A Class of Generalized Dynamic Correlation Models

He, Zhongfang

Royal Bank of Canada

10 February 2018

A Class of Generalized Dynamic Correlation Models

Zhongfang He[*]

This version: February 10, 2018

Abstract

This paper proposes a class of parametric correlation models that apply a two-layer autoregressive-moving-average structure to the dynamics of correlation matrices. The proposed model contains the Dynamic Conditional Correlation model of Engle (2002) and the Varying Correlation model of Tse and Tsui (2002) as special cases and offers greater flexibility in a parsimonious way. Performance of the proposed model is illustrated in a simulation exercise and an application to the U.S. stock indices.

Key words: ARMA, Bayes, MCMC, multivariate GARCH, time series

1   Introduction

Modeling the time-varying covariances between variables is an importance problem in econometrics and empirical finance. An effective modeling strategy is to decompose the covariance matrix into variances and correlations (Bollerslev (1990), Barnard et al. (2000)). Univariate parametric (e.g. GARCH models of Engle (1982) and Bollerslev (1986)) or non-parametric (e.g. Gaussian process as in Lan et al. (2017)) models are then fitted to the component of variances while the correlation matrix is modeled separately. This paper adopts this separation strategy and focuses on modeling the correlation matrix.

For multivariate economic and financial time series data, the family of multivariate GARCH models has been a popular parametric tool to study their second order moments. Two popular models of dynamic correlations in the multivariate GARCH family are the Dynamic Conditional Correlation (DCC) model of Engle (2002) and the Varying Correlation (VC) model of Tse and Tsui (2002). Both models apply an autoregressive-moving-average (ARMA) structure to either the correlation matrix itself or an auxiliary variable to effectively capture the dynamics of the correlations. A key advantage of these two models is their parsimonious model structure which enables their scalability to large dimension correlation matrices directly or after further extension (e.g. Engle et al. (2017)). Since their introduction, various variants of the DCC and VC models have been proposed in the literature to offer greater flexibility (e.g. Kwan et al. (2009), Cappiello et al. (2006), Aielli

---

[*]Corporate Treasury, Royal Bank of Canada, Email: hezhongfang2004@yahoo.com. The views in this paper are solely the author's responsibility and are not related to the company the author works in.

(2013)). See Bauwens et al. (2006) and Silvennoinen and Terasvirta (2008) for surveys of multivariate GARCH models.

In this paper, we propose a class of generalized dynamic correlation (GDC) models to unify and generalize the DCC and VC models. A two-layer ARMA structure is applied to model the dynamics of the correlation matrix. At the bottom layer, an ARMA model for an auxiliary matrix variable is used with the MA component being the outer product of the vector of standardized residuals (demeaned and devolatilized data), while at the top layer, an ARMA model for the correlation matrix is applied where the MA component is the normalized auxiliary matrix variable from the bottom layer. This two-layer ARMA structure is flexible and nests the DCC and VC models as special cases. Relative to the DCC model, the GDC model applies a more flexible ARMA structure to the normalized auxiliary matrix variable from the bottom layer to form the correlation matrix, while the DCC model amounts to directly setting the correlation matrix equal to the normalized auxiliary matrix variable. Compared with the VC model, the auxiliary matrix variable to form the ARMA model for the correlation matrix is allowed to have a separate ARMA structure in the GDC model, rather than being modeled by a simple moving average with a pre-determined window as in the VC model. Allowing for greater flexibility in the structure of the correlation dynamics as in the GDC model could potentially offer empirical benefits and be used by researchers to examine if parameter restrictions of particular forms are suitable for the data series under study. The performance of the proposed GDC model is analyzed in a simulation exercise and an application to the U.S. stock indices.

The remainder of the paper is organized as follows. Section 2 provides the details of the GDC model. Estimation of the GDC model by the Markov chain Monte Carlo (MCMC) method is discussed in Section 3. The simulation exercise is presented in Section 4. The application to the U.S. stock indices is provided in Section 5. Section 6 concludes. Additional details are provided in the appendices.

2 The Model

Assume the n-by-1 vector of variables y_t follows a multivariate normal distribution $y_t \sim N(\mu_t, \Sigma_t)$, $t = 1, 2, ..., T$. In cases where a fixed number or a Dirichlet process mixture of multivariate normal distributions is used to relax the distributional assumption, the proposed GDC model could be applied to the individual components in the mixture.

The mean μ_t could be modeled by an ARMA process or n independent Gaussian processes nonparametrically. In this section, we assume that the mean μ_t is known and focus on modeling the covariance matrix Σ_t. Each element in the covariance matrix Σ_t is decomposed as $\sigma_{ij,t} = \rho_{ij,t}\sqrt{\sigma_{ii,t}\sigma_{jj,t}}$, where $\rho_{ij,t}$ is the correlation coefficient, $i, j = 1, 2, ..., n$. The correlation matrix is denoted as R_t with its (i, j) element being $\rho_{ij,t}$.

The variances $\sigma_{ii,t}$, $i = 1, 2, ..., n$, in Σ_t could be modeled either parametrically or nonparametrically (e.g. Lan et al. (2017)). In this paper, we model the variances as independent GARCH processes:

$$\sigma_{ii,t} = (1 - \alpha_{V,i} - \beta_{V,i})\sigma_{ii} + \alpha_{V,i}(y_{i,t-1} - \mu_{i,t-1})^2 + \beta_{V,i}\sigma_{ii,t-1} \tag{1}$$

where the parameters $\alpha_{V,i}$ and $\beta_{V,i}$ are both non-negative scalars and satisfy $\alpha_{V,i} + \beta_{V,i} \leq 1$. The parameter $\sigma_{i,i}$ is the unconditional mean of $(y_{i,t-1} - \mu_{i,t-1})^2$ for $i = 1, 2, ..., n$.

The correlation matrix R_t is modeled through a two-layer ARMA structure. At the bottom layer, an auxiliary n-by-n matrix variable X_t is introduced and follows the equation:

$$X_t = S_X^* + A_X \odot (u_{t-1} u_{t-1}') + B_X \odot X_{t-1} \tag{2}$$

where the variable u_t is a n-by-1 vector of the standardized residuals with the i-th element being $(y_{i,t} - \mu_{i,t})/\sqrt{\sigma_{ii,t}}$. The parameters S_X^*, A_X and B_X are all n-by-n positive semi-definite matrices. The operator $\odot$ denotes the element-by-element multiplication or the Hadamard product. The Schur product theorem ensures that the Hadamard product of two positive semi-definite matrices remains a positive semi-definite matrix. It is straightforward to see that the resulting auxiliary variable X_t is positive semi-definite given a starting value of a positive semi-definite matrix X_1.

