

Computational finance: correlation, volatility, and markets

Katherine Bennett Ensor^{1*} and Ginger M. Koev^{2,3}

Financial data by nature are inter-related and should be analyzed using multivariate methods. Many models exist for the joint analysis of multiple financial instruments. Early models often assumed some type of constant behavior between the instruments over the time period of analysis. But today, time-varying covariance models are a key component of financial time series analysis leading to a deeper understanding of changing market conditions. Models for cointegration of financial data quickly grow in their complexity and parameters, and 20 years of research offers a variety of solutions to this complexity. After a short introduction of univariate volatility models, this article begins with the basic multivariate formulation for time series covariance modeling and moves to leading time series tools that address this complexity. Coupling these models with regime switching via a Markov process extends the features that can be understood from market behavior. We ground this review in an example of modeling the covariance of securities within sectors and sectors within markets, with dynamics that allow for two different market regimes. Specifically, we simultaneously model individual daily stock data that belong to one of three market sectors and examine the behavior of the market as a whole as well as the behavior of the market sectors over time. A motivation for this characterization concerns portfolio diversification and stock anomalies, and we capture the changing comovement of stocks within and between sectors as market conditions change. For example, some of these market conditions include market crashes or collapses and common external influences. © 2014 The Authors. *WIREs Computational Statistics* published by Wiley Periodicals, Inc.

How to cite this article:

WIREs Comput Stat 2014, 6:326–340. doi: 10.1002/wics.1323

Keywords: co-volatility forecasting; dynamic conditional correlation; GARCH/MGARCH; regime switching; stock volatility

INTRODUCTION

Modeling the daily volatility (variance or standard deviation) of an asset return is an important step in estimating how much risk a particular asset carries. However, the variance is not directly observable from

a time series because there is only one observation at each time point. The seminal paper by Engle¹ which introduced the Autoregressive Conditional Heteroscedasticity (ARCH) model allowed researchers to obtain such an estimate. Since then, many more developments in modeling heteroscedasticity or the changing variance of the process have been made, such as the generalization to the (G)ARCH model which includes lagged components of squared innovations in the model.² Furthermore, it is widely accepted that financial volatilities are correlated across assets and markets, and models have been developed which make use of this fact. The multivariate Generalized

*Correspondence to: ensor@rice.edu

¹Department of Statistics, Rice University, Houston, TX, USA

²Market Analysis & Planning, HP, Palo Alto, CA, USA

³San Jose State and Northwestern University, San Jose, CA, USA

Conflict of interest: The authors have declared no conflicts of interest for this article.

Autoregressive Conditional Heteroscedasticity (MGARCH) model is used for studying the relationships between the volatilities and covolatilities of multiple stocks. An excellent survey on MGARCH models (with over 1000 citations) is provided by Bauwens.³ Highlighted are important early contributions to the vast MGARCH literature including Refs 4–13, and 14.

Challenges of MGARCH modeling include the fact that the covariance matrix must be positive definite at every time point as well as the obvious curse of dimensionality leading to a large number of parameters for even simplistic model formulations. A very popular MGARCH model from the early era that solves these issues is the constant conditional correlation (CCC) model.⁶ The CCC model decomposes the conditional covariances of the returns into the conditional correlations and conditional standard deviations. The basic premise of the CCC model is that the conditional correlation matrix between returns is constant over time but the univariate conditional standard deviations change over time. This changing structure is captured using time series models for the conditional standard deviations.

Of further importance is recognition that markets structurally change over time and that any model or system must adapt to this change. Dynamic conditional correlation (DCC) models are designed to adapt to market conditions. Variants of the DCC models have become one of the most popular financial time series methods in the recent literature due to this need to capture changing market dynamics. In this article, we highlight some of the new findings of this important class of MGARCH models. Another strategy to adapt to changing market dynamics is through Markov switching models where model parameterizations change with market conditions. In the univariate setting, GARCH models with varying parameters, such as the regime switching model by Fleming and Kirby, and Klaassen^{15,16} are shown to improve volatility forecasting. Similar improvements are observed in regime switching models for multivariate covolatility models,^{17–20} where copulas are used to capture the multivariate dependence. A regime switching MGARCH model is highlighted in this article, with an example of stock market returns through the late 1990s and early 2000s, including the fall of ENRON. Structural change from the subprime crises of 2008 is a current topic in many studies (e.g., Refs 21, 22, and 23) and our methodologies hold up well under these dynamics. For example, we are able to foreshadow events such as the decline of Lehmann Brothers with large lead times. However, for this review manuscript we focus our examples on the earlier time period, which ends with the 2001

financial crisis. We first give a brief introduction to the univariate GARCH model followed by the focus of this advanced review, namely the multivariate version or MGARCH model with some of the key extensions including the regime switching model. We follow the extensive model exploration section with a section explaining estimation of one specific model, the hierarchical regime switching dynamic covariance (HRSDC) model, on which our example is based. We then present an example that highlights the utility of a hierarchical MGARCH model used to model the covariance structure within and between market sectors where the correlation structure changes with market conditions. We conclude with a discussion of general computing resources and general comments.

A MOTIVATING EXAMPLE: THE DYNAMICS OF SECTORS AND STOCKS WITHIN SECTORS

Consider daily stock returns of the largest 21 companies in three different market sectors, of Energy, Financial, and Technology for the time period January 2, 1998 to December 31, 2001. Figure 1 displays the data for a time period of 6 months. Each cluster includes the 21 time series plots for a particular sector during the 6-month period. Included in Figure 2 are returns for two different 6-day periods. Each of the two panels demonstrates a different type of comovement between the stocks. The left panel demonstrates a time period when the stock returns are following each other closely. The right panel demonstrates a time period with less correlation. *We are interested in detecting this difference.* If we can model this different structure and forecast the changes, we can use our results to answer specific questions about market dynamics as well as how individual stocks behave with respect to these dynamics.

Modeling the multivariate covariance or the covolatility is the objective of the general class of MGARCH models. We begin this review with a basic discussion of univariate GARCH models and then move in to the complexities of expanding to multiple time series.

