



TI 2007-025/4

Tinbergen Institute Discussion Paper

Dynamic Correlations and Optimal Hedge Ratios

*Charles S. Bos**

Phillip Gould

Vrije Universiteit Amsterdam

* Tinbergen Institute

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31

1018 WB Amsterdam

The Netherlands

Tel.: +31(0)20 551 3500

Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50

3062 PA Rotterdam

The Netherlands

Tel.: +31(0)10 408 8900

Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at
<http://www.tinbergen.nl>.

Dynamic Correlations and Optimal Hedge Ratios

Charles S. Bos* and Phillip Gould
Department of Econometrics & O.R.,
Vrije Universiteit Amsterdam, The Netherlands

February 18, 2007

Abstract

The focus of this article is using dynamic correlation models for the calculation of minimum variance hedge ratios between pairs of assets. Finding an optimal hedge requires not only knowledge of the variability of both assets, but also of the co-movement between the two assets. For this purpose, use is made of industry standard methods, like the naive hedging or the CAPM approach, more advanced GARCH techniques including estimating BEKK or DCC models and alternatively through the use of unobserved components models. This last set comprises models with stochastically varying variances and/or correlations, denoted by the TVR, SCSV and DCSV models, and an approximation to these with a single-source-of-error setup.

Modelling the correlation explicitly is shown to produce the best hedges when applied to the simulated data. For financial time series on the daily S&P 500 cash versus futures returns, and also on weekly S&P 500 versus FTSE 100 returns, the correlations are compared to a realised correlation measure, extracted from high frequency data.

Apart from the comparison of correlations, the reduction in portfolio variance produced by different hedging strategies is examined. The data suggests that the most important factor in reducing portfolio variance is the use of a flexible model for time varying volatility, rather than capturing time variation in correlations. GARCH-based models with time varying correlation are found to perform not as good on the present set of measures as the stochastic volatility models, with or without dynamic correlation.

Keywords: Dynamic correlation; multivariate GARCH; stochastic volatility; hedge ratio.

JEL classification: C32, C52, G11

1 Introduction

When combining multiple financial assets into a portfolio, the portfolio return is driven by the dynamics of the underlying asset returns. The asset returns by themselves are volatile, with periods of relatively high and low volatility interchanging. Likewise, when two assets are correlated through time, it can be expected that this correlation is changing over time as well. Taken together, these observations imply that a portfolio of multiple assets which intends to minimise the total risk will have a dynamic ratio of the underlying assets, to account for the relative volatility and the changes in correlations of the assets.

*Corresponding author: C.S. Bos, Department of Econometrics & O.R., Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. Email: cbos@feweb.vu.nl.

Apart from the construction of optimal portfolios, understanding the dynamics of the volatility and correlation between asset returns is also important for calculations of the Value-at-Risk (VaR, Jorion 1997), or for application of the Capital Asset Pricing Model (CAPM), in which the risk of a stock relative to a market index is assessed (Faff, Hillier, and Hillier 2000).

With the increasing availability of data and computing power, researchers are able to use a range of models and techniques to estimate heteroskedastic time series. The (G)ARCH framework of Engle (1982) and Bollerslev (1986) has been popular in a univariate setting. The stochastic volatility (SV) approach, introduced into the econometric literature by Harvey, Ruiz, and Shephard (1994) and Jacquier, Polson, and Rossi (1994) can also be employed, either using a (quasi) maximum likelihood or a Bayesian approach for the estimation.

In the multivariate GARCH (Bollerslev, Engle, and Wooldridge 1988) and SV (Harvey, Ruiz, and Shephard 1994; Yu and Meyer 2006) literature, correlations are often assumed to be either fixed or follow deterministically from the time variation in variance, in order to simplify the models. Lately, more attention is being paid to freely time varying, dynamic correlations. One difficulty involved in estimating multivariate GARCH models is ensuring positive-definiteness of the covariance matrix. While the BEKK parameterisation of Engle (1995) ensures this condition is met, it imposes restrictions on the GARCH structure. The non-linear constraints on the parameter vector also make estimation more difficult. The dynamic conditional correlation (DCC) model of Engle (2002) adapts GARCH models specifically for the estimation of time varying correlations, while also restricting the more general structure. Park and Switzer (1995) estimate models assuming constant correlations and use them to test for time varying hedge ratios in index cash and futures markets. Recently, Pelletier (2006) compares the DCC model to a regime switching dynamic correlation model, with an ARMACH (see Taylor 1986) structure for the variance process.

As in the case of the GARCH model, it is hard to generalise the SV model to allow for time varying correlations between multiple assets; each possible choice for the parameterisation implies a certain restriction in either the space of the possible covariances or correlations. Also, allowing e.g. all correlations to evolve dynamically over time, can lead to a high number of parameters, even for a relatively low number of assets. Therefore, in this article we first limit ourselves to the bivariate case. Yu and Meyer (2006) provide a first attempt at estimating time varying correlations between assets. Our setup, developed independently, corresponds largely in the specification of the model.

This article intends to take a range of models used in practise and possible extensions which have not yet been applied for estimating the correlation and variance of two assets. With these models, their ability to recover the underlying correlation of the assets is compared, and also the practical effectiveness of the hedging strategies is measured through the standard deviations of hedge portfolio returns. Such a metric is not only relevant for hedge portfolios consisting of cash and futures contracts, but also for constructing and analysing CAPM betas, which are estimated in the same way as time varying hedge ratios.

Section 2 introduces the concepts behind time varying hedge ratios and betas. The different models for time varying correlations are explained in Section 3. While the dynamic correlation stochastic volatility model appeared once before, in Yu and Meyer (2006), the variant using a single-source-of-error (Ord, Snyder, Koehler, Hyndman, and Leeds 2005) approach is entirely new to the literature. Section 4 starts off with a description of a method to construct a realised correlation measure, taking into account the possible occurrence of jumps, and the overlap between the trading times of the two assets. Afterwards, four data

series are described in Section 4.2, covering data which (i) is simulated, with smoothly varying correlation, (ii) is simulated, with sudden jumps in correlation, (iii) covers the standard CAPM case of daily S&P 500 cash and futures returns and (iv) uses a longer, weekly series of S&P 500 and FTSE 100 returns. The data description is followed by a comparison of the performance of the models in reconstructing the simulated or realised correlation, and the resulting hedged portfolio standard deviation, in Sections 4.3–4.4. Finally, Section 5 gathers the main findings of the comparisons.

