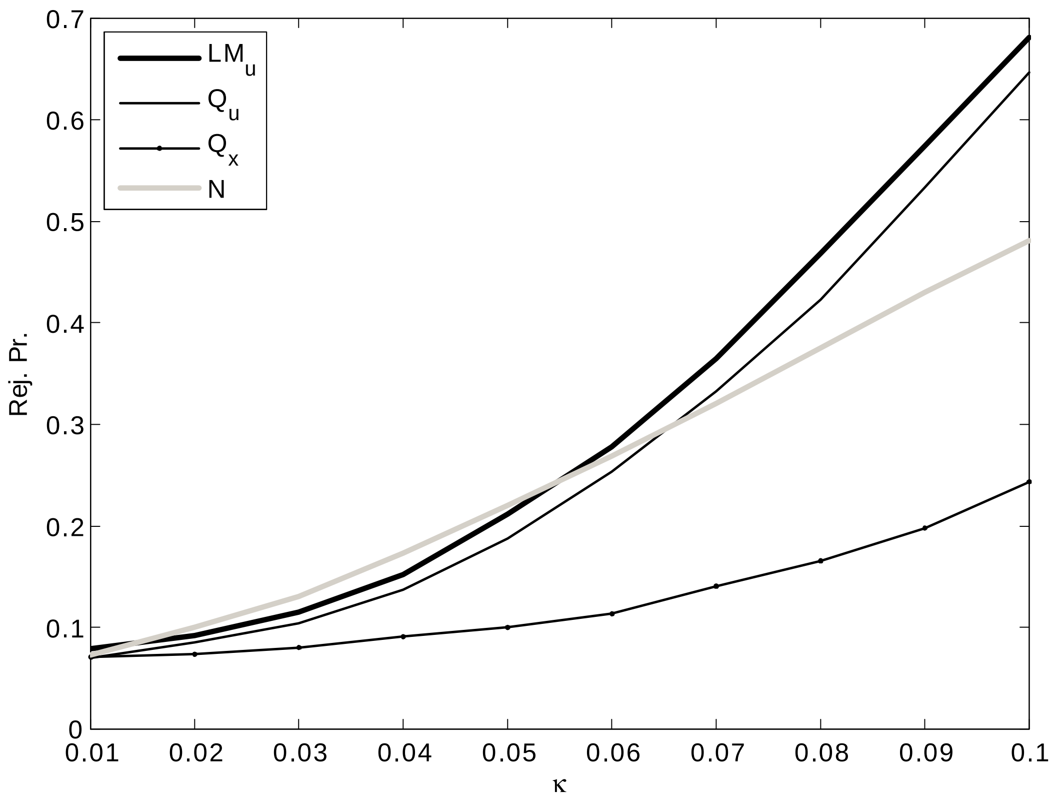

Figure 3: Power Comparison across κ with $\phi = 0.9$, $\omega = 0.05$, $T = 500$.