MODELING VOLATILITY AND COVOLATILITY

Volatility or the movement in stock prices has many definitions in financial analysis. We define volatility as the conditional variance of the daily stock return prices, conditioned on information prior to that day.




FIGURE 1 | Daily stock returns of 63 series spanning three sectors.

As our time series realization is daily stock returns, the conditional variance is not directly observed and cannot be directly estimated because we have a single observation. One could use information on a finer timescale to estimate the variance, e.g., using intraday stock prices to estimate the daily variability in the returns, or using daily returns to estimate the volatility of a monthly return process. These empirical estimates of volatility, often called *realized volatility* require careful examination due to bias from market microstructure and other factors that get rolled up into the estimate (for more on realized volatility, see Refs 24–30, and 31). One might also estimate the volatility of an equity by reverse engineering a fixed model for price, such as the Black-Scholes option pricing model which has as one of its parameters the conditional variance of the underlying process. This resulting volatility estimate is referred to as *implied volatility* and has implicit assumptions, not the least of which is a strong belief in the underlying option pricing model.

The time series solution to the challenge of estimating volatility, is to model the daily conditional

variance or standard deviation directly (covariance in the multivariate case). Specifically, consider the univariate stochastic process $\{y_t\}$, with $t \in T$, where the time index $T = \mathbb{Z}$, the set of integers. The volatility is defined as the conditional variance, or $\sigma_t^2 = \text{Var}[y_t | y_j, j < t]$. The two basic classes of models for the conditional variance bifurcate based on whether a fixed or stochastic equation is used to describe the evolution in volatility. The GARCH/MGARCH class of models fall in the first category of a fixed equation, but as we highlight in this overview there are natural perturbations that expand this class of models to changing market dynamics.

The univariate GARCH model for the process $\{y_t\}$ is defined as

$$y_t = \mu_t(\theta) + \epsilon_t, \quad \epsilon_t = \sigma_t(\theta) z_t.$$

where $\mu_t(\theta)$ is the process mean conditional on the information up to time t , θ are mean model parameters, σ_t is the conditional standard deviation of the process, and the innovation process z_t has $E[z_t] = 0$, $\text{Var}[z_t] = 1$ and is uncorrelated in time. The random




FIGURE 2 | Daily stock returns of 63 series spanning three sectors. High- and low-correlation regimes are depicted. The estimated regime probabilities, rounded to 1 decimal place, are given at the top of each graph.

process represented by the conditional variance σ_t^2 is modeled as a function of itself at previous time points as well as a function of the squared values of the mean adjusted returns at previous time points. Specifically, the classic GARCH model is given by

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^p \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^q \beta_j \sigma_{t-j}^2$$

where $\alpha_0 > 0$, $\alpha_i \geq 0$, $\beta_j \geq 0$, and $\sum_{i=1}^{\max(p,q)} (\alpha_i + \beta_j) < 1$ for mean reversion of the filter. One might also model the conditional standard deviation rather than the conditional variance, as in our example of this review.

There are numerous and important perturbations of the classic GARCH model that are important for modeling financial returns. Some of the most important improvements include: (1) using different distributions for the process z_t (e.g., using a t -distribution with low degrees of freedom can capture the heavier tails of the volatility), (2) asymmetric modeling of upside and downside volatility (e.g., this

can be accomplished through the exponential GARCH model), and (3) extending the basic GARCH model to a random coefficient model. For a full exposition of GARCH modeling, we refer the reader to Refs 17, 32, and 33.

The simple GARCH model for the conditional variance is scaled to the multivariate setting under numerous strategies. The modeling complexities increase exponentially as the number of series increases and various methods to handle this complexity result in different modeling paradigms. To understand these complexities, we outline some of the classes of MGARCH models that result from different approaches to mastering the complexity. We conclude with a real example using one approach. A comprehensive comparison of all the approaches is beyond the scope of this overview.

The basic MGARCH process is defined as follows. Suppose we have a vector stochastic process $\{y_t\}$ of dimension $K \times 1$ governed by the following equations:

$$y_t = \mu_t(\theta) + \epsilon_t,$$

where $\mu_t(\theta)$ is the mean vector conditional on the information up to time t and parameters θ , and

$$\epsilon_t = H_t^{1/2}(\theta) z_t,$$

where $H_t^{1/2}(\theta)$ is a $K \times K$ positive definite matrix. The stochastic process z_t has the following properties:

$$E[z_t] = 0$$

$$\text{Var}[z_t] = I_K,$$

with I_K denoting the $K \times K$ identity matrix and z_t is uncorrelated over time. The matrix $H_t^{1/2}$ is a $K \times K$ positive definite matrix such that H_t is the conditional covariance matrix of y_t , again conditional on the information of the process up to time t .

Variations in MGARCH models result from how one approaches modeling H_t , the conditional covariance.

MGARCH Models for the Conditional Covariance

MGARCH models that are direct generalizations of the univariate GARCH model include the VEC (name initiates from the *vectorization* of the covariance matrix) model of Ref 5 and the BEKK (named after the originators Baba, Engle, Kraft, and Kroner) model of Ref 9. In short, the VEC model vectorizes H_t and models the vector of conditional covariances in an ARMAX (multivariate ARMA) fashion. Although a natural extension to the univariate GARCH model, issues arise with ensuring a positive definite conditional covariance matrix as well as the escalating numbers of parameters as the dimension increases. The BEKK model reformulates this approach by writing a multivariate model for H_t without the vectorization of the matrix, guaranteeing a positive definite matrix at each time point and reducing the number of parameters. The BEKK model is a special case of the VEC model encompassing all the practical VEC formulations.⁹ The number of parameters in the BEKK model remains high, however, especially as the model complexity increases. Due to the high number of unknown parameters, these models are rarely used when the number of series is larger than three or four.³ For a detailed review of these methods, see Ref 17.