Let X_t^* be the normalized X_t with the (i,j) element being:

$$X_{ij,t}^* = \frac{X_{ij,t}}{\sqrt{X_{ii,t} X_{jj,t}}}$$

The top layer of the GDC model follows the equation:

$$R_t = (ll' - A_R - B_R) \odot S_R + A_R \odot X_t^* + B_R \odot R_{t-1} \tag{3}$$

where l is a n-by-1 vector of ones. The parameters A_R and B_R are both n-by-n positive semi-definite matrices and satisfy that $(ll' - A_R - B_R)$ is positive semi-definite. The parameter S_R is a n-by-n matrix and is the **normalized** unconditional mean of the outer product of the standardized residual vector u_t. It follows that the resulting correlation matrix R_t is positive semi-definite with unitary diagonal elements given a starting value of a positive semi-definite matrix R_1 with unitary diagonal elements. If either A_R, B_R or $(ll' - A_R - B_R)$ is positive definite, the correlation matrix R_t is positive definite as well. See Ding and Engle (2001) for detailed discussions of the properties of matrix ARMA structures as Equation (3).

In practice, estimating the matrix parameters S_X, A_X, B_X, A_R and B_R is likely to be overwhelming for even modest size correlation matrices. If there is a prior partition of the elements in u_t into groups, individual parameters of these matrices within the same group could be restricted to be identical to reduce the number of free parameters. Without such partition, the scalar parameter version of the GDC model is:

$$\begin{aligned} X_t &= \omega_X S_X + \alpha_X (u_{t-1} u_{t-1}') + \beta_X X_{t-1} \\ X_t^* &= normalize(X_t) \\ R_t &= (1 - \alpha_R - \beta_R) S_R + \alpha_R X_t^* + \beta_R R_{t-1} \end{aligned} \tag{4}$$

where the parameters ω_X, α_X, β_X, α_R and β_R are non-negative scalars and satisfy $\alpha_R + \beta_R \leq 1$. The parameter S_X is a n-by-n matrix and is the unconditional mean of the outer

product of the standardized residual vector u_t. Further parsimony can be achieved via the unconditional-mean targeting for the auxiliary variable X_t:

$$\omega_X = 1 - \alpha_X - \beta_X$$
$$\alpha_X + \beta_X \leq 1 \tag{5}$$

The scalar parameter version GDC model (Equation (4)) combined with the unconditional-mean targeting (Equation (5)) is termed the GDC model hereafter for expositional convenience. An integrated version of the GDC model (IGDC hereafter) further imposes the constraint $\alpha_R + \beta_R = 1$ and reduces one more free parameter than the GDC model.

If one imposes the parameter restrictions $\alpha_R = 1$ and $\beta_R = 0$ on Equation (4), the GDC model is reduced to the DCC model of Engle (2002). On the other hand, the VC model of Tse and Tsui (2002) is equivalent to imposing the following parameter restrictions on the equation of the auxiliary variable X_t in Equation (4):

$$\omega_X = 0$$
$$\alpha_X = \frac{1}{M} 1_{\{lag \leq M\}}$$
$$\beta_X = 1 \tag{6}$$

where the parameter M is a pre-determined positive integer and the indicator $1_{\{lag \leq M\}} = 1$ if u_{t-j} satisfies $j \leq M$ and 0 otherwise. Effectively the parameter restrictions imposed by the VC model reduces the auxiliary variable X_t to be a M-period simple moving average of the outer product $u_t u_t'$.

The different model structures of the GDC, DCC and VC models imply interesting differences in how information in the standardized residuals is incorporated to model the dynamics of correlations. Table 1 shows a comparison of the GDC, DCC and VC models by representing the ARMA structures in MA form. At the bottom layer of forming an auxiliary matrix variable from historical standardized residuals, the VC model uses a simple moving average with a pre-determined window while the GDC and DCC models apply an ARMA structure and hence enable more flexible use of all historical standardized residuals. On the other hand, at the top layer for the correlation matrix, the DCC model simply sets the correlation matrix equal the contemporaneous normalized auxiliary matrix variable from the bottom layer, while the GDC and VC models allows the correlation matrix to take an ARMA structure and hence utilize both contemporaneous and past normalized auxiliary matrix variables.

In terms of the number of free parameters for modeling the correlations, the DCC model uses 2 parameters α_X and β_X, while the VC model uses 3 parameters α_R, β_R and the window length parameter M, which is pre-determined by the researchers and is often set to equal the dimension of the data vector y_t. The GDC model uses 4 parameters α_X, β_X, α_R and β_R to model the correlation matrix and hence is at slight disadvantage in terms of model parsimony as the cost of greater flexibility. The IGDC model achieves greater parsimony by using 3 parameters α_X, β_X and α_R at the expense of being less flexible than the GDC model.

By recasting the DCC and VC models as restricted versions of the GDC framework, one empirical strategy could be using the GDC model as the default and perform additional

statistical tests, either by Bayesian or frequentist approaches, to assess the appropriateness of the parameter restrictions the DCC or the VC model imposes in empirical applications on a case by case basis. Given the empirical success of the DCC and VC models, it is possible that the parameter restrictions they impose fit important empirical features of economic and financial time series data and, in such cases, these restricted versions could potentially outperform the unrestricted GDC model through greater model parsimony.

Table 1: Comparing the Structures of Correlation Models

	Dynamics of Auxiliary Variable X_t	Dynamics of Correlation Matrix R_t (X_t^* is normalized X_t)
GDC	$A_t + \alpha_X u_{t-1}u'_{t-1} + ... + \alpha_X \beta_X^{t-2} u_1 u'_1$	$B_t + \alpha_R X_t^* + ... + \alpha_R \beta_R^{t-2} X_2^*$
IGDC	$A_t + \alpha_X u_{t-1}u'_{t-1} + ... + \alpha_X \beta_X^{t-2} u_1 u'_1$	$C_t + \alpha_R X_t^* + ... + \alpha_R (1-\alpha_R)^{t-2} X_2^*$
DCC	$A_t + \alpha_X u_{t-1}u'_{t-1} + ... + \alpha_X \beta_X^{t-2} u_1 u'_1$	X_t^*
VC	$\frac{1}{M}(u_{t-1}u'_{t-1} + ... + u_{t-M}u'_{t-M})$	$B_t + \alpha_R X_t^* + ... + \alpha_R \beta_R^{t-2} X_2^*$

Note: The intercepts A_t, B_t and C_t are deterministic functions of the time index t and starting values X_1 and R_1: $A_t = \left(\frac{(1-\alpha_X-\beta_X)(1-\beta_X^{t-1})}{1-\beta_X} S_X + \beta_X^{t-1} X_1 \right)$, $B_t = \left(\frac{(1-\alpha_R-\beta_R)(1-\beta_R^{t-1})}{1-\beta_R} S_R + \beta_R^{t-1} R_1 \right)$ and C_t equals B_t by replacing β_R with $1 - \alpha_R$.