2 Time Varying Betas through optimising utility

There is a substantial body of literature dedicated to estimating time varying optimal hedge ratios for simple portfolios consisting of two assets. Brooks, Henry, and Persaud (2002) specifically analyse hedge portfolios consisting of positions in cash and futures markets, while Cho and Engle (1999) use similar models to investigate asymmetric time variation in CAPM betas. Though the assets could be a pair of cash and futures, this is not necessary, and a similar hedge could be set up e.g. between the price of kerosene and crude oil futures (Cobbs and Wolf 2004).

The general setup of the problem is as follows. For any two assets at time t , the investor holds β_t units of asset 2 per unit of asset 1, with prices $p_{1,t}$ and $p_{2,t}$ respectively. The continuously compounded returns for the two assets are defined as $r_{i,t} = \log p_{i,t}/p_{i,t-1}$, $i = 1, 2$ with variances $\sigma_{i,t}^2$ and covariance $\sigma_{12,t} = \rho_{12,t}\sigma_{1,t}\sigma_{2,t}$.

The rational investor seeks to maximise utility, which is a function of portfolio returns $r_{p,t} = r_{1,t} - \beta r_{2,t}$ and portfolio variance $\sigma_{p,t}^2$. For portfolio returns, the negative sign before β is used because investors are said to take a short position on e.g. the futures contract and a long position on the underlying stock.

Written in general form the relevant utility function, assuming a standard risk-averse setting with risk-aversion parameter ψ , is:

$$\begin{aligned} U [E_{t-1}(r_{p,t}), \sigma_{p,t}^2] &= E_{t-1}[r_{p,t}] - \psi\sigma_{p,t}^2 \\ &\approx 0 - \psi(\sigma_{1,t}^2 + \beta_{t-1}^2\sigma_{2,t}^2 - 2\beta_{t-1}\sigma_{12,t}) \end{aligned}$$

where the last equality is strict when we assume efficient markets and forego discussions on the cost-of-carry and possibly positive expected returns for risk-bearing assets. In this case, a rational agent will seek to maximise this utility function with respect to β_{t-1} , which is equivalent to minimising the variance of portfolio returns. The derived optimal value is

$$\beta_{t-1}^* = \frac{\sigma_{12,t}}{\sigma_{2,t}^2}. \tag{1}$$

In the time varying hedge ratio literature, the assumption of constant variances is relaxed, which then leads to time variation in the optimal¹ hedge ratio. As discussed in the previous section, different restrictions — such as constant correlations — are sometimes used to facilitate estimation. The next sections describe a range of modelling decisions which imply each their own ‘optimal’ hedging ratio.

¹Note how the precise implication of ‘optimal’ in relation to the hedge ratio depends on the model assumptions.

3 Modelling Correlation

As the focus of the article is on the relationship between two correlated assets over time, the models considered should explicitly be able to accurately track the evolution of correlation coefficients.

First, Section 3.1 provides three baseline models, which look directly at a linear relationship between the returns of the two assets. Subsequent sections try to model the variance processes of each of the series separately in more detail, and combine the series through a (time varying) correlation process. For these latter sections, the focus is on the demeaned bivariate returns process $r_t = (r_{1t}, r_{2t})'$. For simplicity, this bivariate returns process is supposed to be normally distributed according to

$$r_t \sim \mathcal{N}(0, H_t), \tag{2}$$

$$H_t = \begin{pmatrix} \sigma_{1t}^2 & \sigma_{12t} \\ \sigma_{12t} & \sigma_{2t}^2 \end{pmatrix} \equiv \begin{pmatrix} \sigma_{1t}^2 & \rho_t \sigma_{1t} \sigma_{2t} \\ \rho_t \sigma_{1t} \sigma_{2t} & \sigma_{2t}^2 \end{pmatrix}, \tag{3}$$

with H_t the variance matrix. Subsequent sections may either use the specification in terms of the variances/covariances, or the specification using the correlation ρ_t , depending on the model at hand.

Notice that the first moment of the data is not modelled at all, as the focus of this article is solely on the effects of variations and covariations. Likewise, in order to see differences in modelling assumptions for the correlations and variances more clearly, the simplifying assumption of (conditionally) normal returns is made.

In turn, the following sections introduce the naive, constant-hedge CAPM and time varying regression as baseline models (Section 3.1), the BEKK and DCC version of multivariate GARCH (Section 3.2) and the exact bivariate dynamic correlation SV vs. the QML approach using a single source of error variant (Section 3.3).

3.1 Baseline models

For comparison with the models with fully time varying correlations, three baseline models are used. The simplest is the ‘naive’ model, which hedges the assets fully, fixing $\beta \equiv 1$ throughout the sample.

Of course, two distinct assets are never fully correlated, and such a ‘naive’ hedge is not ideal. The standard approach in the CAPM literature is to look at the time-invariant regression

$$r_{1,t} = \alpha + \beta r_{2,t} + \epsilon_t, \tag{4}$$

relating a stock asset with return r_1 to a market factor (r_2). The CAPM β is usually estimated simply by OLS, though it is hard to assume homogeneity of the variance of ϵ in this equation.

As a third baseline, some flexibility can be added to the CAPM model by allowing β to be time varying, as the optimal hedge ratio can also be time varying (see also Equation (1)). Writing the regression in terms of an unobserved random walk β_t gives

$$\begin{aligned} r_{1,t} &= \alpha + \beta_t r_{2,t} + \epsilon_t, & \epsilon_t &\sim \mathcal{N}(0, \sigma_\epsilon^2), \\ \beta_{t+1} &= \beta_t + \eta_t, & \eta_t &\sim \mathcal{N}(0, \sigma_\eta^2). \end{aligned} \tag{5}$$

Such a model in state space (Durbin and Koopman 2001) can be estimated using the Kalman equations, allowing the optimal time varying β_t to be extracted. This model will be indicated as the time varying regression (TVR) model.

Each of these three models makes its own set of assumptions on the underlying data generating process. Especially, all three baseline models effectively assume constant variances for both assets, with only the TVR model allowing for some variation in the relation between the two. Accordingly, each of the resulting hedging decisions will be more or less ‘optimal’, as shall be seen in the applications.

3.2 GARCH Analysis

After the inception of the GARCH models (Engle 1982; Bollerslev 1986), multivariate variants quickly started appearing. A most general version, the multivariate *vec* GARCH model introduced in Bollerslev, Engle, and Wooldridge (1988) relates all elements of the covariance matrix H_t of the returns equation (2) to all cross-products of returns in the form

$$\text{vec}(H_t) = \text{vec}(C) + A \text{vec}(r_{t-1}r'_{t-1}) + B \text{vec}(H_{t-1}). \quad (7)$$

As there are here no explicit restrictions on A, B and C which guarantee that the covariance matrix H_t is positive definite, several variants have been introduced. Two of these variants are considered in this paper.