A second group of MGARCH models that directly address the issue of escalating model parameters are factor models. Factor models reduce the

parameterization due to a common dynamic structure they impose on all the elements of H_t . These factor models capture the comovements of the multiple series by a small number of common underlying variables.⁸ There exist several variants of the MGARCH factor models including the full-factor multivariate GARCH (FF-MGARCH) model of Ref 14 where:

$$H_t = W \Sigma_t W',$$

and W is a $K \times K$ triangular matrix with ones on the diagonal and the matrix $\Sigma_t = \text{diag}(\sigma_{1,t}^2, \dots, \sigma_{K,t}^2)$ where $\sigma_{i,t}^2$ is the conditional variance of the i th factor. In this formulation, the i th element of $W^{-1}\epsilon_t$ can be separately defined as a univariate GARCH model. An overview of general time series factor models is found in Ref 34.

A third group of MGARCH models are composed of linear combinations of several univariate models for the covariances. This class includes models such as the orthogonal-GARCH (O-GARCH) and latent factor models. In the O-GARCH model,^{4,11} the $K \times K$ time-varying variance matrix H_t is generated by $m \leq K$ univariate GARCH models. Parameter identifiability becomes an issue in the case of a Gaussian model but is mitigated when mixtures of normal distributions are considered. We refer the interested reader to Ref 17 or Ref 35 for more detail.

MGARCH Models for Conditional Correlation and Univariate Variance (or Standard Deviation)

Finally, consider the class of MGARCH models that include nonlinear combinations of univariate GARCH models. These models allow analysts to specify univariate conditional variances (or standard deviations as in our example) and the conditional correlation matrix of the multiple series. This group includes the important CCC model of Ref 6 which is defined as follows:

$$H_t = D_t R D_t = \left(\rho_{ij} \sqrt{h_{iit} h_{jtt}} \right),$$

where

$$D_t = \text{diag} \left(h_{11t}^{1/2}, \dots, h_{KKt}^{1/2} \right), \tag{1}$$

The diagonal element of H_t or univariate conditional variance h_{iit} can be defined as any univariate GARCH model. The constant (in time) correlation matrix $R = [\rho_{ij}]$ is a symmetric positive definite matrix with $\rho_{ii} = 1, \forall i$.

Time-Dependent Conditional Correlation Models

When the assumption that the conditional correlations are constant is relaxed as in Refs 7, 10, and 13 we refer to the model as a DCC model. Then the question becomes, how to model the changing correlation structure? One strategy is that of Ref 7 that uses the Fisher transformation of the correlation coefficient, guaranteeing the positive definiteness of the conditional correlation matrix. However, this model formulation is limited to a bivariate process.

A DCC model useful for high-dimensional time series is the one of Ref 13 and is defined as:

$$H_t = D_t R_t D_t, \tag{2}$$

where D_t is defined in Eq. (1) and h_{iit} can be defined as any univariate GARCH model, and

$$R_t = (1 - \alpha - \beta) R + \alpha \Psi_{t-1} + \beta R_{t-1}. \tag{3}$$

where α and β are nonnegative parameters satisfying $\alpha + \beta < 1$, R is a symmetric $K \times K$ positive definite parameter matrix with diagonal elements equal to 1, and Ψ_{t-1} is the $K \times K$ correlation matrix of ϵ_τ for $\tau = t - M, t - M + 1, \dots, t - 1$. Thus, the ij th element of Ψ_{t-1} is:

$$\Psi_{ij,t-1} = \frac{\sum_{m=1}^M u_{i,t-m} u_{j,t-m}}{\sqrt{\left(\sum_{i=1}^M u_{i,t-m}^2\right) \left(\sum_{j=1}^M u_{j,t-m}^2\right)}}$$

where $u_{it} = \epsilon_{it} / \sqrt{h_{iit}}$, representing a standardization of ϵ_t . The matrix Ψ_{t-1} can be expressed as:

$$\Psi_{t-1} = B_{t-1}^{-1} L_{t-1} L'_{t-1} B_{t-1}^{-1},$$

where B_{t-1} is a $K \times K$ diagonal matrix. The i th diagonal element is given by

$$\left(\sum_{b=1}^M u_{i,t-b}^2\right)^{1/2}$$

and $L_{t-1} = (\mathbf{u}_{t-1}, \dots, \mathbf{u}_{t-M})$ is a $K \times M$ matrix, with $\mathbf{u}_t = (u_{1t}, \dots, u_{Kt})'$. Also, $M \geq K$ to ensure that R_t is positive definite. In essence, Eq. (3) is an autoregressive moving average approximation to the changing correlation structure.

The DCC model of Ref 10 has the same form for H_t (see Eq. (2)) but with the time-varying correlation matrix given by:

$$R_t = \text{diag} \left(q_{11,t}^{-1/2}, \dots, q_{KK,t}^{-1/2} \right) \\ Q_t \text{diag} \left(q_{11,t}^{-1/2}, \dots, q_{KK,t}^{-1/2} \right).$$

The $K \times K$ symmetric positive definite matrix $Q_t = (q_{ij,t})$ defines the evolution of the correlation and is given by:

$$Q_t = (1 - \alpha - \beta) \bar{Q} + \alpha \mathbf{u}_{t-1} \mathbf{u}'_{t-1} + \beta Q_{t-1},$$

with $u_{it} = \epsilon_{it} / \sqrt{h_{iit}}$. In this setting, \bar{Q} is the $K \times K$ unconditional variance matrix of \mathbf{u}_t , and α and β are nonnegative scalar parameters satisfying $\alpha + \beta < 1$. Again, the fundamental issue is how we choose to model H_t through different parameterizations of the correlation matrix R_t . In this case, the time-varying correlation R_t is a weighted average of the unconditional variance, the observed squared values, and the previous conditional variance of the standardized conditional standard deviations. The matrix Q_t is not a correlation matrix, so we rescale by the diagonal elements of Q_t as we are modeling the correlation R_t .

The flexibility of the DCC model is extended even further by Ref 36 providing a rich parameterization for the correlation process based on *a priori* knowledge of asset grouping. An MGARCH strategy for clustering time series, absent *a priori* knowledge, is developed in Ref 37. Recently, a multiplicative DCC model is used to model the dynamic volatility and correlation structure of electricity futures.³⁸ The authors are able to identify different dynamics for long- and short-term contracts, resulting in stronger forecasting performance when compared with the standard DCC model. In addition to providing another good overview of general MGARCH models, DCC models are extended to include realized measures of correlation in Ref 39.