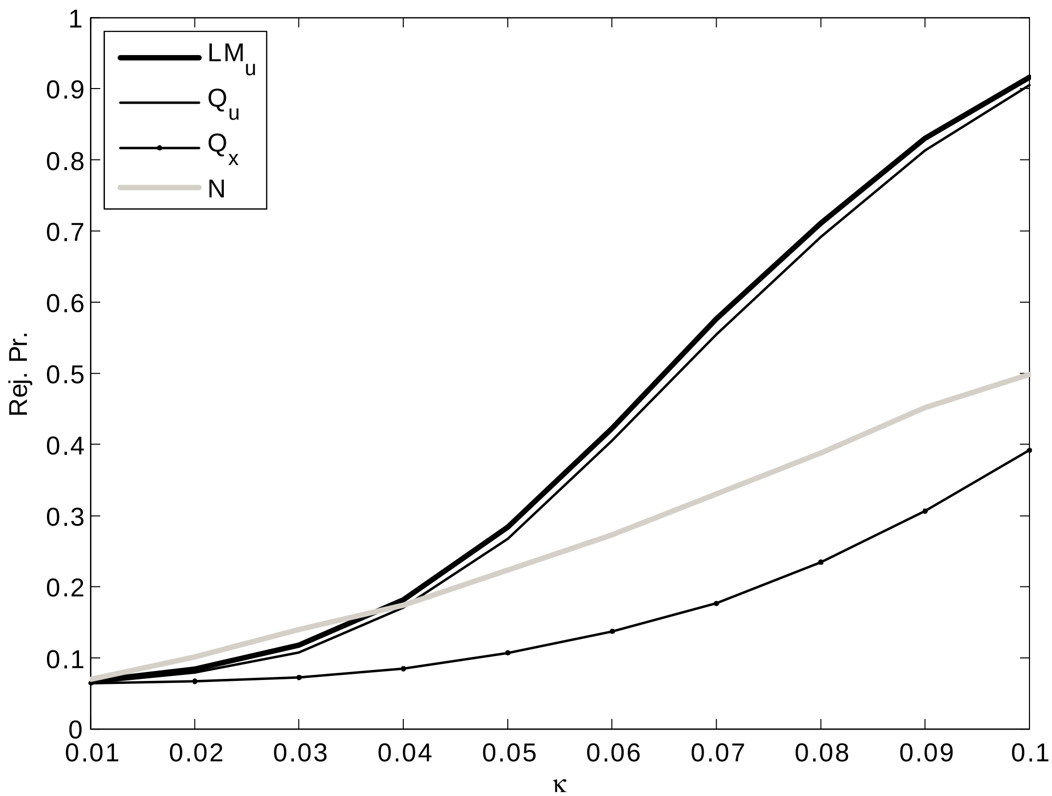

Figure 4: Power Comparison across κ with $\phi = 0.9$, $\omega = 0.05$, $T = 1000$.

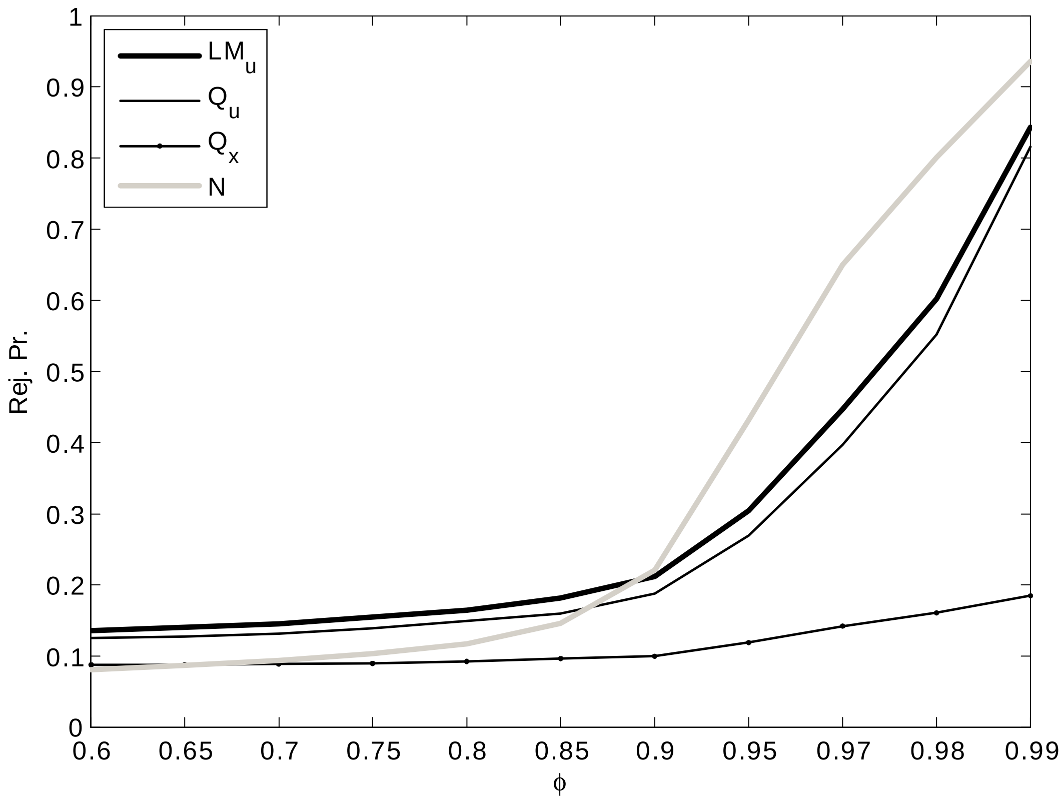

Figure 5: Power Comparison across κ with $\kappa = 0.05$, $\omega = 0.05$, $T = 500$.