The BEKK model

One of the earlier is the BEKK GARCH model, see e.g. Engle (1995). Here, instead of specifying separate equations for all covariance matrix elements, a matrix formulation is used which ensures from the start the positive definiteness of H_t at all time points. The BEKK updating equations for the covariance matrix are of the form

$$H_t = C^*C^* + A^*r_{t-1}r'_{t-1}A^* + B^*H_{t-1}B^* \quad (7')$$

where A^* and B^* are symmetric matrices, and C^* is upper-diagonal.

In the bivariate case, this leads to 11 parameters requiring estimation, instead of 21 parameters for the corresponding *vec* specification (7). A nonlinear restriction on the parameters of A^* and B^* is needed to ensure that the eigenvalues of $A^* + B^*$ lie within the unit circle, for stationarity.

The DCC model

Where the BEKK model describes a process for the variances and covariances of the returns, the Dynamic Conditional Correlation model (Engle 2002, DCC) explicitly recognises that the interesting dynamics lie in the realm of the correlations instead of the covariances, and has become a popular multivariate GARCH representation. The specification of the covariance

matrix is

$$H_t = D_t R_t D_t, \quad (7'')$$

$$D_t^2 = \text{diag}[\omega_i] + \text{diag}[\kappa_i] \circ r_{t-1} r'_{t-1} + \text{diag}[\lambda_i] \circ D_{t-1}^2, \quad (8)$$

$$u_t = D_t^{-1} r_t, \quad (9)$$

$$Q_t = (1 - \lambda)(u_{t-1} u'_{t-1}) + \lambda Q_{t-1}, \quad (10)$$

$$R_t = \text{diag}[Q_t]^{-1} Q_t \text{diag}[Q_t]^{-1}, \quad (11)$$

where D_t is a diagonal matrix with the variances of all of the components and R_t is a correlation matrix. Positive definiteness of R_t is ensured by transformation (11). All other matrices define the dynamics of correlations and variances in a similar way to the BEKK GARCH specification. The updating specification (10) for Q_t implies that the (transformed) correlations are integrated of the first order, although AR(1) or other specifications can also be used.

Estimation of BEKK and DCC models

Estimation of the parameters of the BEKK model is done by optimising the likelihood function. Given a vector of parameters θ , the likelihood of the data can be calculated. Using either standard Gauss-Newton techniques or applying some more robust estimation methods (as the likelihood function may be not very well behaved in practical cases), the likelihood can be optimised over θ .

For the DCC model a similar approach can be used. However, it can be shown that the likelihood function L separates into a part L_v pertaining to the volatility components, and a part L_c containing only parameters for the correlation (see Engle 2002 for details). The two parts can be optimised separately, effectively breaking down one high-dimensional optimisation into two, simpler, low-dimensional problems.

3.3 Unobserved components for Stochastic Volatility models

The multivariate extension of the SV model, for the constant correlation case is already proposed in Harvey, Ruiz, and Shephard (1994). Also a multi-factor SV model is a common extension, but to get to a separate time varying correlation factor is more difficult. To our knowledge, the bivariate Dynamic Correlation Stochastic Volatility (DCSV) model used here is only presented before in Yu and Meyer (2006). Fixing the dynamic correlation back to the static case leads to the Static Correlation SV model (SCSV), which is used for comparison.

Apart from these stochastic volatility models, a quasi maximum likelihood approach using a Single Source of Error Model (SSOE, see Ord, Snyder, Koehler, Hyndman, and Leeds 2005) is presented in a subsequent section.

Dynamic Correlation SV

The starting point for the DCSV model is again the returns equation (2), stating that $r_t \sim \mathcal{N}(0, H_t)$, with the correlation element in the variance matrix H_t of (3) modelled as

$$\rho_t \equiv \frac{\exp(q_t) - 1}{\exp(q_t) + 1} = 2 \frac{1}{\exp - q_t + 1} - 1, \quad (12)$$

$$q_{t+1} = q_t + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2). \quad (13)$$

Both log-price returns have SV-type variances, according to

$$\sigma_{it}^2 = \exp(h_{it}) \tag{14}$$

$$h_{it+1} = \gamma_i + \phi(h_{it} - \gamma_i) + \xi_{it} \qquad \xi_t = \mathcal{N}(0, \Sigma_\xi) \tag{15}$$

The two returns are related through the correlation coefficient ρ_t , which varies over time according to a transformed random walk (13). The transformation in (12), which is a rescaled Sigmoid function, assures that the correlation is bounded in $(-1, 1)$.

As there is little information on the value of ρ_t at each time point t , it does not seem useful to specify a more refined process for the correlation process, or a different transformation. As a simplification, fixing $\sigma_\eta \equiv 0$ gets us back to the static correlation model with dual stochastic volatility (SCSV), which is used below for comparison.

The two SV innovations ξ_{it} could display correlation as well, though in this article we limit ourselves to the situation where $\Sigma_\xi \equiv \text{diag}(\sigma_{\xi,1}^2, \sigma_{\xi,2}^2)$. First of all, the novelty of introducing dynamic correlation in a stochastic volatility model already imposes sufficient technical difficulties. Secondly, it is at present unclear how well identified a correlation coefficient between two innovation processes of unobserved volatility components would be.

A QML Approximation using SSOE

The major drawback of the DCSV model of the previous section is that it is fully nonlinear, based on unobserved components. This leads to the need for advanced Bayesian sampling techniques (more on this topic below) to get to an estimate of parameters, variances and correlations.

In the article by Harvey, Ruiz, and Shephard (1994), a quasi maximum likelihood (QML) approach was used to estimate the stochastic volatility model. Their approach entailed linearising the model to get to a model which can be estimated using standard Kalman filtering techniques (Durbin and Koopman 2001). The linearised version of the univariate SV model for r_t reads

$$\ln r_t^2 = h_t + \nu_t \tag{16}$$

$$h_{t+1} = \gamma + \phi(h_t - \gamma) + \xi_t \tag{15'}$$

where $\nu_t = \ln u_t^2$, $u_t \sim \mathcal{N}(0, 1)$, has a non-standard density. The density of ν_t is a transformation from the standard normal density of u_t , with mean -1.27 and variance $\pi^2/2$. When the Kalman equations are used in the above setting, effectively a normal approximation to the density of ν_t is used, and hence the resulting estimator is only a quasi-maximum likelihood estimator. On the other hand, the approach has the appeal of being both intuitive and easy to estimate, using the linear Kalman filter. The approach is not efficient, so results will be an approximation of the outcomes of e.g. a Bayesian estimation approach of the exact model.