The DCC models allow for time-varying correlation structure but this structure is allowed to change at each time point through a structured model formulation. The bias of estimation in the DCC model for high dimensions is noted by multiple authors (e.g., Ref 40). In a recent study,⁴¹ the inconsistent feature of the traditional two-step estimation of the DCC model is examined with the authors noting that the consistency of the local correlation structure estimate is not substantiated which then impacts the estimation and interpretation of the dynamic correlation structure. The suggested fix is formulation of the

correlation process as a linear MGARCH process; this formulation is referred to as the cDCC estimator and requires only a marginal increase in computational time for large systems.⁴¹ The cDCC model is extended to robust forecasting of DCCs in Ref 42. In essence, strategies are introduced which bound the influence of large innovations on the estimates of the cDCC model parameters building off of previous work in the univariate setting by Ref 43 and the use of M-estimates for the correlation dynamics.

The MGARCH model of Ref 12 also allows the correlations between the series to be time-varying using a two-step approach by addressing the multivariate dependence first. The first step removes unconditional correlation by taking principal components (PCs) of the data. The conditional means and variances of each PC are then modeled by a univariate GARCH model. To construct the MGARCH model, the PC mapping is reversed using the estimated PCs from the univariate GARCH models rather than the data-based principal components. The inverse of the PC construction is used to transform the conditional moments of the PCs into the conditional mean and variance of the data series themselves. This formulation of the MGARCH model is easy to estimate as it only involves estimates of univariate GARCH models following a classic PC of the original series.

Market-Dependent Regime Switching Dynamic Correlation Model

We are seeking to model a system that is constant within blocks of time but changes as the market dynamics change. The Regime Switching for Dynamic Correlations (RSDC) model of Ref 19 allows time-varying correlation between the series by allowing the system to switch between regimes. The covariances in the system are decomposed into correlations and standard deviations, and the correlation matrix follows a regime-switching model, i.e., the correlations are constant within regime but different across regimes. A Markov chain governs the transitions between the regimes. The model is described as follows. Again, assume that K -variate process \mathbf{Y}_t has the form:

$$\mathbf{Y}_t = H_t^{1/2} \mathbf{U}_t,$$

where \mathbf{U}_t is an i.i.d. $(0, I_k)$ process. That is, each element in \mathbf{U}_t has zero mean and unit variance. Also, each series is uncorrelated with all remaining series. As before, the time-varying covariance matrix H_t can be decomposed as:

$$H_t = S_t \Gamma_t S_t \tag{4}$$

where S_t is a diagonal matrix composed of the standard deviations $s_{k,t}$, for $k = 1, \dots, K$ and the matrix Γ_t contains the correlations (we switch notation from D_t and R_t in the previous section to limit confusion on the estimator we develop in the next section for the RSDC model). Specifically, Γ_t follows a regime-switching model:

$$\Gamma_t = \sum_{r=1}^R I_{\{\Delta_t=r\}} \Gamma_r$$

where Δ_t is an unobserved Markov process independent of \mathbf{U}_t that can take R possible values ($\Delta_t = 1, 2, \dots, R$) and I is the indicator function. The $K \times K$ matrices Γ_r are correlation matrices with $\Gamma_r \neq \Gamma_{r'}$ for $r \neq r'$. The probability law governing Δ_t is defined by its transition probability matrix Π . The probability of going from regime i in period t to regime j in period $t + 1$ is denoted by $\pi_{i,j}$ and the limiting probability of being in regime n is π_n . We assume that the Markov chain is ergodic and irreducible.

The RSDC model has several good properties that make it appropriate for our example problem. First, we expect that our system will be in different regimes depending on the state of the market, and that there is a structural change in market dynamics. The resulting regime-switching structure improves our understanding of this structural change. Secondly, we can obtain estimates of the parameters, even when the number of time series is large. Our example involves detecting changes in an overall market, which in turn includes many time series within multiple sectors. Another property of the RSDC model is that it is easy to impose that the variance matrices are positive semidefinite. Finally, from the model for the univariate volatility we obtain multistep ahead predictions and couple these predictions with forecasted regimes, thereby obtaining forecasts of the covariance matrix.

The example developed in this study is the adaptation of Pelletier's model to one which is appropriate to the problem of sector relationships and securities nested within the sectors. We are interested in detecting changes in market behavior as determined by the dependence structure between market sectors. Instead of using one measurement of market sector behavior (such as an index fund), we use multiple measurements of market sector behavior given by the multiple stock return series. We then estimate correlations both between and within the sectors using the multiple measurements. Thus, each $\Gamma_r, r = 1, \dots, R$, becomes a block matrix whose elements include estimates of the average within- and between-sector correlations. For our example which includes three market sectors, each

Γ_r has six parameters (three within-sector correlations and three between-sector correlations).

ESTIMATION OF THE HIERARCHICAL RSDC MODEL

For our example, we focus on the regime-switching dynamic correlation model, adapting this model to a hierarchical investigation to handle the covolatility of sectors and stocks within sectors. We refer to this model as the hierarchical regime-switching dynamic correlation (HRSDC) model. Estimation of the RSDC model uses a two-step estimation procedure as outlined in Ref 19. In Ref 17, we adapt this estimation procedure, along with programs initially developed by Pelletier, to accommodate multiple measurements for the same quantity as present in our hierarchical example (i.e., the HRSDC model).

To proceed, denote the vector of model parameters as θ , split them into two groups $\theta = (\theta_1, \theta_2)$, where θ_1 and θ_2 contain the parameters estimated in the first and the second step of the algorithm, respectively. In the first step, we estimate the univariate volatility models using the asymmetric power ARCH (apARCH) model in which a power transform of the conditional standard deviation (rather than variance) is modeled as follows:

$$s_t^\delta = \omega + \sum_{i=1}^q (\tilde{\alpha}_i |y_{t-i}| - \gamma_i \epsilon_{t-i})^\delta + \sum_{j=1}^p \beta_j s_{t-j}^\delta \quad (5)$$

with $\tilde{\alpha}_i = \alpha_i/E|\tilde{u}_t|$. For the ARMACH model (see Ref 44), $\delta = 1$, but the general asymmetric power ARCH (apARCH) allows for different power transformations. The asymmetry parameter γ captures the different upside and downside volatility. It is noted in Ref 44 that the conditional standard deviation often exhibits stronger autocorrelation structure than the conditional variance. However, we could choose the best fitting univariate volatility model at this stage and the R package *rugarch*⁴⁵ provides multiple choices.