3 Estimation, Forecast and Model Comparison

Estimating the GDC and IGDC models does not require any specialized method beyond the conventional ones for estimating the DCC and VC models. The maximum likelihood approaches in Engle (2002) and Tse and Tsui (2002) could be directly applied to estimate the GDC models. For the ease of finite-sample inference, we use the Bayesian Markov chain Monte Carlo (MCMC) method to estimate the GDC models in this paper. A standard Metropolis-within-Gibbs sampler is developed and is sketched below. The full details of the algorithm can be found in Appendix A.

For convenience of developing the proposal distributions in the Metropolis-Hasings (MH) steps, we reparametize the GDC model by $\gamma_X = \alpha_X + \beta_X$, $\lambda_X = \frac{\alpha_X}{\gamma_X}$, $\gamma_R = \alpha_R + \beta_R$, $\lambda_R = \frac{\alpha_R}{\gamma_R}$ and specify independent Beta prior distributions for γ_X, λ_X, γ_R and λ_R. Let θ_C denote the vector collecting parameters for the model of the correlation matrix and $p(\theta_C)$ denote their joint prior distribution.

Similarly for the GARCH models of the variances, the reparametization is $\gamma_{V,i} = \alpha_{V,i} + \beta_{V,i}$ and $\lambda_{V,i} = \frac{\alpha_{V,i}}{\gamma_{V,i}}$ for $i = 1, 2, ..., n$ with independent Beta prior distributions. Let θ_V denote the vector collecting parameters for the model of the variances and $p(\theta_V)$ denote their joint prior distribution.

Denote the estimation sample as Y_T containing $y_1, y_2, ..., y_T$. The posterior distribution of the model parameters is $p(\theta_C, \theta_V | Y_T)$. The Gibbs sampler iterates sampling over the two conditional posterior distributions $p(\theta_C | \theta_V, Y_T)$ and $p(\theta_V | \theta_C, Y_T)$, which satisfy:

$$p(\theta_C | \theta_V, Y_T) \propto p(\theta_C | \theta_V) p(Y_T | \theta_C, \theta_V)$$

and

$$p(\theta_V|\theta_C, Y_T) \propto p(\theta_V|\theta_C)p(Y_T|\theta_C, \theta_V)$$

Given the independent prior distributions, we further have $p(\theta_C|\theta_V) = p(\theta_C)$ and $p(\theta_V|\theta_C) = p(\theta_V)$.

To sample from the conditional posterior distributions $p(\theta_C|\theta_V, Y_T)$ and $p(\theta_V|\theta_C, Y_T)$, a MH step is applied for each individual parameter within θ_C and θ_V with a generic random-walk proposal of a normal distribution truncated between 0 and 1. The details can be found in Appendix A.

Given a large number K of draws $\{\theta_C^{(i)}, \theta_V^{(i)}\}_{i=1}^K$ from the posterior distribution $p(\theta_C, \theta_V|Y_T)$, the predictive distribution for the next-period observation y_{T+1} is:

$$
\begin{aligned}
p(y_{T+1}|Y_T) &= \int p(y_{T+1}|\theta_C, \theta_V, Y_T)p(\theta_C, \theta_V|Y_T)d\theta_C d\theta_V \\
&\approx \frac{1}{K}\sum_{i=1}^K p(y_{T+1}|\theta_C^{(i)}, \theta_V^{(i)}, Y_T) \quad (7)
\end{aligned}
$$

where the density function $p(y_{T+1}|\theta_C, \theta_V, Y_T)$ is $N(\mu_{T+1}, \Sigma_{T+1})$ with the covariance matrix Σ_{T+1} calculated by iterating forward the GARCH equations for the variances (Equation (1)) and the GDC equations for the correlation matrix (Equations (4) and (5)).

To compare the performance of two correlation models M_1 and M_2, one could use the predictive likelihoods over the test sample y_{T+1}, y_{T+2}, ..., y_{T+H} under the two models $p(y_{T+1}, y_{T+2}, ..., y_{T+H}|Y_T, M_1)$ and $p(y_{T+1}, y_{T+2}, ..., y_{T+H}|Y_T, M_2)$. It is straightforward to apply the decomposition $p(y_{T+1}, y_{T+2}, ..., y_{T+H}|Y_T) = \prod_{h=1}^H p(y_{T+h}|Y_T, y_{T+1}, ..., y_{T+h-1})$ where each component $p(y_{T+h}|Y_T, y_{T+1}, ..., y_{T+h-1})$ can be calculated by Equation (7) with expanded training samples. The difference between the log of the predictive likelihoods of M_1 and M_2 is the log Bayes factor and is the tool used in this paper for model comparison. A positive log Bayes factor suggests that the data favors the model M_1. The larger the log Bayes factor, the stronger evidence there is to favor the model M_1 over the alternative model M_2.

4   Simulation Study

A simulation exercise is conducted to study the performance of the GDC models. The data generating process (DGP) is a tri-variate normal distribution $y_t \sim N(0, \Sigma_t)$ for $t = 1, 2, ..., 200$. The variances are generated by the GARCH process $\sigma_{ii,t} = 0.01 + 0.2y_{i,t-1}^2 + 0.7\sigma_{ii,t-1}$ with $\sigma_{ii,1} = 0.1$ for $i = 1, 2, 3$. The correlation matrix is $R_t = 1_{\{t \le 70\}}R_1 + 1_{\{70 < t \le 135\}}R_2 + 1_{\{t > 135\}}R_3$ with:

$$
R_1 = \begin{pmatrix} 1 & -0.7 & 0.7 \\ -0.7 & 1 & 0 \\ 0.7 & 0 & 1 \end{pmatrix}
$$

$$R_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0.7 \\ 0 & 0.7 & 1 \end{pmatrix}$$

$$R_3 = \begin{pmatrix} 1 & 0.7 & -0.7 \\ 0.7 & 1 & -0.7 \\ -0.7 & -0.7 & 1 \end{pmatrix}$$

The (1,2) and (1,3) elements of the correlation matrix R_t are step-wise increasing and decreasing functions of t respectively, while the (2,3) element of R_t is a step-wise concave function of t. This setup of the correlation matrix is to test if the GDC models are able to estimate the various shapes of the time-varying correlations based on modest size data samples. We generate one set of data $\{y_t\}_{t=1}^{200}$ from the DGP as the training sample for the correlation models.

The mean μ_t in the models is assumed to be a constant and has the prior $N(0, 100)$. The priors for the parameters γ_X, γ_R, $\gamma_{V,1}$, $\gamma_{V,2}$ and $\gamma_{V,3}$ are specified as Beta$(9, 1)$ and favor values close to one. For the parameters λ_X, $\lambda_{V,1}$, $\lambda_{V,2}$ and $\lambda_{V,3}$, their priors are Beta$(1, 4)$ and favor values close to zero. We experiment plotting the likelihoods of the GDC and IGDC models as functions of the parameter λ_R, conditional on plausible values of the other model parameters, and find that the maximum often occurs at points towards one. Therefore for the GDC and IGDC models, the prior of the parameter λ_R is Beta$(9, 1)$ and favors values close to one. For the VC model, the prior for λ_R is Beta$(1, 4)$. The difference in the prior λ_R between the GDC models and the VC model reflects the fact that the VC model only captures the autoregressive dynamics at the top layer and hence its coefficient β_R on the AR term R_{t-1} tends to be close to one while its coefficient α_R on the MA term X_t^* is smaller. On the contrary, the GDC models already apply an ARMA structure at the bottom layer and hence the AR coefficient at the top layer would be much lower than the case in the VC model.