When the correlations ρ_t are introduced as in the DCSV model, a new source of non-linearity appears. A general approach to cope with non-linear state space models is the single source of error (SSOE) framework, presented in detail by Ord, Snyder, Koehler, Hyndman, and Leeds (2005). Here, instead of allowing separate independent disturbances e_t and ξ_t as in the equations above, in the multiple source of error approach, full correlation between these disturbances is assumed, taking $\xi_t \equiv \psi \nu_{t-1}$ for some constant parameter ψ which is to be estimated. Note how this provides a closer link to the GARCH approach, where also the innovation in the return equation drives the changes in the variance equation.

To link together a bivariate system of observations and the correlation sequence, the full model is specified as

$$r_{1t} = \alpha + \beta_t r_{2t} + \nu_{0t}, \quad (5')$$

$$\ln r_{it}^2 = h_{it} + \nu_{it}, \quad i = 1, 2 \quad (16')$$

$$\beta_t = \rho_t \sqrt{\frac{\exp(h_{1t})}{\exp(h_{2t})}}, \quad \rho_t = \frac{\exp(q_t) - 1}{\exp(q_t) + 1}, \quad (12')$$

$$q_{t+1} = q_t + \psi_0 \nu_{0t}, \quad (13')$$

$$h_{it+1} = \gamma_i + \phi_i (h_{it} - \gamma_i) + \psi_i \nu_{it}. \quad (15'')$$

In this setup, the time varying regression (5) of Section 2 is used to obtain more information on the dynamics of the variance and covariance terms. This equation is combined with two QML observation equations, (16'), while the three state equations describing q_t and h_{it} are linked using smoothing parameters ψ_0, ψ_1, ψ_2 to the disturbances of the observation equations.

Estimation of the SV models

The advantage of the SSOE-SV models is that they are able to handle non-linearity easily, as the above set of equations can be solved recursively for a fixed set of parameters. This allows the likelihood function to be calculated in a direct manner, optimising over the parameters using standard methods.

One drawback is the lack of a smoother for SSOE models, so this approach is most suitable for forecasting. Secondly, in a linear setup there is a close correspondence between the outcome of a multiple source of error model and its SSOE counterpart; for the non-linear case, there is no such clear correspondence, and results will depend on the precise implementation of the links between the disturbances in the model.

Estimation of the DCSV model is more demanding, as the likelihood is not available in closed form without integrating out the unobserved components $h_{it}, i = 1, 2$ and $\rho_t, t = 1, \dots, T$ for volatility and correlation. Therefore, estimation is done here using a Bayesian approach with data augmentation, using a Markov chain Monte Carlo method.

The algorithm proceeds, after initialising the parameters $\theta = (\phi_i, \gamma_i, \sigma_{i,\xi}, \sigma_\eta), i = 1, 2$ and states, by iterating over the following steps:

- i Sampling a new vector of ρ , by successively sampling $\rho_t | \rho_{t-1}, \rho_{t+1}, y_t, h_t, \theta$, for $t = 1, \dots, T$. As this density is not available in closed form, instead a random walk Metropolis-Hastings step is used with to sample from the posterior density

$$P(q_t | q_{t-1}, q_{t+1}, y_t, h_t, \theta) \propto P(q_t | q_{t-1}, q_{t+1}, \sigma_\eta) \times \mathcal{L}(y_t; h_t, \rho_t(q_t)).$$

Both latter densities are simple normals, such that sampling a new value of q_t is not difficult. After sampling q_t , it is transformed back to ρ_t ;

- ii Sampling two new vectors of h jointly, from the density of $h_t | h_{t-1}, h_{t+1}, y_t, \rho_t, \theta$. Again, the density is not easily tractable in closed form, but the posterior is again a combination of the density of $h_t | h_{t-1}, h_{t+1}, \theta$ and the likelihood of the present observation $\mathcal{L}(y_t | h_t, \rho_t, \theta)$.

iii Assuming a prior $\pi(\sigma_\eta) \sim \text{IG-1}(\alpha_\eta, \beta_\eta)$, the posterior of σ_η is

$$P(\sigma_\eta|\rho) \sim \text{IG-1} \left(\alpha = \frac{T-1}{2} + \alpha_\eta, \beta = \left(\sum_{t=2}^T \frac{(q_t - q_{t-1})^2}{2} + \frac{1}{\beta_\eta} \right)^{-1} \right)$$

iv The parameters $\gamma_i|h_i$ follow a simple normal density, assuming a normal prior with mean μ_γ and variance σ_γ^2 ;

v The remaining parameters $\phi_i, \sigma_i \xi$ are sampled per asset i using another Metropolis-Hastings step, with a random walk normal candidate density.

This algorithm results, after performing a sufficient number of iterations, in a sample from the posterior density of the parameters θ and states h, ρ . For the analysis in subsequent sections, the posterior mode of the parameters θ estimated over the available sample is used as input for a particle filter (Pitt and Shephard 1999), to extract filtered estimates of the states $\rho_t, h_t|y_1, \dots, y_t$, conditioning only on past and present information. This gives a more fair comparison than using the output of the MCMC chain, which describes the distribution of the states $\rho_t, h_t|y_1, \dots, y_T$ conditioning on the full data set.

Table 1: Prior densities and moments

Parameter	Density	Prior parameters		μ	σ_π
σ_η	IG-1	$\alpha_\eta = 1.3$	$\beta_\eta = 100$	0.13	0.13
γ_i	\mathcal{N}	$\mu_\gamma = 0$	$\sigma_\gamma = 2$	0	2
ϕ_i	Beta	$\alpha_\phi = 12$	$\beta_\phi = 3$	0.8	0.1
$\sigma_i \xi$	IG-1	$\alpha_\xi = 1.3$	$\beta_\xi = 300$	0.075	0.075

A Bayesian procedure needs some prior information on the parameters. These priors were fixed using information on the expected size of the parameters based on related research, and taken with sufficient spread to allow the present data to decide on the location of the posterior. The parameters, prior density family, prior parameters and moments of the prior are given in Table 1. For all data series, these same priors are used.

Alternatively, Sandmann and Koopman (1998), with further extensions by Jungbacker and Koopman (2005), provide a classical maximum likelihood approach for models with stochastic variance, applying importance sampling to simulate out the unobserved states. This method could be extended to include the time varying correlation as well.

4 Evaluating Model Performance

In this section the performance of the models is evaluated in four different settings. In Section 4.4 the portfolio standard deviation resulting from using the different modelling setups is compared. Section 4.1 describes an alternative manner of comparing the models by the construction of a realised covariation measure, based on theory by Andersen and Bollerslev (1998) and Barndorff-Nielsen and Shephard (2004a). The steps needed to extract the measure for data sets in practice are introduced.