Figures 5 and 6 show the power of the tests as the parameter ϕ is varied across the set (0.6, 0.99) with sample sizes T = 500 and T = 1000 respectively. Once again, the score-based tests, including the Nyblom test, perform best overall with the gap increasing with ϕ. When $T = 500$, the Q_u and LM_u tests are beaten by the Nyblom test for $\phi > 0.9$, but the break-even value of ϕ rises to around 0.95 when $T = 1000$; compare similar findings in Harvey and Streibel (1998).

Rejection probabilities for the BK, Tse and cRR tests are little better, and sometimes worse, than those for the Q_x test. Results are available on request.

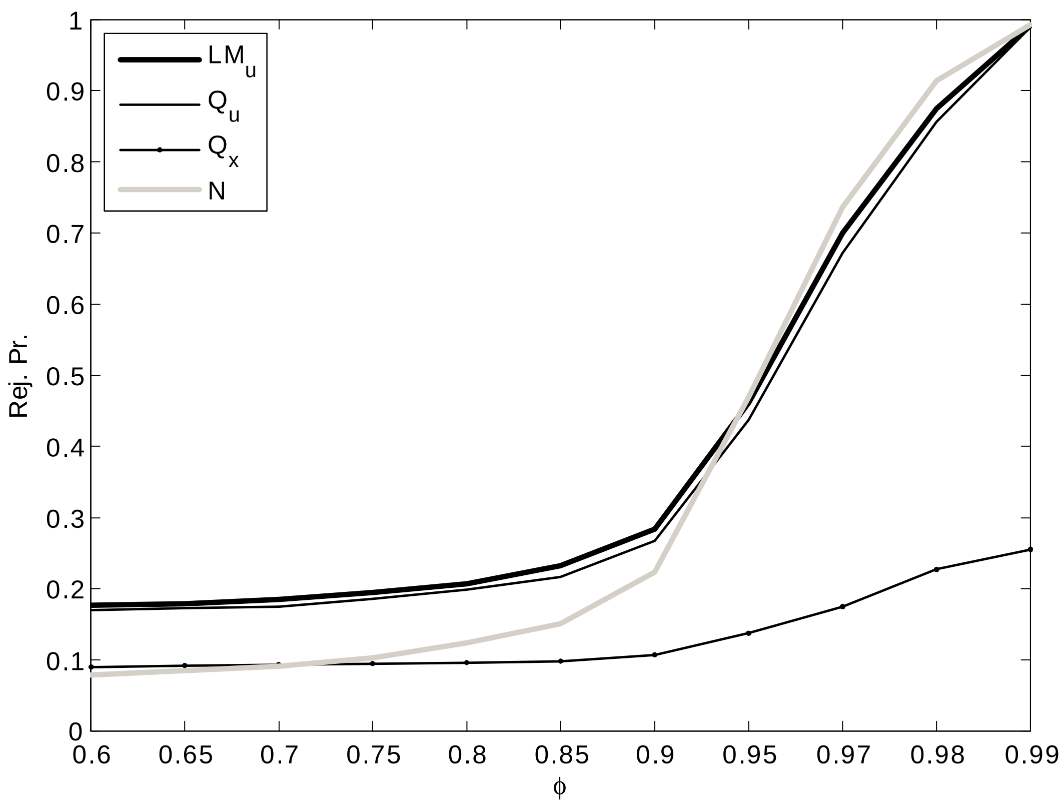

Figure 6: Power Comparison across κ with $\kappa = 0.05$, $\omega = 0.05$, $T = 1000$.

4.2 Stochastic Correlation

In the second model the correlation to be driven by an unobserved components Gaussian autoregressive process,

$$\gamma_t = \omega(1-\phi) + \phi\gamma_{t-1} + \kappa\eta_t, \quad \eta_t \sim NID(0,1), \tag{14}$$

in which the correlations were again constrained to lie in the range $(-1, 1)$ by using a transformation of the form (1). We set values for $\omega = \{0, 0.4\}$, $\phi = \{0.8, 0.95\}$ and $\kappa = \{0.1, 0.15, 0.2\}$, but the time-varying volatility is as in the DCS model.