Next, denote QL_1 as the log-likelihood where the correlation matrix is taken to be an identity matrix:

$$QL_1(\theta_1; Y) = -\frac{1}{2} \sum_{t=1}^T (K \log(2\pi) + 2 \log(|S_t|) + U_t' U_t).$$

As this is just the sum of K univariate log-likelihoods, maximizing it is equivalent to maximizing (separately) each univariate log-likelihood.

For the second step of the estimation, we need to maximize QL_2 which is the log-likelihood given θ_1 :

$$QL_2(\theta_2; Y, \theta_1) = -\frac{1}{2} \sum_{t=1}^T (K \log(2\pi) + 2 \log(|\Gamma_t|) + U_t' \Gamma_t^{-1} U_t).$$

We can maximize QL_2 using the EM algorithm. We make use of the results in Ref 46 using the particle filter because the Markov chain Δ_t is unobserved. To make inference on the state of the Markov chain, we compute values for $\hat{\xi}_{t|t}$ and η_t which are defined as follows: Let $\hat{\xi}_{t|t}$ be an $R \times 1$ vector containing the probability of being in each regime at time t conditional on observations up to time t . Thus, $\hat{\xi}_{t|t}$ is an $R \times 1$ vector containing the elements $P\{\Delta_t = j | U_T; \theta_2\}$. Let η_t be the $R \times 1$ vector whose j th element is the density of U_t conditional on past observations and being in the j th regime at time t . To obtain these estimates, we use the steps outlined below. The differences in our estimation procedure and in that of Pelletier's can be observed in steps 1 and 2 below, whereas steps 3–5 remain identical. We compute the starting values in a different manner due to the fact that we have multiple observations for estimating between- and within-sector correlations. Also, we found it useful to work on the log scale due to the high dimensionality of our problem.

Estimation algorithm:

1. Choose a starting value for $\hat{\xi}_{1|0}$ and $\theta_2 = (\Pi, \Gamma_1, \dots, \Gamma_R)$.

We set each element in $\hat{\xi}_{1|0} = 1/R$. To obtain starting values for the correlation matrices in each regime, we compute an estimate of the correlation matrix of the data using the numerically accurate corrected two-pass method described in Ref 47. We will denote the empirical estimator as $\hat{\Sigma}$. Next, we set $\Gamma_1 = 0.8\hat{\Sigma}$ and $\Gamma_2 = 1.2\hat{\Sigma}$. We initially set the diagonal elements of Π equal to 0.7 as we expect some persistence in the Markov chain. We tried several reasonable starting values and obtained convergence to the same values. Thus, the starting values did not affect our results—only the computing time.

2. Compute η_t by evaluating the multivariate normal density of U_t conditional on past observations and being in each regime at time t or,

$$\tilde{\eta}_{r,t} = -\frac{1}{2} \left(K \log(2\pi) - \log|\hat{\Gamma}_r| - U_t' \hat{\Gamma}_r^{-1} U_t \right),$$

$$r = 1, \dots, R$$

with $\eta_{r,t} = e^{\tilde{\eta}_{r,t}}$.

Then compute iteratively for $t = 1, \dots, T$

$$\hat{\xi}_{t|t} = \frac{\left(\hat{\xi}_{t|t-1} \odot \eta_t\right)}{1' \left(\hat{\xi}_{t|t-1} \odot \eta_t\right)}$$

$$\hat{\xi}_{t+1|t} = \Pi \cdot \hat{\xi}_{t|t}.$$

We use \odot to denote element-by-element multiplication. The $R \times 1$ vector $\hat{\xi}_{t+1|t}$ contains the probabilities of being in each regime at time $t + 1$ conditional on observations up to time t .

3. Compute the smoothed-state variables starting from $t = T - 1, \dots, 1$ as follows:

$$\hat{\xi}_{t|T} = \hat{\xi}_{t|t} \odot \left\{ \Pi' \left[\hat{\xi}_{t+1|T} (\div) \hat{\xi}_{t+1|t} \right] \right\}.$$

We use (\div) to denote element-by-element division.

4. Update θ_2 as follows:

$$\hat{\Pi}_{i,j} = \frac{\sum_{t=2}^T P \left[\Delta_t = j, \Delta_{t-1} = i | \hat{U}_T; \hat{\theta}_2 \right]}{\sum_{t=2}^T P \left[\Delta_{t-1} = i | \hat{U}_T; \hat{\theta}_2 \right]}$$

$$\tilde{\Gamma}_r = \frac{\sum_{t=1}^T \left(\hat{U}_t \hat{U}_t' \right) P \left[\Delta_t = r | \hat{U}_T; \hat{\theta}_2 \right]}{\sum_{t=1}^T P \left[\Delta_t = r | \hat{U}_T; \hat{\theta}_2 \right]}.$$

We rescale the $\tilde{\Gamma}_r$ matrices, so that we have ones on the diagonal as this does not happen naturally; we let $\tilde{\Gamma}_r$ denote the estimate before rescaling and $\hat{\Gamma}_r$ represent the estimate after rescaling. Finally, we compute $\hat{\Gamma}_r$ as:

$$\hat{\Gamma}_t = D_r^{-1} \tilde{\Gamma}_r D_r^{-1}$$

where D_r is a diagonal matrix with $\sqrt{\tilde{\Gamma}_{i,i,r}}$ on row i and column i . This rescaling ensures that we have ones on the diagonal and off-diagonal elements between -1 and 1 .