With this realised covariation measure introduced, Section 4.2 introduces two data generating processes for analysis of simulated series in a controlled setting, and two sets of financial

data. It is followed by Section 4.3, which starts the evaluation of the model performance comparing the estimated correlation with either the simulated or the realised correlation. The aforementioned Section 4.4 compares the performance on the basis of the portfolio standard deviation.

4.1 Realised covariation

The models presented in Section 3 deliver an estimate of the variances and correlations between two assets. In a simulation exercise, these estimates can be compared with the true series underlying the simulated returns, but in applied work, no immediate comparison is possible. Hence, an estimate of possible time-varying variances and correlations is useful for judging whether a model's estimate of these quantities can be considered acceptable.

In this paper an approach based on realised volatilities and realised correlations is taken. Realised volatility is a concept discussed in e.g. Andersen and Bollerslev (1998), extended to realised covariation in Barndorff-Nielsen and Shephard (2004a).

When M intra-day² (demeaned) returns $y_{j,i}, j = 1, \dots, M$ are available on day i , then the realised covariation matrix for day i is defined as

$$RV_i = [y_M^*]_i = \sum_{j=1}^M y_{j,i} y'_{j,i}.$$

The diagonal elements of the realised covariation matrix equal the original measure of realised variation, whereas the off-diagonal elements can be used to construct an estimate of the correlation between assets within a day.

Several considerations are necessary for controlling the precision of the correlation estimates. First of all, obviously the returns within the vector $y_{j,i}$ should cover the same intra-day period. For assets traded on geographically distinct markets, only intra-day returns of the overlapping time period can be used.

Secondly, to obtain a good estimate of the daily covariation or correlation, the number M of intra-day returns should be large enough. This can be troublesome due to data availability as high-frequency data sets can be hard to come by, especially over longer time horizons.

Thirdly, the intra-day returns should not be of too high frequency either. When the frequency is increased beyond 5-minute data, more intervals with few or no trades can be found, leading to either a missing observation or a zero return, depending on the interpretation. Other micro-structure type effects may occur as well, hampering the estimation of the daily correlation. See e.g. Ait-Sahalia, Mykland, and Zhang (2005) for a discussion of the optimal sampling frequency in the presence of micro-structure noise.

A further consideration is the possible occurrence of jumps within series of asset prices. Though on a daily basis a price process may seem smooth, intra-daily it is found that sudden jumps can occur. If such a jump would occur in only one of the two series, it can seriously alter the estimate of the correlation between the series, and therefore it would be convenient to detect if jumps occur.

The theory of jumps and their detection is described in articles of Andersen, Bollerslev, and Diebold (2005) and Barndorff-Nielsen and Shephard (2006). A possible procedure is to

²In the following sections, apart from daily covariation also weekly covariation is estimated. For simplicity of argument, this section uses a day as the period of interest.

calculate the bipower variation of series a in the vector y as

$$\text{BPV}_{i,a} = \sum_{j=2}^M |y_{j,i,a}| |y_{j-1,i,a}| / \mu^2, \quad \mu \equiv \sqrt{2/\pi}.$$

The bipower variation gives an estimate of the integrated variance excluding possible jumps, whereas the realised variation still includes those jumps. Hence, the difference between the diagonal elements of RV and BPV, following Barndorff-Nielsen and Shephard (2004b), can be used to estimate the jump process free of negative squared jumps as

$$J_{i,a} \equiv \max((\text{diag RV}_i)_a - \text{BPV}_{i,a}, 0)$$

To test whether a jump is significant, Barndorff-Nielsen and Shephard (2006) use (among others) a test statistic which they call the feasible linear jump statistic, reading

$$\hat{G}_{i,a} = \frac{(\text{diag RV}_i)_a - \text{BPV}_{i,a}}{\sqrt{\delta \theta \text{QPV}_{i,a}}} \xrightarrow{H_0} \mathcal{N}(0, 1)$$

under the hypothesis H_0 of no jump. $\text{QPV}_{i,a}$ in this formula is the realised quadpower variation,

$$\text{QPV}_{i,a} = \frac{1}{\delta} \sum_{j=4}^M |y_{j,i,a}| |y_{j-1,i,a}| |y_{j-2,i,a}| |y_{j-3,i,a}| / \mu^4,$$

and $\theta \equiv \pi^2/4 + \pi - 5$.

To summarise, in following sections on financial data a realised correlation measure will be constructed by

- using data on the 5-minute frequency, treating periods in which no trades occurred as missing;
- applying above procedure for testing for the occurrence of jumps in all elements of the vector of returns y , with a significance level of $\alpha = 0.05$; when a jump is detected, the largest absolute return in the series is deleted, and the testing procedure is repeated until no further jumps are found;
- skipping data on the first 15 minutes of each trading day, to remove initial effects;
- for days with too few intra-day returns, setting the realised correlation measure to a missing.

Clearly, these choices of operationalising the construction of an approximate measure of correlation are debatable, and that a resulting series of correlations should be used as a guideline. It was however found that this procedure is robust against e.g. changing the significance level, the amount of data disregarded from the beginning of the day of returns etc. Including also observations detected as jumps by the above procedure leads to a more erratic behaviour of the estimate of the realised correlations, with more days displaying lower correlation than average.

4.2 Data

As mentioned before, four sets of data are used in subsequent section for comparing the performance of different models in recovering the underlying correlation, and in constructing a minimum-variance hedge portfolio.

The first two data sets are simulated series. In Engle (2002), a range of different GARCH models are compared in their ability to extract a time varying correlation structure from bivariate returns. The variances in the covariance matrix H_t of Equation (3) are specified as

$$\begin{aligned}\sigma_{1,t}^2 &= .01 + .05y_{1,t-1}^2 + .94\sigma_{1,t-1}^2, \\ \sigma_{2,t}^2 &= .5 + .2y_{2,t-1}^2 + .5\sigma_{2,t-1}^2,\end{aligned}$$

implying a GARCH process for both returns. For the time varying correlation, several deterministic processes are proposed. Here we look at two cases,

$$\begin{aligned}\rho_t &= .5 + .4 \cos(2\pi t/200), && \text{Slow sine} \\ \rho_t &= \text{mod}(t/200), && \text{Ramp}\end{aligned}$$

either providing a slow sine wave for the correlation, or a ramp structure.

Data is generated for a period of $N = 1000$ observations. Of these 1000 observations, and of the two applied data sets, 90% is used for estimating the parameters, while 10% is preserved as a hold-out period for assessing out-of-sample performance.