The resulting priors for the parameters α_X, $\alpha_{V,1}$, $\alpha_{V,2}$ and $\alpha_{V,3}$ have the 10[th], 50[th] and 90[th] percentiles of about 0.02, 0.14 and 0.39, while the resulting priors for the parameters β_X, $\beta_{V,1}$, $\beta_{V,2}$ and $\beta_{V,3}$ have the 10[th], 50[th] and 90[th] percentiles of about 0.49, 0.75 and 0.91. For the parameter pair (α_R, β_R), the 10[th], 50[th] and 90[th] percentiles are (0.65, 0.01), (0.83, 0.07) and (0.94, 0.20) respectively in the GDC model and are (0.02, 0.49), (0.14, 0.75) and (0.39, 0.91) respectively in the VC model. These priors are deemed reasonably dispersed over empirically plausible ranges.

The proposal distribution for the MH steps is a generic random-walk normal distribution truncated between 0 and 1. The proposal distributions are tuned to result in acceptance rates between 17% and 43%. The draws of all correlation models appear to stabilize quickly after a couple of initial draws. We keep a total of 1,000 posterior draws after thinning every 10-th draw for analysis after a burn-in of 200. The Geweke convergence test (Geweke (1992)) is performed to check the convergence of the posterior draws. The details of the Geweke test are provided in Appendix B. The resulting p-values of the Geweke convergence tests are all greater than 5% and suggest no rejection of the null hypothesis of convergence.
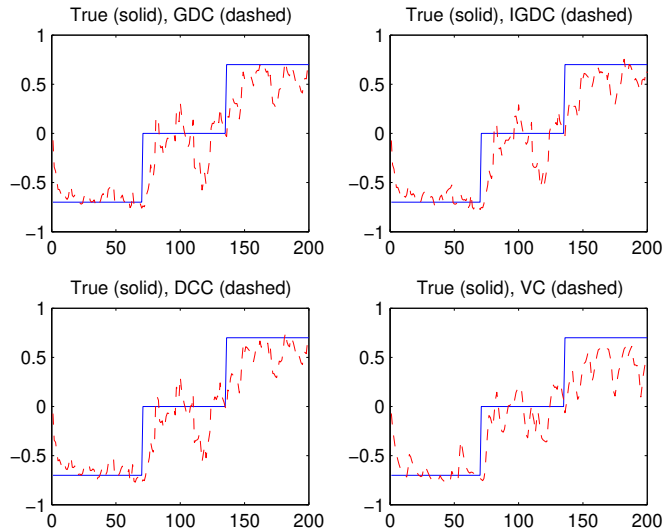
Figure 1: The (1,2) Element of the Correlation Matrix: Simulation Study

The true values of the mean and GARCH parameters are all within their respective estimated 95% credible intervals. The detailed parameter estimates are available upon request. We focus on the estimates of the correlations. Figures 1, 2 and 3 plot the posterior means of the (1,2), (1,3) and (2,3) elements of the correlation matrix estimated from the GDC, IGDC, DCC and VC models along with their true values. The correlation estimates from the four models are very close. The VC model appears to produce slightly more volatile estimates of correlations than the other three models.

Table 2 shows the root mean squared errors (RMSE) of the three correlation estimates from the four correlation models respectively. The GDC model generally has the smallest RMSE for all the three correlation estimates, while the VC model tends to have the largest RMSE. Nevertheless the differences in the RMSE measures are minimal across the four models. We conclude that the four correlation models produce comparable performance in this simulation exercise. Despite that all the four models are misspecified in this simulation exercise, their correlation estimates are able to track the shapes of the true correlations based on a relatively short sample of data, which is an encouraging indication of their usefulness in empirical applications.

5   Empirical Application

As an empirical illustration, we apply the GDC models to daily log returns of the NASDAQ composite, S&P500 and Dow Jones Industrial Average indices. All data are from the FRED database of the federal reserve bank of St. Louis. The data sample is from January 3, 2017 to February 14, 2018 with a total of 282 daily observations. In estimation, the log returns are scaled up by 100. Figure 4 plots the daily return data. Frequent simultaneous declines
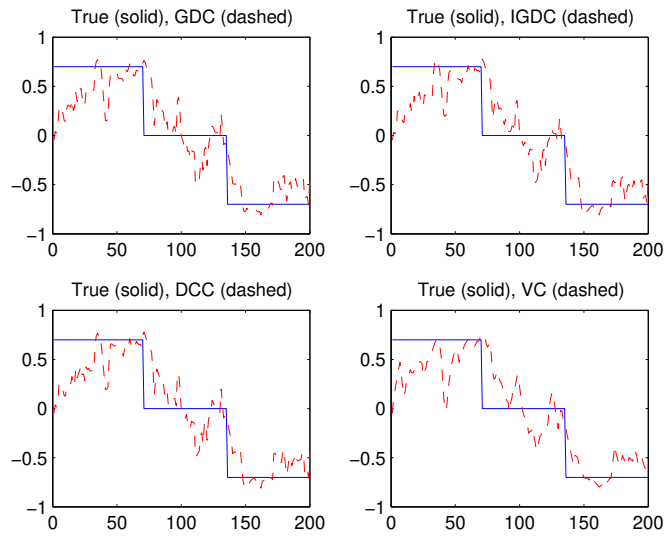
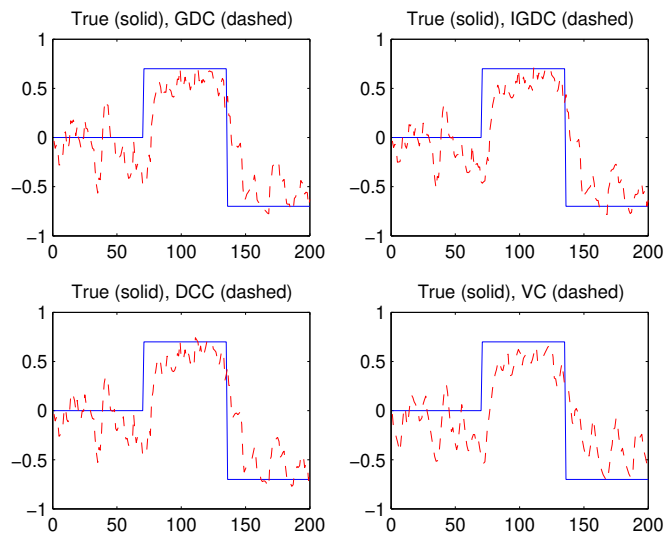Figure 2: The (1,3) Element of the Correlation Matrix: Simulation Study