Table 3 shows the estimated rejection probabilities for various values of the parameters ω, ϕ and κ at sample sizes of 500 and 1000. The findings from the previous sub-section generally carry over to this setting. Contrasting the first three rows ($\omega = 0$) with the last three rows ($\omega = 0.4$) of both Panel A and Panel B shows that the powers of the score-based tests increase with the level of correlation, ω, whereas that of the Q_x test deteriorates, as do the powers of the Tse and cRR tests. The one exception is the BK test which for this particular model, but not for the others, does rather well. The Q_u and LM_u tests dominate the Nyblom test, even for $T = 500$; this was not the case for the DCS model as reported in table 2. Finally, the relative performance of the score-based tests improves as ϕ increases from 0.8 to 0.95, which is consistent with Figures 5 and 6.

Table 3: Power Comparison for Stochastic Correlation Model

Panel A:	T = *500*								
ω	ϕ	κ	LM_u	Q_u	N	Q_x	BK	Tse	cRR
0	0.8	0.2	33.2	31.7	20.2	32.0	47.4	24.7	29.2
0	0.95	0.1	42.5	40.7	48.0	39.6	40.3	27.5	35.9
0	0.95	0.15	80.5	79.4	70.6	77.3	79.2	60.2	75.1
0.4	0.8	0.2	39.9	37.8	22.9	20.7	59.2	24.2	20.0
0.4	0.95	0.1	50.7	48.1	52.5	26.4	50.9	27.2	24.9
0.4	0.95	0.15	84.9	83.2	74.3	61.5	87.0	56.0	59.7
Panel B:	T = *1000*								
ω	ϕ	κ	LM_u	Q_u	N	Q_x	BK	Tse	cRR
0	0.8	0.2	54.7	53.8	19.6	54.4	72.8	38.6	55.0
0	0.95	0.1	68.2	67.0	51.3	66.3	65.4	43.8	65.9
0	0.95	0.15	97.8	97.6	75.6	97.4	97.4	82.4	97.3
0.4	0.8	0.2	62.2	61.3	23.3	34.5	85.3	34.3	35.2
0.4	0.95	0.1	76.7	75.1	56.1	44.3	78.2	40.0	44.6
0.4	0.95	0.15	98.7	98.5	79.5	88.3	98.9	76.5	88.2

4.3 Diagonal Vech-GARCH Model

Because the new tests are derived within the framework of a DCS model for changing correlation, it could be argued that the results of sub-section 4.1 and, to a lesser extent, those of sub-section 4.2, are weighted in favor of them. We therefore consider a third model in which the dynamics are moment-based. In this diagonal vech GARCH model the covariances are generated by a dynamic equation which is similar in form to that of the equations for variance. Thus

$$\begin{aligned} \sigma_{12t} &= \delta + \beta\sigma_{12t-1} + \alpha y_{1t-1}y_{2t-1}, \qquad t = 2, ..., T, \\ \sigma_{it}^2 &= \delta_i + \beta_i\sigma_{it-1}^2 + \alpha_i y_{it-1}^2, \quad i = 1, 2. \end{aligned}$$

Table 4 shows the rejection probabilities with $\beta = 0.8$ and $\alpha = 0.05$, that is a persistence of 0.85, together with $\delta_i = 0.05$, $\beta_i = 0.8$ and $\alpha_i = 0.15$, $i = 1, 2$. Generally speaking, the findings from the DCS simulations carry over to this setting. In particular, the powers of the score-based tests increase with an increase in the unconditional level of correlation, driven by ν, whereas the power of the moment-based Q_x test deteriorates, as do the powers of the other tests based on the product of standardised residuals. Other results, not reported here, confirm that, as expected, power increases as α gets bigger.

Table 4: Power Comparison for Vech GARCH Model

	Parameter			Test						
	δ	β	α	LM_u	Q_u	N	Q_x	BK	Tse	cRR
$T = 500$	0.02	0.8	0.05	21.5	20.4	14.3	15.5	5.5	15.3	14.0
	0.05	0.8	0.05	35.9	34.1	27.0	8.1	10.3	9.0	7.6
$T = 1000$	0.02	0.8	0.05	34.3	33.5	15.9	26.1	6.2	23.0	25.8
	0.05	0.8	0.05	55.9	54.8	29.1	10.0	13.2	9.7	10.0

4.4 Bivariate t-distribution

The above simulations are for Gaussian models and, as such, demonstrate the advantages of using the scores for a changing correlation test. However, the conditional distributions of financial asset returns are often heavy-tailed and a t-distribution is usually a better option. For modeling volatility, the DCS approach leads to an EGARCH model in which the dynamics of the logarithm of scale, λ, are driven by

$$u_t = \frac{(\nu+1)(y_t-\mu)^2}{\nu \exp(2\lambda_{t|t-1}) + (y_t-\mu)^2} - 1, \quad \nu > 0. \tag{16}$$

Because u_t is a linear function of a beta distribution at the true parameter values, the model is known as Beta-t-EGARCH; see Harvey (2013, ch 4). The fact that the score function is bounded has the practical effect of moderating the influence of outliers.