The latter two data sets compare the S&P 500 index to either its own future, or to the FTSE 100 index. In the first case, daily data over the period January 7, 1998–December 28, 2006 is used, for a total of $N = 2266$ observations. The data are obtained from DiskTrading³, and concern the daily index or future price at closing time. For this same period, data is also available from the same source at the 5 minute frequency, to construct the realised correlation measure as presented in Section 4.1. This data set allows to investigate the common situation of an investor trying to hedge his or her risk using a future on the same stock.

The second applied data set compares the S&P 500 as traded in New York to the FTSE 100 index, traded in London. From Yahoo Finance⁴, a longer series of daily prices is available. Of these, the Friday closing price (or last day of the week, in case the market is closed on Friday) is taken to construct a weekly time series over the period January 4, 1985–December 31, 2006, for a total of $N = 1147$ observations. Such a longer time period, with assets trading on two geographically distinct markets, can display the robustness of the models to more variation in correlation. It however also hampers the construction of a reasonable realised correlation measure: For the FTSE index, we only have available data from DiskTrading at the 5 minutes frequency for the period of August 21, 1998 until December 10, 2004. The realised correlation series can be extended until the end of December 2006 by using a high-frequency future series on the FTSE 100, which is available from September 2002–December 2006.⁵ Over this subperiod a comparison of the realised correlation to the model based correlation can be made, again remembering that the realised correlation measure is a rough approximation to the true underlying correlation.

³See <http://www.disktrading.com>.

⁴See <http://finance.yahoo.com>, symbols SNP:~GSPC and FSI:~FTSE.

⁵Effectively, over the period where both high frequency FTSE and future data is available, both are used for calculating a realised correlation measure against the S&P 500 series. Both correlations line up reasonably well, apart from a roughly constant difference as the future is less correlated with the S&P 500 series as the FTSE index itself. For the latter part of the sample, the correlation between FTSE future and the S&P 500 is adapted for this difference.

Table 2: Data availability

Data	In-sample	Out-of-sample	Correlation
Sine	1-900	901-1000	1-1000
Ramp	1-900	901-1000	1-1000
S&P 500-Future	1998/1/2-2006/2/3	2006/2/6-2006/12/28	1998/1/2-2006/12/28
S&P 500-FTSE	1985/1/4-2004/10/15	2004/10/22-2006/12/29	1998/8/21-2004/12/10 + 2002/9/27-2006/12/29

As a summary, Table 2 reiterates the availability of the data for the four data sets, and the periods which are used for estimation, ie. the first 90% of the data, the out-of-sample period, and the period for which a correlation measure is available.

4.3 Recovering the correlation structure

Before moving on to the hedging performance of the different model, in the next section, here the question is whether these models are able to reconstruct the correlation sequence from the bivariate data series.

Table 3: MSE of correlation measure, for simulated data

	MSE($\hat{\rho}$) Sine		MSE($\hat{\rho}$) Ramp	
	in-sample	out-of-sample	in-sample	out-of-sample
Naive	32.96	33.40	36.38	8.46
CAPM	8.12	8.10	8.41	8.64
TVR	4.80	4.72	7.31	3.10
BEKK	3.39	3.67	5.27	3.44
DCC	3.11	2.78	4.78	2.18
DCSV	3.36	4.25	6.32	2.62
SCSV	8.18	8.21	8.45	7.96
SSOESV	4.77	5.23	6.90	4.15

Note: The table reports the average mean squared error ($\times 100$) of the estimated correlation, both in-sample (the first 90% of the observations) and out-of-sample (over the last 10% of the simulated data), over 100 iterations of generating data from the DGP and estimating the models.

For the simulated data sets, the true correlation is known. Table 3 displays the average mean squared error (MSE) for the models, and the percentage difference with the best model. The MSE was averaged over 100 repetitions of generating data and reestimating the models. For the models estimated using the Gibbs sampler with data augmentation, a posterior sample of 1000 iterations is collected after allowing a burn-in period of 100 iterations.

Clearly, fixing the correlation at 1 as the Naive model does, is not a good way to recover the true correlation. The fixed correlations of the CAPM and SCSV models are likewise not good at estimating either a sine or a ramp structure.

Closest correspondence between the true and estimated correlations is found with the DCC model, with the DCSV and BEKK following closely behind. This indicates that modelling the correlation seriously does indeed pay off. The approximate SSOESV model is behaving roughly on par with the TVR model, slightly better even in sample.

For the applications of measuring correlation between the S&P 500 index and either its future or the FTSE 100 index, the MCMC method was run for the DCSV and SCSV models until a sample of size 10.000 parameter vectors was collected, after a burn-in period of 1.000 iterations. On this data, one has to take into account that the measure of comparison, the realised correlation, is only an approximate measure of correlation, and that even this approximate measure is only available for a subperiod in the case of the S&P 500 vs the FTSE (see Table 2). Therefore, Table 4 should only be used as a reference point, and not for drawing hard conclusions as to which model best fits reality.

Table 4: MSE of correlation measure vs realised correlation, for financial data

	MSE($\hat{\rho}$) S&P vs Future		MSE($\hat{\rho}$) S&P vs FTSE	
	in-sample	out-of-sample	in-sample	out-of-sample
Naive	3.08	0.56	21.11	12.63
CAPM	2.23	0.23	2.98	1.82
TVR	1.96	0.54	7.70	3.28
BEKK	2.24	0.24	5.56	2.96
DCC	2.10	0.21	4.42	1.54
DCSV	2.04	0.25	4.43	1.70
SCSV	2.33	0.26	2.82	2.19
SSOESV	2.03	0.14	2.59	2.04

Note: The table reports mean squared error ($\times 100$) of the estimated correlation, as compared with the realised correlation calculated on 5-minutes returns, both in-sample (the first 90% of the observations) and out-of-sample (over the last 10% of the data), for as far as the high frequency returns are available.

As the correlation between S&P 500 and its future is high, especially on a daily basis, taking the naive stance of putting it to one is not that terribly bad, as seen in the first two columns of Table 4. The CAPM and SCSV models fix the correlation at values of around 0.97, so the MSE for these models is not much different from the one of the naive model. The other models manage to lower their MSEs somewhat, down to 1.96 for the time varying regression model. For the S&P 500 vs FTSE results, the CAPM model gives lowest MSE. To understand where these results come from, it is instructive to discuss plots of the estimated correlation, and the way they relate to the realised correlation over the sample.