Figure 3: The (2,3) Element of the Correlation Matrix: Simulation Study

Table 2: Comparing the Correlation Estimates of Correlation Models

	The (1,2) Element of Correlation Matrix	The (1,3) Element of Correlation Matrix	The (2,3) Element of Correlation Matrix
GDC:	0.2398	0.2752	0.3285
IGDC:	0.2408	0.2819	0.3284
DCC:	0.2413	0.2829	0.3287
VC:	0.2662	0.2759	0.3757

Note: This table compares the root mean squared errors of the posterior means of the correlation estimates from the GDC, IGDC, DCC and VC models relative to the true values of the correlations.
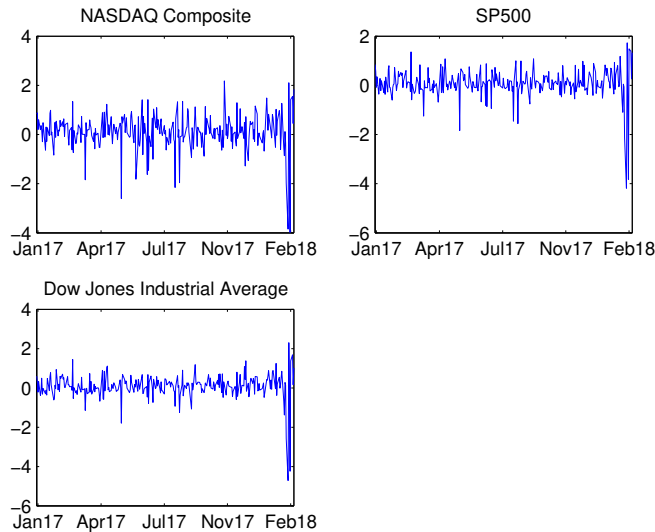

Figure 4: Daily Log Returns of NASDAQ Composite, S&P500 and Dow Jones Industrial Average Indices (in Percent)

and spikes of the three daily return series can be observed in the data. In particular, all three daily return series experienced 5- to 8-standard-deviation declines on February 5 and 8 of 2018. We estimate both the GDC and IGDC models for the stock index data, along with the DCC and VC models for comparison. We model the mean process of the models by a constant. The priors for the parameters are the same as those in the simulation exercise of Section 4. We consider that these priors cover the empirically plausible ranges of the parameters' values while being reasonably dispersed. In the estimation, the truncated-normal random-walk proposals are tuned to result in acceptance rates between 15% and 43% for all parameters. After a burn-in of 500 draws, we keep 20,000 subsequent draws for analysis. The Geweke convergence tests of the posterior draws return p-values greater than 5% and hence indicate no rejection of the null hypothesis of convergence.

Estimates of the means of the models are basically the sample average daily returns and are not reported here. The estimates of the GARCH processes are very close in the three

models. The posterior means of $\alpha_{V,i}$ and $\beta_{V,i}$, $i = 1, 2, 3$, are about 0.2 and 0.7 respectively with 90% credible intervals of about (0.01, 0.50) and (0.40, 0.95). Details of the estimates of the means and GARCH parameters in the models are available upon request.

We focus on the estimates of the correlation models. Table 3 provides the posterior means of the correlation model parameters along with their 90% credible intervals. Figures 5 and 6 compare the posterior distributions of the parameters in the four correlation models. For the IGDC model, posterior draws of β_R is obtained by calculating $1 - \alpha_R$. There is considerable uncertainty in the parameter values reflected by their relatively wide credible intervals. The posterior distributions of the parameters β_X and α_R in the IGDC model appear to have bimodal shapes. For the parameter β_X in the IGDC model, the two modes appear to be one slightly below 0.5 and the other one slightly above 0.5, while the two modes for α_R in th IGDC model appears to be one slight above 0.5 and the other one close to one. Between the GDC and IGDC models, the IGDC model estimates tend to have a smaller MA coefficient α_X at the bottom layer and a larger MA coefficient α_R and AR coefficient β_R at the top layer than the GDC model estimates. The estimated persistence level $\alpha_R + \beta_R$ of the top layer in the GDC model is well below one.

Table 3: Parameter Estimates of Correlation Models

	Posterior Mean	5-th Percentile	95-th Percentile
GDC:			
α_R	0.52	0.23	0.85
β_R	0.18	0.01	0.46
α_X	0.26	0.02	0.64
β_X	0.48	0.11	0.82
IGDC:			
α_R	0.67	0.34	0.97
α_X	0.16	0.01	0.47
β_X	0.47	0.11	0.83
DCC:			
α_X	0.08	0.01	0.20
β_X	0.54	0.15	0.86
VC:			
α_R	0.06	0.01	0.14
β_R	0.55	0.23	0.87

Note: This table provides the posterior means and the 90% credible intervals of the parameters in the correlation component of the GDC, IGDC, DCC and VC models. The data is the daily log returns of the NASDAQ composite, S&P500 and Dow Jones Industrial Average indices from January 3, 2017 to February 14, 2018 with a total of 282 daily observations. The data are scaled by 100. For the VC model, the length of the moving average window is 3.

The DCC model estimates have a smaller MA coefficient α_X than the IGDC model estimates. But the estimated AR coefficient β_X is higher in the DCC model than the IGDC model. The estimated persistence levels $\alpha_X + \beta_X$ of the bottom layer in the two models are very close and are both below that of the GDC model.
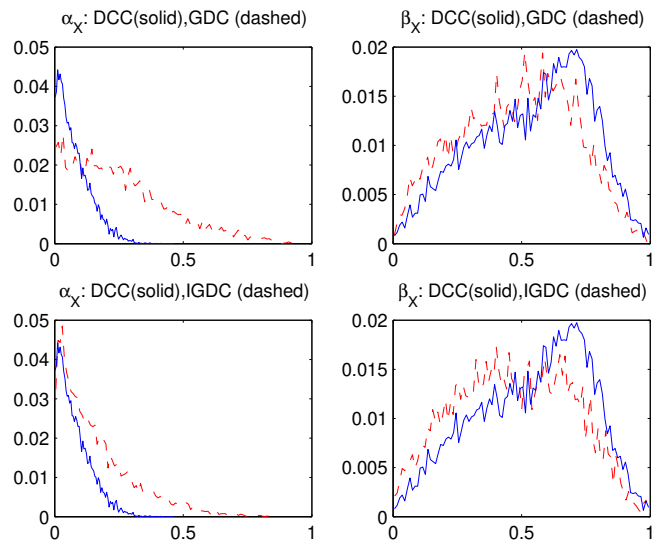
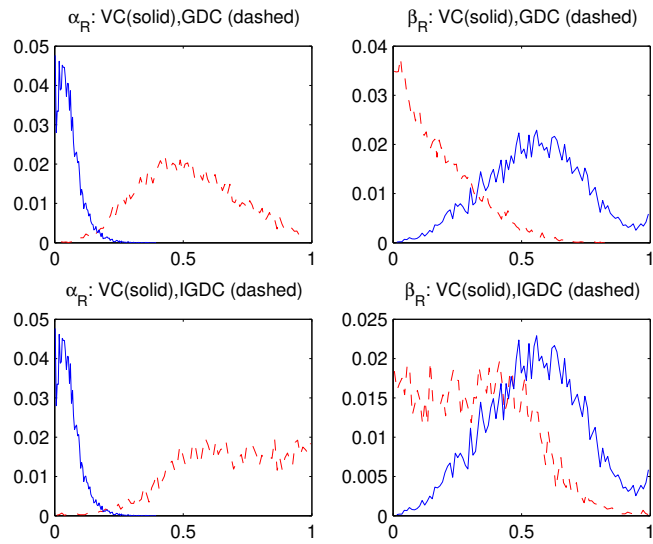Figure 5: Posterior Distributions of Correlation Model Parameters: Bottom Layer