Rather than carrying out a full LM test for a bivariate t-distribution, the scores in (8) are amended by modifying the observations as in (10), where the standardized observations are obtained by fitting univariate Beta-t-EGARCH models. The scores with respect to correlation are then constructed by estimating the correlation and degrees of freedom in a bivariate t model. Table 5 compares the performance of the resulting portmanteau test, denoted $Q_u(t)$, with that of the Gaussian test portmanteau test studied in the previous subsections. The simulations estimate size with 10,000 replications and power with 5,000. Volatility was generated from Beta-t-EGARCH models with $\nu = 8$ and parameters $\omega_i = 0$, $\phi_i = 0.95$ and $\kappa_i = 0.1$, for $i = 1, 2$. The first two rows of the table show the size of the tests for two levels of correlation. Both tests are slightly oversized, though reasonably close to the nominal 5% level, with the discrepancy decreasing when T rises to 1000. The difference between the Gaussian and t-based tests is much more evident when considering power: the rejection probabilities for $Q_u(t)$ are much higher.

Table 5: Size and Power for a Student t-distribution

			$T = 500$		$T = 1000$	
ω	ϕ	κ	Q_u	$Q_u(t)$	Q_u	$Q_u(t)$
0	1	0	6.2	6.3	5.1	5.9
0.4	1	0	7.1	6.2	6.4	5.9
0	0.8	0.05	19.6	31.4	32.6	53.8
0	0.8	0.1	66.2	89.6	92.1	99.7
0	0.95	0.05	37.7	57.4	64.4	87.9
0	0.95	0.1	88.9	98.6	98.9	100
0.4	0.8	0.05	25.2	33.3	38.1	56.5
0.4	0.8	0.1	67.3	89.4	89.6	99.7
0.4	0.95	0.05	46.9	61.7	70.6	89.9
0.4	0.95	0.1	89.7	98.9	98.7	100

5 Application: Hong Kong and South Korea Stock Indices

To demonstrate the effectiveness of the proposed test statistics, we examine the stability of the correlation between daily local currency returns of the Hong Kong (Hang Seng) and South Korean (SET) stock indices from 2/1/1984 to 27/11/2007. Because of the length of the series ($T = 6237$) and the occurrence of several major events in this time frame[6], we also consider a shorter window between 1/1/2004 and 27/11/2007 ($T = 1019$): this provides a tougher challenge for detecting changing correlation[7].

Table 6 presents the results for score and moment-based tests constructed using volatility-corrected residuals. For the full sample there is strong evidence for time varying correlation. The prob-values for all score-based tests are essentially zero. The moment-based test is slightly less conclusive in that it fails to reject at the 1% level of significance. The higher values of the score-based tests are consistent with the local power and Monte Carlo results because the unconditional correlation over the full sample is 0.24. The

[6] This data was modeled in Harvey (2010) by means of a time varying copula. As noted there, the sample includes i) Black Monday, October 19th, 1987; ii) the speculative attack on the Hong Kong dollar on 20th October 20th, 1997; and iii) the High Technology Crash of October 2nd, 2000.

[7] Data including the recent financial crisis are not considered due to the likelihood that contagion would lead to sharp changes in the correlation structure that would be easy for all tests to identify.

LM_u statistic is much larger than Q_u, although the latter is still very big in absolute terms. More significant is the much larger value of the $Q_u(t)$ statistic, which presumably stems from the fact that the estimated degrees of freedom is 4.7.

Table 6: Tests against Changing Correlation for Hong Kong and South Korean Stock Markets

Sample	LM_u	Q_u	$Q_u(t)$	Q_x	N
2/1/84 - 27/11/07	341.53	285.13	552.81	6.39	34.57
	(0)	(0)	(0)	(1.21)	(<0.1)
1/1/04 - 27/11/07	4.25	4.13	4.64	0.56	1.14
	(3.93)	(4.20)	(3.12)	(45.4)	(< 1.0)

Note: P-values are in parentheses. For the N test these are based values tabulated in Nyblom (1989).