Figure 1 displays the estimated correlations, and from it is seen how the results especially for the TVR, DCC, DCSV and SSOESV models are not so different: A consistently high correlation is found, such that also the MSE will be similar between these models. Comparing the correlations of the models which are closest to the realised correlation with this correlation measure itself, in Figure 2, shows that throughout the sample, the realised correlation ends up consistently lower than the model-based measures.

Also, the realised correlation as extracted from the high frequency returns is very volatile, even after removing obvious jumps from the time series. Throughout 2000-2001, the number of trades in the data set on the futures is going down, to jump up again at the start of 2002. This is apparent from a distinctly higher estimate of the correlation from this moment onwards. Changing the frequency to minute-by-minute data does not significantly alter these results. This can serve as a warning message, that finding a time-varying estimate for the correlation between two high frequency time series may be difficult, and one would do better to judge the model performance on the basis of e.g. the portfolio hedging in the next section.




Figure 1: Model based correlation measures of S&P 500 vs its future




Figure 2: Correlation of S&P 500 vs its future, for the TVR, DCSV and SSOESV models, as compared to the realised correlation

For columns 3 and 4 of Table 4, remember that the MSE of the correlation is only computed over the period August 1998–December 2006, where the high frequency data for the S&P 500 and the FTSE indices was available. Over this period, the correlation resulting from the SSOESV model seems closer to the realised correlation, though out-of-sample the DCC and DCSV model seem to perform better again.

Figures 3–4 give corresponding plots on the estimated correlations, with for the period starting in August 1998 the comparison with the realised correlation. In this case, there is less discrepancy between the realised and model-based measures. The three model-based measures follow the increase in correlation between 2000–2002, with a slight decline indicated by the DCC and DCSV measures after 2003.




Figure 3: Model based correlation measures of S&P 500 vs FTSE 100




Figure 4: Correlation of S&P 500 vs its future, for the DCC, DCSV and SSOESV models, as compared to the realised correlation, over the subperiod starting in August 1998

4.4 Lowering the portfolio risk

The main idea behind the CAPM is to construct a portfolio which would lower the risk of holding the asset by itself, through hedging. Both in the case of simulated data as in the applications, such a portfolio based on the different models can be constructed to get to the lowest possible portfolio standard deviation. Tables 5–6 display this standard deviation, both in- and out-of-sample, together with the percentage difference with the lowest standard deviation.

For building a low-variance portfolio, it clearly is important to know not only the correlation between the two assets, but also the relative variances. For the simulated data series, the lowest risk is obtained using the DCSV model, both in- and out-of-sample, apart from one virtual tie with the SSOESV in the out-of-sample sine data case. For the more extreme ramp data, the SSOESV is not behaving as good: A detailed look at the estimated corre-

Table 5: Portfolio standard deviation, simulated data

	$\sigma(r_p)$ Sine				$\sigma(r_p)$ Ramp			
	in-sample		out-of-sample		in-sample		out-of-sample	
Naive	1.036	[39.7]	1.061	[38.5]	1.056	[40.4]	0.763	[26.0]
CAPM	0.824	[11.0]	0.845	[10.3]	0.834	[11.0]	0.695	[14.9]
TVR	0.870	[17.3]	0.824	[7.5]	0.821	[9.2]	0.656	[8.4]
BEKK	0.804	[8.3]	0.843	[10.0]	0.818	[8.8]	0.675	[11.6]
DCC	0.801	[7.9]	0.816	[6.4]	0.820	[9.1]	0.660	[9.1]
DCSV	0.742	[0.0]	0.768	[0.2]	0.752	[0.0]	0.605	[0.0]
SCSV	0.814	[9.8]	0.842	[9.9]	0.824	[9.6]	0.680	[12.5]
SSOESV	0.747	[0.7]	0.766	[0.0]	0.880	[17.1]	0.645	[6.7]

Note: The table reports the portfolio standard deviation, both in-sample (first 90% of the data) and out-of-sample, with the percentage difference with the lowest standard deviation reported between square brackets.

Table 6: Portfolio standard deviation, financial data

	$\sigma(r_p)$ S&P vs Future				$\sigma(r_p)$ S&P vs FTSE			
	in-sample		out-of-sample		in-sample		out-of-sample	
Naive	0.300	[13.9]	0.169	[7.5]	1.963	[18.7]	1.074	[13.4]
CAPM	0.294	[11.8]	0.169	[7.7]	1.731	[4.6]	0.984	[3.9]
TVR	0.295	[11.9]	0.168	[7.5]	1.704	[3.0]	0.984	[3.9]
BEKK	0.295	[11.9]	0.171	[9.3]	1.704	[3.0]	1.030	[8.8]
DCC	0.302	[14.8]	0.170	[8.2]	1.706	[3.1]	0.997	[5.3]
DCSV	0.268	[1.7]	0.159	[1.3]	1.657	[0.1]	0.947	[0.0]
SCSV	0.263	[0.0]	0.157	[0.0]	1.665	[0.6]	0.969	[2.4]
SSOESV	0.292	[10.8]	0.175	[11.7]	1.654	[0.0]	0.966	[2.1]

Note: See Table 5 for an explanation of the entries in the table.

lations indicates that the SSOESV model can sometimes run into trouble when estimating low correlations, possibly leading to explosive behaviour. In more realistic cases, this should however not be a problem.

Note how the BEKK, DCC and SCSV models consistently lag the DCSV model, at around 9% of a difference in standard deviation, and how the simple CAPM model doesn't do so bad, with a standard deviation of the portfolio which lies on average 11% above the result of the DCSV model, in the case of the simulated data.

More interesting are the real-world results for the financial data hedging, in Table 6. For the S&P 500 against the futures data, in the first four columns, the true SV models take the lead, where the generalisation from the SCSV to time varying correlation does not seem to be as important as the introduction of stochastic volatility itself. The other models result in rather similar portfolio standard deviations, at least for this data set, both in- and out-of-sample.

For the S&P 500 vs FTSE data, both DCSV and SSOESV track the volatility and correlation well, resulting in a low hedged portfolio variance. The multivariate GARCH models, BEKK and DCC, do not deliver any advantage as compared to either the simpler TVR model, and perform out-of-sample even worse than the simplest CAPM.

It is important to note that the closest correspondence between the correlation estimate with the true or (estimated) realised correlation is not predictive of optimal variance minimisation for the hedged portfolio.

5 Conclusions

This article studied the possibility of building a hedge portfolio out of two assets, using industry-standard approaches like the CAPM or with a naive hedge, and compared these to more advanced approaches applying BEKK and DCC GARCH-type models. Furthermore, novel methods are proposed, like allowing for time variation in the CAPM using a time varying regression model, or using an unobserved components approach with stochastic volatility and stochastic correlation. For comparison, also a static correlation SV and an approximate single-source-of-error-SV model are introduced, of which the latter is entirely novel.