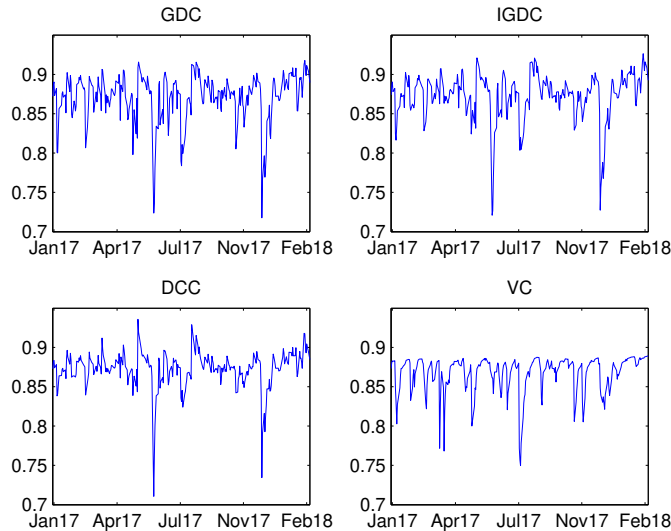Figure 6: Posterior Distributions of Correlation Model Parameters: Top Layer

Figure 7: Estimated Correlation Between the Daily Returns of NASDAQ Composite and S&P500 Indices

Comparing with the VC model, the estimated persistence level $\alpha_R + \beta_R$ of the top layer in the GDC model is higher than the VC model. The GDC model has a higher MA coefficient α_R than its AR coefficient β_R while the opposite is true in the VC model estimates. A possible reason, as discussed in Section 4, is that the GDC model already applies an ARMA structure at the bottom layer and hence the persistence level at the top layer could be lower than the VC model as the bottom layer of the GDC model helps absorb some persistence in the correlations.

Figures 7, 8 and 9 plot the posterior means of the correlations between the pairs (NASDAQ, S&P500), (NASDAQ, DowJones) and (S&P500, DowJones) estimated from the GDC, IGDC, DCC and VC models. The correlation estimates form the GDC, IGDC and DCC models are very close with inter-model correlations above 0.95, while the correlation estimates from the VC model are noticeably different than the other three models. In particular, for the date of February 5, 2018 when all three return series had extreme declines, the GDC, IGDC and DCC models generate noticeably higher correlation estimates than previous days while the correlation estimates from the VC model appear less responsive. Among the GDC, IGDC and DCC models, the correlation estimates from the DCC model tend to have more extreme spikes than the two GDC models, while the two GDC models tend to produce sharper declines in correlation estimates than the DCC model. Nevertheless the difference in the correlation estimates between the GDC models and the DCC model is small.
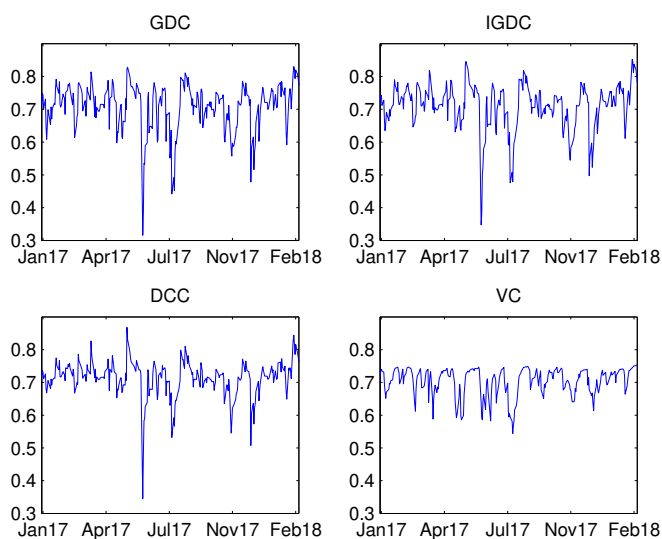
13

Figure 8: Estimated Correlation Between the Daily Returns of NASDAQ Composite and Dow Jones Industrial Average Indices
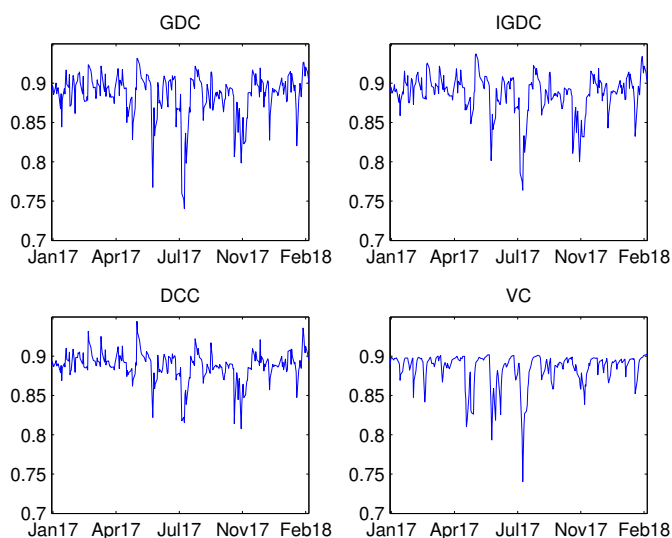

Figure 9: Estimated Correlation Between the Daily Returns of S&P500 and Dow Jones Industrial Average Indices

## 5.1 Predictive Likelihoods and Model Comparison

The predictive likelihoods discussed in Section 3 are used to compare the performance of the GDC, IGDC, DCC and VC models. The test sample is from January 2 to February 14, 2018 with a total of 31 observations. A recursive out-of-sample forecast exercise is performed over the test sample. We first use the sample from January 3, 2017 to December 29, 2017 to estimate the four models and calculate the one-day-ahead predictive likelihood (Equation (7)) for the first observation in the test sample, i.e. January 2, 2018. Next the training sample is expanded to include the data of January 2, 2018 and all the four models are re-estimated based on the expanded training sample. The one-day-ahead predictive likelihood for the second observation in the test sample (January 3, 2018) is calculated. This procedure is iterated forward and produces a 31-by-1 vector of the one-day-ahead predictive likelihoods over all the observations in the test sample.