Now, we will replace the starting value $\hat{\xi}_{1|0}$ with its smoothed estimate $\hat{\xi}_{1|T}$ (its estimate conditional on information from $t = 1, \dots, T$). We do this as it has been shown in³⁰ that the MLE estimate for $\hat{\xi}_{1|0}$ is given by $\hat{\xi}_{1|T}$.

This makes the choice of the starting value for $\hat{\xi}_{1|0}$ less important as we are replacing it by the MLE in the next iteration.

5. Repeat steps 2 through 4 until convergence occurs (maximum element of the difference in successive parameter estimates is small).

Under the usual assumptions of quasi-maximum likelihood estimation, the two-step maximum likelihood estimates obtained from the procedure detailed above are consistent and asymptotically normal.¹⁹

SIMULATION STUDY

We conducted a small simulation study of our implementation of the RSDC model for two regimes, grounding our simulation on the model estimates similar to our example that follows in the next section. The modest study yields good results as shown in Tables 1 and 2. For the simulation, we use the simpler ARMACH model with asymmetry parameter $\gamma = 0$ and power parameter $\delta = 1$ in Eq. (5). That is, $K = 63, T = 1004, R = 2$, and the correlation model parameters from the Results section. For the univariate volatility models, we simulate from a GARCH(1,1) model: $s_{k,t}^2 = 0.0005 + 0.6\epsilon_{k,t-1}^2 + 0.1s_{k,t-1}^2$. These GARCH(1,1) parameters were chosen to obtain the same level of volatility as in our stock sector data. In out-of-sample predictions, we predict the correct regime at 76% of the time points in our simulation.

Example

To demonstrate the practical application of MGARCH models and specifically the HRSDC model, we return to our example of 63 large cap stocks spanning three different market sectors, Energy, Financial, and Technology, for the time period January 2, 1998 to December 31, 2001. To obtain stocks from these specific sectors, we used the Standard Industrial Classification Code List published by the U.S. Securities and Exchange Commission (SEC).^a The aim of this study was to use the 20 companies with the largest market capitalization (on January 6, 1998) in each sector. However, due to missing data issues and classification code errors, we considered up to the largest 28 companies. For each sector, we used the largest 21 companies with the correct industry classification code and no missing data. Thus, our system includes 63 series of daily stock returns for the time period January 2, 1998 to December 31, 2001. We define the daily stock return at time t as $Y_t = (\text{Price}_t - \text{Price}_{t-1}) / \text{Price}_{t-1}$.

Figure 1 displays the return series for a time period of 5 months. Each cluster includes the 21 time series plots for a particular sector during the 6-month period. Included in Figure 2 are the same return series

TABLE 1 | Correlation and Transition Probability Estimates in Simulation Study Based on 25 Realizations

Correlation Between	True R1	Est. R1	SD	True R2	Est. R2	SD
Energy/Energy	0.354	0.336	0.018	0.644	0.530	0.016
Energy/Financial	0.093	0.010	0.016	0.170	0.129	0.026
Energy/Technical	0.023	0.045	0.014	0.142	0.098	0.028
Financial/Financial	0.480	0.427	0.0167	0.525	0.44	0.022
Financial/Technical	0.119	0.141	0.0124	0.361	0.28	0.024
Technical/Technical	0.149	0.1961	0.021	0.674	0.546	0.016
Transition probability	0.834	0.798	0.017	0.830	0.793	0.021

TABLE 2 | Univariate Volatility Estimates in Simulation Study Based on 25 Realizations

Parameter	True	Estimated	SD
a	0.001	0.001	0.000
α	0.600	0.593	0.069
β	0.100	0.095	0.053

for two different 6-day periods. The two time epochs highlight different relationships between the series, in other words, different patterns of comovement between series within sectors as well as the comovement between sectors.

RESULTS

We fit the HRSDC model to the series of 63 returns, incorporating the fact that there are three sectors each consisting of 21 different companies. The HRSDC model includes two correlation regimes of correlation representing periods of high correlation and periods where the correlation is lower. A two regime model captures the transition between periods of *typical* correlation and periods of *high* correlation, with the latter being a direct result of the systemic risk of markets.^{21,23,48} Again, we model the correlations within each sector as well as the correlation between the sectors in a hierarchical fashion. Tables 3 and 4 include parameter estimates for this model. To ensure proper convergence of the estimation algorithm, we repeated the process with several very different sets of starting values and achieved convergence to the same parameter estimates in each case.

Of additional interest is the relative magnitude of correlations within and between the three sectors. Figure 3 displays the one-step ahead forecasts for the average correlation between and within sectors for a 9-month time period in 1999–2000. The diagonal

TABLE 3 | Correlation and Transition Probability Estimates with Standard Errors

Correlation Between	Regime 1	SE	Regime 2	SE
Energy/Energy	0.354	0.009	0.644	0.005
Energy/Financial	0.093	0.012	0.170	0.022
Energy/Technology	0.023	0.009	0.142	0.014
Financial/Financial	0.480	0.006	0.525	0.007
Financial/Technology	0.119	0.019	0.361	0.018
Technology/Technology	0.149	0.008	0.674	0.007
Transition Probabilities	0.8340	0.1807	0.8305	0.0674

TABLE 4 | Univariate Volatility Parameter Estimates with Standard Errors

Parameter	Average	Average SE
ω	0.002	0.006
$\tilde{\alpha}_1$	0.065	0.056
γ	0.547	0.459
β_1	0.895	0.157
δ	1.170	0.590
t -dist skew	1.084	0.074
t -dist shape	10.89	9.06

elements represent the correlations within sectors where the off-diagonal elements represent the correlations between sectors. Again these correlations are a weighted average of the correlation estimates for the two regimes with the weights determined by the probability of being in each regime. The correlation between sectors are not near one but do increase during the higher correlation regimes. The strongest cross-sector correlation is observed between the financial and technology sectors. Although the magnitudes of the within- and between-sector correlations differ, the pattern in the forecasts is identical because this




FIGURE 3 | One step ahead forecasts for average correlations between sectors. The scale for each plot is 0 to 0.8.

pattern is driven by the probability of the two fixed regimes.