The results for the shorter sub-sample show an even more striking difference between the score and moment-based tests. Whereas the moment-based test fails to reject the null hypothesis of constant correlation at any reasonable significance level, suggesting a period of stability during 2004-2007, the score-based tests demonstrate their higher power by rejecting at the 5% level of significance. This discrepancy is once again explained by the unconditional correlation, which is now 0.61. As before the biggest score-based test statistic is $Q_u(t)$; the degrees of freedom is now 5.99.

Figure 7 plots the time-varying correlation over the full sample when estimated with a bivariate t DCS model. Considerable short run variation is evident throughout, but there is a clear increase in the level, starting in the late 1990s. In the sub-sample after 2004, there is considerable movement in the correlation, which ranges from 0.50 to 0.85. Nevertheless, only the score-based tests are able to detect these changes.

6 Conclusion

The proposed test for time-varying correlation is relatively simple. First standardize the two series by dividing by the scale given by fitting univariate volatility models, preferably Beta-t-EGARCH, to each series. Then construct the scores with respect to correlation by estimating the correlation and degrees of freedom in a bivariate t model. The simple portmanteau statistic,

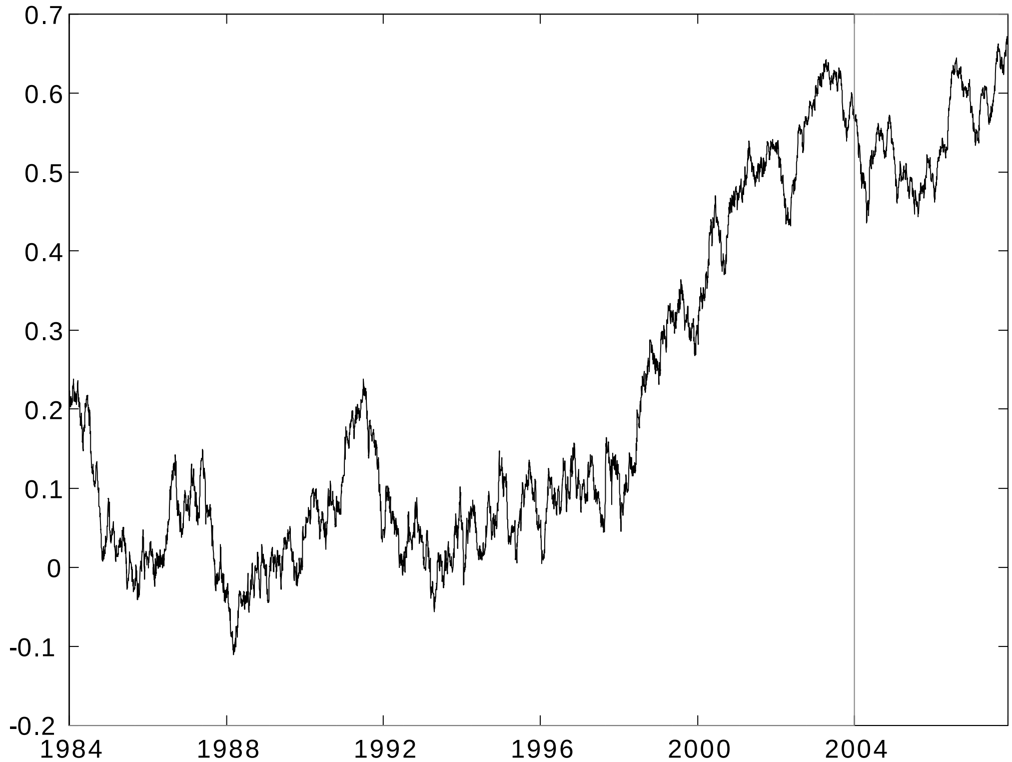

Figure 7: Time Varying Correlation for Hong Kong and Korean Stock Market Indices - 2/1/1984 to 27/11/2007.

in the Ljung-Box form, is constructed with the number of lags chosen by an information criterion.

The simulation results show that there is little to be gained by making the correction demanded by the full LM test. Indeed, the LM test is more oversized than the portmanteau test when the number of lags is selected by an information criterion. The Nyblom test is a good option when the changes in correlation are thought to be very persistent. What is very clear from the simulations is that tests based only on cross-products of residuals are almost always dominated by the score-based tests, with the difference in power increasing as the underlying correlation moves away from zero and often being very considerable.