The cumulative log Bayes factors for the GDP and IGDC models against the DCC and VC models are calculated based on the one-day-ahead predictive likelihoods as discussed in Section 3. One unique pattern of the test sample is the unusually sharp declines (5 to 8 standard deviations away from historical average) of all three stock index returns on February 5 and 8 of 2018 and, to a less extent (3 to 4 standard deviations away from historical average), on February 2 of 2018. The one-day-ahead predictive likelihoods from all the four correlation models deteriorate significantly on these three days, reflecting the inadequacy of the multivariate normality assumption for stock returns on extremely volatile periods. To isolate the impact of these three days of extreme stock returns on model comparison, we discuss the patterns of the cumulative log Bayes factors with and without the three unusual days.

Figure 10 plots the cumulative log Bayes factors along with the daily log returns in the full test sample. The sharp declines of the stock returns on February 2, 5 and 8 of 2018 are evident in the plot. For the pair of the GDC and IGDC models, the GDC model appears gradually accumulating gains in Bayes factor over the IGDC model until February 1 of 2018 right before the sharp decline of the stock returns on February 2 of 2018 and deteriorates thereafter. One possible reason could be that the GDC model puts a positive weight on the constant matrix S_R at the top layer and hence might dampen the impact of the extreme returns on correlations relative to the IGDC model.

The advantage of the GDC and IGDC models relative to the VC model is evident in the right bottom panel of Figure 10. Both the GDC and IGDC models steadily accumulate gains in Bayes factors over the VC model. Relative to the DCC model, the IGDC model shows largely equal predictive performance until after Feb 8 of 2018, the last of the three days of extreme stock returns, and begins to accumulate gains in Bayes factor thereafter. The performance of the GDC model relative to the DCC model is more mixed. Similar to its pattern against the IGDC model, the GDC model appears gradually accumulating gains in Bayes factor over the DCC model until February 1 of 2018 and deteriorates during the period of extreme stock returns from February 2 to 8 of 2018. After these days of extreme stock returns, the relative performance of the GDC model against the DCC model appears stabilized and resumes its upward trend in Bayes factor. As discussed in comparing the GDC and IGDC models, the GDC model might adapt less quickly to extreme returns than
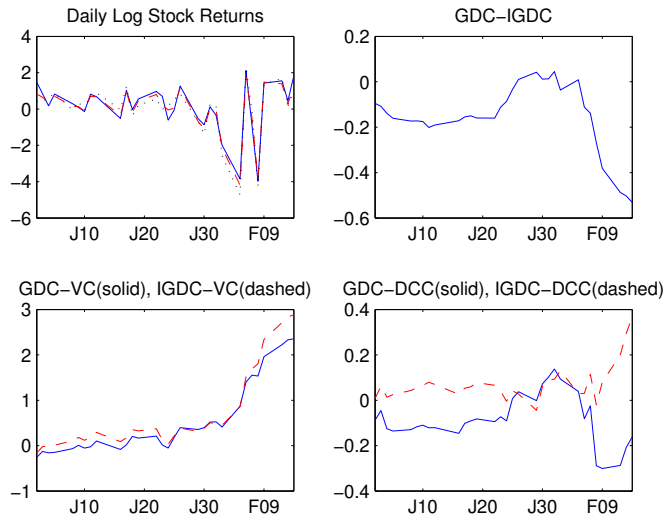
Note: The test sample is from January 2 to February 14, 2018. The left top panel plots the daily log returns of the NASDAQ composite(solid), S&P500(dashed) and Down Jones Industrial Average(dotted) indices in the test sample. The right top panel plots the cumulative log Bayes factor of the GDC model against the IGDC model. The left bottom panel plots the cumulative log Bayes factors of the GDC(solid) and IGDC(dashed) models against the VC model. The right bottom panel plots the cumulative log Bayes factors of the GDC(solid) and IGDC(dashed) models against the DCC model. The x axises of the subplots are the dates in the format MDD.

Figure 10: Cumulative Log Bayes Factor: Stock Indices
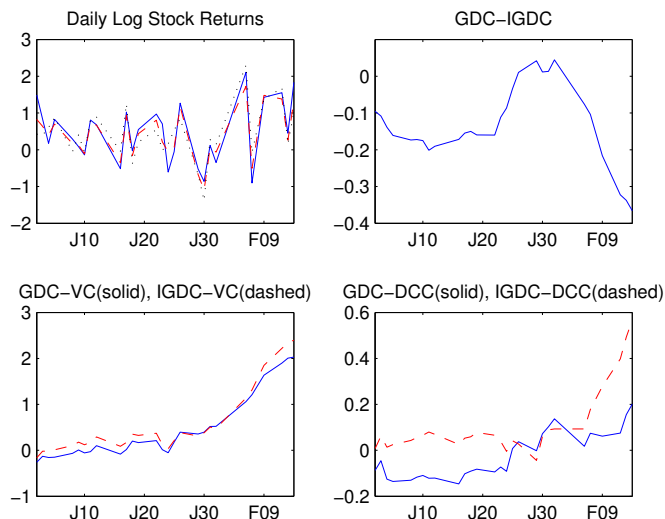
Note: The left top panel plots the daily log returns of the NASDAQ composite(solid), S&P500(dashed) and Down Jones Industrial Average(dotted) indices in the test sample from January 2 to February 14, 2018, excluding the three days February 2, 5 and 8 of 2018. The right top panel plots the cumulative log Bayes factor of the GDC model against the IGDC model. The left bottom panel plots the cumulative log Bayes factors of the GDC(solid) and IGDC(dashed) models against the VC model. The right bottom panel plots the cumulative log Bayes factors of the GDC(solid) and IGDC(dashed) models against the DCC model. The x axises of the subplots are the dates in the format MDD.

Figure 11: Cumulative Log Bayes Factor: Stock Indices (Test Sample Excluding February 2, 5 and 8 of 2018)

the DCC model due to its positive weight on the constant matrix S_R at the top layer.

Figure 11 shows the daily log stock returns and cumulative log Bayes factors of the GDC models in the test sample excluding the three days of extreme stock returns (February 2, 5 and 8 of 2018). That is, the cumulative log Bayes factors are calculated based on only the one-day-ahead predictive likelihoods in the days of the test sample excluding the three days of extreme stock returns. The purpose is to isolate the impact of the extreme stock returns and to compare the models' predictive performances in more "normal" time periods. The patterns of the relative performance between the pairs of GDC vs. IGDC, GDC/IGDC vs. VC and IGDC vs. DCC remain largely the same as in Figure 10. For the pair of GDC vs. DCC, the log Bayes factor shows a generally upward trend favoring the GDC model throughout the test days.