Let us revisit the the issue of differing comovement behavior types. In Figure 2, we displayed two panels: one in which the stock returns followed each other closely and the other in which the stock returns exhibited less correlation. Our model is detecting these differing comovement behaviors by predicting with high probability the regime with higher correlation when the stocks are following each other closely and predicting with high probability the regime with lower correlation when the stocks exhibit much less comovement.

Using the HRSDC Methodology to Detect Anomalous Behavior

With the HRSDC model, we can estimate (and forecast) the covariance matrix, H_t , at time t by combining the estimates (or forecasts) for the standard deviations for each asset and the correlation matrix as in Equation 2. Note that in our example the correlation structure bounces between two different matrices; however, the modeled univariate standard deviations will change dramatically consistent with the observed behavior of the returns for the underlying security. As a graphical display of the changes in the overall

market structure, we examine the biplot of the associated principal components of the estimate of the covariance structure, in other words, our model of the $K \times K$ covariance matrix. For each time step, we retain sufficient principal components to explain 90% of the total variation of the processes. The evolution of the principal components directs us to unusual events such as a market crash or stock anomalies. The evolution of such structures can identify favorable market conditions for trading strategies, ultimately leading to performance improvements, see Ref 17 for more details.

A well-known outlier in this data is the Enron collapse. A time line of the Enron collapse is well-documented due to the litigation between the U.S. SEC and various Enron executives. On October 16, 2001, Enron announced allegedly ‘nonrecurring’ losses of approximately \$1 billion. When examining the daily biplots for our model, we first detect Enron as a well-identified outlier on October 4, 2001, as shown in Figure 4. The biplots look very similar with respect to Enron’s effects until the end of our sample or December 31, 2001. To remind the reader of the timeline of the Enron collapse, on October 29 and November 1, 2001, the two leading credit rating agencies downgraded Enron’s credit rating. On November 8, 2001, Enron announced its intention to restate its financial statements for 1997 through 2000




FIGURE 4 | Biplot for covariance matrix of October 4, 2001.

and the first and second quarters of 2001 to reduce previously reported net income by an aggregate of \$586 million. On November 21, 2001, Enron's credit rating was downgraded to 'junk' status. On December 2, 2001, Enron filed for bankruptcy, making its stock, which less than a year earlier had been trading at over \$80 per share, virtually worthless. We compared our model's outlier detection ability to a model using a weighted estimate of the covariance matrix using a moving window of 70 trading days, with weights based on the minimum volume ellipsoid estimator. The biplots for this estimator detected Enron as an outlier on November 28, 2001, well after this detection could be of any use to the investor.

There were other, albeit less well-known, outliers that were detected through our methodology, which again relies on the HRSDC model of the covolatility (see Refs 17 and 18 for a full explanation of all examples). For example, on June 2, 2000, we identified an unusual pattern in the returns for Silicon Graphics, Inc. (SGI) (Figure 5). On further

investigation of the company's filings with the U. S. SEC,^b we learned that there was an 8-K filing on June 9, 2000 concerning a spin-off of the company. The 8-K filings are defined as 'current report filings' and are required when a company has specific 'trigger' events transpire. These reports must be filed within a few days of the event, and the report must include details about the event, how it affects the company, and its impact on shareholders. This filing was made 1 week after we identified SGI as an outlier in terms of its behavior with respect to the 'system' of stocks.

In total, we identified 16 anomalies during our study period. Of the 16 anomalies identified, we were able to find explanations for 13 through subsequent SEC filings. Again, these detections could be useful to an investor as they were detected before the the filings with the SEC. Furthermore, when this methodology is implemented on the more recent market collapse of 2008 we are able to detect the fall of Lehmann Brothers with the same degree of forewarning. This detection is based purely on price movement of individual securities when viewed as a system of securities with

available through R to easily implement the dynamic and regime-switching conditional correlation models. These models are an important component of a financial time series suite of tools to understand stock behavior under changing market structure.

NOTES

^a <http://www.sec.gov/info/edgar/siccodes.htm>.

^b <http://www.sec.gov/edgar/searchedgar/companysearch.html>.