In order to judge the relative performance of the models, two criteria are considered: Either the precision with which the correlation is estimated, or, more importantly, the resulting variability of the hedged portfolio.

For simulated data, the more elaborate, non-approximative models BEKK, DCC and DCSV perform better in recouping the correlation, especially when the underlying correlation comes from a smooth process like the simulated sine. For the financial series it is found that the construction of a reliable correlation measure based on intraday returns is not as straightforward as could be hoped for.

Using an advanced method of constructing a measure of realised correlation, adapting for jumps in the return process, taking great care in only using returns which correspond to the same time period for both series, leaves a rather volatile correlation series which seems hampered by a lack of clean data, especially for the S&P 500 futures data pre-2002. Therefore, reported measures on the MSE of the correlation (which maybe should be called a mean squared difference with the realised correlation), is not taken too seriously in comparing the models using real data.

A more definite result of the computation of model-based correlations is that for the S&P 500 vs its future the correlation has been high, around 0.97, and nearly constant, throughout the entire sample period 1998–2006. The model-based estimates agree in this respect to a high extent.

For the longer time series concerning the S&P 500 and FTSE indices, the models agree that overall correlation increased, from 0.4-0.6 in 1984 up to 0.8 at the end of 2006. However the path of the correlation is estimated differently for the models used.

This distinction between the models is highlighted further in the resulting variability of the optimal portfolio that can be constructed using the models. For both the simulated and financial data, models of the SV class are preferred for lowering the variability. For the S&P 500-futures data, the static correlation SV model results in lowest variability of the hedged portfolio. As the correlation between these assets was found to be nearly constant throughout the sample, intricate modelling seems not to help much, apart from getting a good grasp on the volatility in each of the series. For the other series, DCSV or even the approximate SSOESV performs better, also lowering the standard deviation of the portfolio by 3-14.8% as compared to the GARCH approaches.

The fact that this latter result also holds for the simulated data is quite remarkable: The DGP of the volatilities in the simulation was of the GARCH-type, with only the correlation

simulated in a way differing from the BEKK or DCC specifications. Even so, the fully misspecified DCSV model delivers a better performing hedged portfolio.

Only bivariate models have been applied in this article, in order to grasp the importance of modelling correlation correctly in a simple hedge portfolio of two assets. The extension to multiple assets, though interesting in its own right, would obfuscate one of the main findings that, at least for the series considered here, applying a flexible volatility model is at least as important as allowing the correlation to change over time. However, an extension to multivariate DCC-GARCH vs. DCSV along the lines proposed in Yu and Meyer (2006), effectively following the framework of Engle (2002), could be an interesting possibility.

Furthermore, the approach taken in this article of extracting a measure of realised correlation was a very practical approach: Available data was used as best as possible, but further detail in results and further finetuning of the algorithm, with possibly a combination of different frequencies for computing the correlation measure, could deliver a better comparative measure.

References

- Aït-Sahalia, Y., P. A. Mykland, and L. Zhang (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies* 18(2), 351–416.
- Andersen, T. G. and T. Bollerslev (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39(4), 885–905.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold (2005). Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility. Working paper 11775, NBER. Forthcoming in *Review of Economics and Statistics*.
- Barndorff-Nielsen, O. E. and N. Shephard (2004a). Econometric analysis of realised covariation: High frequency based covariance, regression and correlation in financial economics. *Econometrica* 72(3), 885–925.
- Barndorff-Nielsen, O. E. and N. Shephard (2004b). Power and bipower variation with stochastic volatility and jumps (with discussion). *Journal of Financial Econometrics* 2(1), 1–48.
- Barndorff-Nielsen, O. E. and N. Shephard (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Economics* 4(1), 1–30.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3), 307–327.
- Bollerslev, T., R. F. Engle, and J. M. Wooldridge (1988). A capital asset pricing model with time-varying covariances. *The Journal of Political Economy* 96(1), 116–131.
- Brooks, C., O. T. Henry, and G. Persaud (2002). The effect of asymmetries on optimal hedge ratios. *Journal of Business* 75(2), 333–352.
- Cho, Y.-H. and R. F. Engle (1999). Time-varying betas and asymmetric effects of news: Empirical analysis of blue chip stocks. Working paper W7330, NBER.

- Cobbs, R. and A. Wolf (2004). Jet fuel hedging strategies: Options available for airlines and a survey of industry practices. Unpublished manuscript, http://www.kellogg.northwestern.edu/research/fimrc/papers/jet_fuel.pdf.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflations. *Econometrica* 50, 987–1008.
- Engle, R. F. (1995). *ARCH: Selected Readings*. Advanced Texts in Econometrics. Oxford: Oxford University Press.
- Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* 20(3), 339–350.
- Faff, R. W., D. Hillier, and J. Hillier (2000). Time varying beta risk: An analysis of alternative modelling techniques. *Journal of Business Finance and Accounting* 27, 523–554. Naar verwijzing Phillip Gould.
- Harvey, A. C., E. Ruiz, and N. Shephard (1994). Multivariate stochastic variance models. *Review of Economic Studies* 61(2), 247–264.
- Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models (with discussion). *Journal of Business and Economic Statistics* 12(4), 371–417.
- Jorion, P. (1997). *Value at Risk: The New Benchmark for Controlling Market Risk*. New York: McGraw-Hill.
- Jungbacker, B. and S. J. Koopman (2005). On importance sampling for state space models. Discussion paper 05-117/4, Tinbergen Institute.
- Ord, J. K., R. D. Snyder, A. B. Koehler, R. J. Hyndman, and M. Leeds (2005). Time series forecasting: The case for the single source of error state space approach. Working Paper 7/05, Monash Econometrics and Business Statistics.
- Park, T. and L. Switzer (1995). Time-varying distribution and the optimal hedge ratios for stock index futures. *Applied Financial Economics* 5, 131–137.
- Pelletier, D. (2006). Regime switching for dynamic correlations. *Journal of Econometrics* 131(1-2), 445–473.
- Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* 94(446), 590–599.
- Sandmann, G. and S. J. Koopman (1998). Estimation of stochastic volatility models via Monte Carlo maximum likelihood. *Journal of Econometrics* 87, 271–301.
- Taylor, S. J. (1986). *Modelling Financial Time Series*. Chichester: Wiley.
- Yu, J. and R. Meyer (2006). Multivariate stochastic volatility models: Bayesian estimation and model comparison. *Econometric Reviews* 25(2-3), 361–384.