Further development of tests developed from DCS models, for example tests against time variation in copulas, seems to be a fruitful avenue for future research.

Acknowledgements

Some of the results in this paper were presented at the Time Series Workshop at the University of La Laguna, Tenerife, in January 2014 and at the Econometric Society Australasian meeting in Hobart in July, 2014. We would like to thank participants, particularly Yoosoon Chang, for helpful comments. Thanks also to Jukka Nyblom and our colleagues at Cambridge. We are grateful to the Keynes Fund for financial support.

APPENDIX: Data-driven Q-test

The lag length, P, is selected by the criterion proposed by Escanciano and Lobato (2009), namely

$$P = \min\{p : 1 \leq P \leq d : L_P \geq L_h, \quad h = 1, 2, ...d\},$$

where

$$Q = Q(P) - \pi(p, T, q),$$

d is a fixed upper bound, and $\pi(p, T, q)$ is a penalty term that takes the form

$$\pi(p, T, q) = \begin{cases} p \log T & \text{if} \quad \max\limits_{1 \leq j \leq d} \sqrt{T}|\widetilde{\rho}_j| \quad \leq \quad \sqrt{q \log T} \\ 2p, & \text{if} \quad \max\limits_{1 \leq j \leq d} \sqrt{T}|\widetilde{\rho}_j| \quad > \quad \sqrt{q \log T} \end{cases},$$

where q is some fixed positive number. Escanciano and Lobato (2009) suggest setting $q = 2.4$. Their simulation evidence suggests that the choice of d is not crucial. Here we set $d = 20$.

REFERENCES

Abadir, K. M. and J. R. Magnus, 2005, Matrix Algebra, Cambridge, Cambridge University Press.

Bera, A. K. and S-J. Kim, 2002, Testing constancy of correlation and other specifications of the BGARCH model with an application to international equity returns, Journal of Empirical Finance 9, 171-195.

Bera, A. K. and X-L. Zuo, 1996, Specification test for a linear regression model with ARCH process, Journal of Statistical Planning and Inference 50, 283–308.

Bollerslev, T., 1990, Modeling the coherence in short run nominal exchange rates: A multivariate generalized ARCH model, The Review of Economics and Statistics 72, 498-505.

Calvori, F., Creal, D., Koopman, S.J. and A. Lucas. (2014). Testing for Parameter Instability across Competing Modeling Frameworks. Duisenberg School of Finance - Tinbergen Institute Discussion Paper, TI 14-010/IV/DSF71.

Creal, D., S. J. Koopman and A. Lucas, 2011, A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations, Journal of Business and Economic Statistics 29, 552-563.

Creal, D., S. J. Koopman and A. Lucas, 2013, Generalized autoregressive score model with applications, Journal of Applied Econometrics 28, 777-795.

Escanciano, J. C. and I. N. Lobato, 2009, An automatic portmanteau test for serial correlation, Journal of Econometrics 151, 140-149.

Godfrey, L.G. 1988. Misspecification Tests in Econometrics. Econometric Society Monograph, Cambridge, Cambridge University Press.

Harvey, A. C., 2010, Tracking a changing copula, Journal of Empirical Finance 17, 485-500.

Harvey, A. C., 2013, Dynamic Models for Volatility and Heavy Tails, Econometric Society Monograph, New York, Cambridge University Press.

Harvey, A. C. and M. Streibel, 1998, Testing for a slowly changing level with special reference to stochastic volatility, Journal of Econometrics 87, 167-189.

Lee, J. H. H., 1991, A Lagrange multiplier test for GARCH models, Economics Letters 37, 265-271.

Lumsdaine, R. L., 2009, Correlation, models, and risk management in challenging times, Journal of Financial Econometrics 7, 40-51.

Nyblom, J., 1989, Testing for the constancy of parameters over time, Journal of the American Statistical Association 84, 223-230.

Pierce, D. A., 1982, The asymptotic effect of substituting estimators for parameters in certain types of statistics, Annals of Statistics 10, 475–8.

Tse, Y. K., 2000, A test for constant correlations in a multivariate GARCH model, Journal of Econometrics 98 107-127.

Tse, Y. K., 2002, Residual-based diagnostics for conditional heteroscedasticity models, Econometrics Journal 5 , 358-373.