REFERENCES

- Engle R. Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation. *Econometrica* 1982, 50:987–1008.
- Bollerslev T. Generalized autoregressive conditional heteroskedasticity. *J Econom* 1986, 31:307–327.
- Bauwens L, Laurent S, Rombouts JVK. Multivariate GARCH models: a survey. *J Appl Econom* 2006 CORE Discussion Paper 2003/31, 21:79–109.
- Alexander CO. Multivariate orthogonal factor GARCH, 1997. University of Sussex, Mimeo.
- Bollerslev T, Engle RF, Wooldridge JM. A capital asset pricing model with time varying covariances. *J Polit Econ* 1988, 96:116–131.
- Bollerslev T. Modeling the coherence in short-run nominal exchange rates: a multivariate generalized arch model. *Rev Econ Stat* 1990, 72:498–505.
- Christodoulakis GA, Satchett SE. Correlated arch: modelling the time-varying correlation between financial asset returns. *Eur J Multivariate Anal* 2002, 139:351–370.
- Engle RF, Ng VK, Rothschild M. Asset pricing with a factor-ARCH covariance structure: empirical estimates for treasury bills. *J Econom* 1990, 45:213–238.
- Engle R, Kroner FK. Multivariate simultaneous generalized ARCH. *Econom Theory* 1995, 11:122–150.
- Engle RF. Dynamic conditional correlation—a simple class of multivariate GARCH models. *J Bus Econ Stat* 2002, 20:339–350.
- Kariya T. Mtv model and its application to the prediction of stock prices. In: Pullila T, Puntanen S, eds. *Proceedings of the Second International Tampere Conference in Statistics*, University of Tempere Finland, 1988.
- Klaassen F. Have exchange rates become more closely tied? Evidence from a new multivariate GARCH model, 1999. CentER Discussion Paper Series No. 1999-10, Tilburg University.
- Tse YK, Tsui AKC. A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. *J Bus Econ Stat* 2002, 20:351–362.
- Vrontos ID, Dellaportas P, Politis DN. A full-factor multivariate GARCH model. *Econom J* 2003, 6:311–333.
- Fleming J, Kirby C. Component-driven regime-switching volatility. *J Financ Econom* 2013, 2:263–301.
- Klaassen F. Improving GARCH volatility forecasts with regime-switching GARCH. *Empir Econ* 2002, 27:363–394.
- Davis GM. An examination of some open problems in time series analysis. Ph.D. Dissertation directed by K. B. ensor, Rice University, <http://hdl.handle.net/1911/18887>, 2005.
- Davis GM, Ensor KB. Risk: Diversification and performance precursors for stocks. *Am J Math Manage Sci* 2008, 28:255–273.
- Pelletier D. Regime switching for dynamic correlations. *J Econom* 2006, 131:445–473.
- da Liva Filho OC, Ziegelmann FA, Dueker MJ. Modeling dependence dynamics through copulas with regime switching. *Insur: Math Econ* 2012, 50:346–356.
- Arruda BP, Valls Pereira PL. Analysis of the volatility's dependency structure during the subprime crisis. *Appl Econ* 2013, 45:5031–5045.
- Audrino F. Forecasting correlations during the late-2000s financial crisis: the short-run component, the long-run component, and structural breaks. *Comput Stat Data Anal* 2013, 76:43–60.
- Kritzman M, Li Y, Page S, Rigobon R. Principal components as a measure of systemic risk. Technical Report, MIT Sloan Research Paper No. 4785-10, 2010.
- At-Sahalia Y, Fan J, Xiu D. High-frequency covariance estimates with noisy and asynchronous financial data. *J Am Stat Assoc* 2010, 105:1504–1517.
- Barndorff-Nielsen OE, Hansen PR, Lunde A, Shephard N. Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *J Econom* 2011, 162:149–169.
- Hautsch N, Kyj LM, Oomen RCA. A blocking and regularization approach to high-dimensional realized covariance estimation. *J Appl Econom* 2012, 27:625–645.
- Kyj LM, Ostdiek B, Ensor KB. Realized covariance estimation in dynamic portfolio optimization. Technical Report, Working Paper, Humboldt-Universit, 2009.
- Kyj LM. Estimating realized covariance using high frequency data. PhD Dissertation directed by K. B. Ensor and B. Ostdiek, Rice University. Available at: <http://hdl.handle.net/1911/22195>, 2008.
- Kyj LM, Ostdiek B, Ensor KB. Covariance estimation in dynamic portfolio optimization: a realized single factor

- model. In: *American Finance Association 2010 Atlanta Meetings Paper*, 2010.
30. Wang Y, Zou J. Vast volatility matrix estimation for high-frequency financial data. *Ann Stat* 2010, 38: 943–978.
 31. Zhang L. Estimating covariation: Epps effect, microstructure noise. *J Econom* 2011, 160:33–47.
 32. Peña D, Tiao GC, Tsay RS. *A Course in Time Series Analysis*. New York: John Wiley & Sons, Inc.; 2001.
 33. Tsay RS. *Analysis of Financial Time Series*. 3rd ed. New York: Wiley; 2010.
 34. Ensor KB. Time series factor models. *WIREs Comp Stat* 2013, 5:97–104. doi: 10.1002/wics.1245.
 35. Lanne M, Saikkonen P. A multivariate generalized orthogonal factor GARCH model. *J Bus Econ Stat* 2007, 25:61–75.
 36. Billio M, Caporin M, Gobbo M. Flexible dynamic conditional correlation multivariate GARCH for asset allocation. *Appl Financ Econom Lett* 2006, 2:123–130.
 37. Aielli GP, Caporin M. Variance clustering improved dynamic conditional correlation MGARCH estimators. *Comput Stat Data Anal* 2014, 76:556–576.
 38. Bauwens L, Hafner CM, Pierret D. Multivariate volatility modeling of electricity futures. *J Appl Econom* 2013, 28:743–761.
 39. Yu Y. Multivariate GARCH modeling and comparison to real kernel estimates. Doctoral Dissertation directed by D. Pelletier and D. Pearce, North Carolina State University, 2013.
 40. Hafner CM, Reznikova O. On the estimation of dynamic conditional correlation models. *Comput Stat Data Anal* 2012 1st issue of the Annals of Computational and Financial Econometrics; Sixth Special Issue on Computational Econometrics, 56:3533–3545.
 41. Aielli GP. Dynamic conditional correlation: on properties and estimation. *J Bus Econ Stat* 2013, 31: 282–299.
 42. Boudt K, Danelsson J, Laurent S. Robust forecasting of dynamic conditional correlation GARCH models. *Int J Forecast* 2013, 29:244–257.
 43. Muler N, Yohai V. Robust estimates for GARCH models. *J Stat Plann Inference* 2008, 138:2918–2940.
 44. Taylor S. *Modelling Financial Time Series*. Chichester: John Wiley & Sons, Inc.; 1986.
 45. A. Ghalanos. *rugarch: Univariate garch models*. R package version 1.2-9, 2013. Available at: <http://cran.r-project.org/web/packages/rugarch/>.
 46. Hamilton JD. *Time Series Analysis*. Princeton, NJ: Princeton University Press; 1994.
 47. Chan T, Golub G, LeVeque R. Algorithms for computing the sample variance: analysis and recommendations. *Am Stat* 1983, 37:242–247.
 48. Thomas SJ, Ray BK, Ensor KB. Understanding the volatility and liquidity of corporate bonds through multivariate time series of counts with zero-inflation. Technical Report, Rice University, Department of Statistics, 2013.
 49. Nakatani T. *ccgarch: An R Package for modeling multivariate GARCH models with conditional correlations*, r package version 0.2.1. Available at: <http://cran.r-project.org/web/packages/ccgarch/> edition, 2013.
 50. Ghalanos A. *rmgarch: Multivariate GARCH models*. r package version 1.2-0, 2013. Available at: <http://cran.r-project.org/web/packages/rmgarch/> edition.
 51. SAS Institute Inc.. *Sas Onlinedoc, Version 8*. Cary, NC: SAS Institute Inc.; 2000.
 52. Stata Corp. *Stata Statistical Software: Release 13*. StataCorp LP: College Station, TX; 2013.