6   Conclusion

In this paper, we propose a class of GDC models that applies a two-layer ARMA structure to model the dynamics of a correlation matrix. The popular DCC and VC models can be recast as special cases of the proposed GDP model. The greater flexibility of the GDC model is obtained in a parsimonious way. In the scalar version of the GDC model with

17

unconditional mean targeting, there are two more parameters than the DCC model, while in the integrated version of the GDC model, there is only one more parameter than the DCC model. Therefore the GDC model could be implemented in empirical applications without significantly increasing the computation burden and model complexity, while allowing for greater flexibility than the existing DCC and VC models. Performance of the GDC models is illustrated by a simulation exercise and an empirical application to the U.S. stock indices.

Appendix A: Details of the MCMC Algorithm

Let ϕ be a generic scalar variable representing an arbitrary parameter in θ_C (γ_R, λ_R, γ_X, λ_X) or θ_V ($\gamma_{V,i}$, $\lambda_{V,i}$, $i = 1, 2, ..., n$). The MH step for ϕ is as follows. Given its i-th draw $\phi^{(i)}$, a proposal ϕ^* is simulated from the truncated normal distribution $N(\phi^{(i)}, \tau)1_{\{0<\phi^*<1\}}$. The $(i+1)$-th draw $\phi^{(i+1)}$ equals ϕ^* with the probability $\min(1, x)$ and $\phi^{(i)}$ otherwise, where x is $\frac{p(\phi^*|Y,\theta_{-\phi})TN(\phi^{(i)}|\phi^*,\tau)}{p(\phi^{(i)}|Y,\theta_{-\phi})TN(\phi^*|\phi^{(i)},\tau)}$, $TN(\cdot|z,\tau)$ is the density function of a normal distribution with mean z and variance τ truncated between 0 and 1 and $p(\phi|Y,\theta_{-\phi})$ is the density function of the parameter ϕ conditional on the data and all other model parameters $\theta_{-\phi}$. The posterior density $p(\phi|Y,\theta_{-\phi})$ is calculated by the Bayes rule $p(\phi|Y,\theta_{-\phi}) \propto p(\phi)p(Y|\phi,\theta_{-\phi})$. The prior $p(\phi)$ is discussed in Section 3. To calculate the likelihood $p(Y|\phi,\theta_{-\phi})$, one could iterate the GARCH equations (Equation 1) and the GDC equations (Equation 4 and 5) to calculate the covariance matrix Σ_t. The mean μ_t of the data y_t is assumed to be a vector of constants and is conveniently sampled from univariate regressions of each variable in y_t on a constant, conditional on draws of the covariance matrix Σ_t. The likelihood $p(Y|\phi,\theta_{-\phi})$ can be calculated as the product of multivariate normal density functions with mean μ_t and covariance matrix Σ_t. To avoid numerical overflow, log likelihoods are calculated instead. Given draws of θ_C and θ_V, one could obtain draws of α_R, β_R, α_X, β_X, $\alpha_{V,i}$, $\beta_{V,i}$, $i = 1, 2,$ ..., n by the formula $\alpha = \gamma\lambda$ and $\beta = \gamma - \alpha$. It should be noted that, for the GARCH parameters, we sample them as n univariate GARCH models, which is less efficient than utilizing the correlations between the n data series but is more convenient for use in large dimension applications. Also note that the parameters S_X is calculated as the sample average of the outer product of the standardized residual vector u_t and the parameter S_R is the normalized S_X.

Appendix B: Geweke Test of Convergence

The Geweke test is first proposed in Geweke (1992) as a practical and convenient method to check the convergence of simulation draws. See Chib (2001) for further discussions of convergence tests in general.

The full sample of the draws of a given parameter is divided into three parts: the first 10%, the second 50% and the final 40% of the draws. The intuition is that, if the draws are converged, the first subsample and the final subsample should be approximately independent. By the central limit theorem, the average of the draws in a subsample should follow a normal distribution with the variance calculated by a heteroscedasticity-and-autocorrelation-consistent (HAC) method to account for the possible serial correlations

in the draws. It is straightforward to derive that the difference of the averages in the first and the final subsamples should follow a normal distribution with zero mean and the variance being the sum of the variances in the two subsamples, which is the statistic to perform the Geweke test. If the value of the statistic lies in the extreme tails of its distribution, we conclude that the draws are not converged. Otherwise, we do not reject the hypothesis that the draws are indeed converged. In this paper, we use the Newey-West HAC method to calculate the variances of the subsamples of draws with the truncation lag being the integer part of $4(T/100)^{2/9}$ where T is the number of draws in the subsample.

References

Aielli, G. (2013). Dynamic conditional correlation: On properties and estimation. *Journal of Business and Economic Statistics*, 282–299.

Barnard, J., R. McCulloch, and X. Meng (2000). Modeling covariance matrices in terms of standard deviations and correlations with application to shrinkage. *Statistica Sinica*, 1281–1311.

Bauwens, L., S. Laurent, and J. Rombouts (2006). Multivariate garch models: A survey. *Journal of Applied Econometrics*, 79–109.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics 31*, 307–327.

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

Cappiello, L., R. Engle, and K. Sheppard (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics*, 537–572.

Chib, S. (2001). Markov chain monte carlo methods: Computation and inference. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics, Volume 5*, pp. 3569–3649. Elsevier Science B.V.

Ding, Z. and R. Engle (2001). Large scale conditional covariance matrix modeling, estimation and testing. *Academia Economic Papers*, 157–184.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics 20* (3), 339–350.

Engle, R., O. Ledoit, and M. Wolf (2017). Large dynamic covariance matrices. *Journal of Business and Economic Statistics forthcoming*.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the uk inflation. *Econometrica 50*, 987–1008.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In J. Bernardo, J. Berger, A. Dawid, and A. Smith (Eds.), *Bayesian Statistics*, pp. 169–193. Oxford University Press, New York.

Kwan, W., W. Li, and K. Ng (2009). A multivariate threshold varying conditional correlations model. *Econometric Reviews*, 20–38.

Lan, S., A. Holbrook, G. Elias, N. Fortin, H. Ombao, and B. Shahbaba (2017). Flexible bayesian dynamic modeling of correlation and covariance matrices. arXiv:1711.02869 [stat.ME].

Silvennoinen, A. and T. Terasvirta (2008). Multivariate garch models. In T. Andersen, R. Davis, J. Kreiss, and T. Mikosch (Eds.), *Handbook of Financial Time Series*. New York: Springer.

Tse, Y. K. and A. K. C. Tsui (2002). A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. *Journal of Business & Economic Statistics 20*(3), 351–362.