

A general multivariate threshold GARCH model with dynamic conditional correlations

Francesco Audrino^{a,*} and Fabio Trojani^{a,b}

^aDepartment of Economics, University of St. Gallen

^bSwiss Institute of Banking and Finance, University of St. Gallen

Revised: September 2008

Abstract

We introduce a new multivariate GARCH model with multivariate thresholds in conditional correlations, and develop a two-step estimation procedure that is feasible in large dimensional applications. Optimal threshold functions are estimated endogenously from the data, and the model conditional covariance matrix is ensured to be positive definite. We study the empirical performance of our model in three applications using US stock and bond market data. The conditional volatility of stock returns exhibits GARCH and threshold structures, but the conditional correlation dynamics depend on piecewise constant thresholds only. Similar findings arise for the conditional correlations of government and corporate bond returns. We estimate both threshold and GARCH structures in the conditional correlations of stock and government bond returns. In all applications our model has, in terms of statistical and economic significance, higher forecasting power for conditional correlations compared with several other multivariate GARCH models.

Keywords: Multivariate GARCH models; Dynamic conditional correlations; Tree-structured GARCH models.

JEL codes: C12, C13, C51, C53, C61

*Corresponding address: Institute of Mathematics and Statistics, University of St. Gallen, Bodanstrasse 6, CH-9000 St. Gallen, e-mail: francesco.audrino@unisg.ch, Phone: 0041 71 2242431. E-mail for Fabio Trojani: fabio.trojani@unisg.ch.

1 Introduction

In this paper, we present a new multivariate GARCH model with Dynamic Conditional Correlations (DCC) that extends previous approaches by admitting multivariate thresholds in the conditional volatilities and correlations of multivariate time series. This extension allows us to account for rich asymmetric effects and dependencies of conditional volatilities and correlations, as they are often encountered - for instance - in financial real data applications. Similarly to the classical Engle (2002) DCC-model, our model estimation is numerically feasible in large dimensions. Moreover, the positive definiteness of the conditional covariance matrix is ensured in a natural way by the structure of the model. Finally, thresholds in volatilities and correlations of our model are not fixed ex ante, but are estimated from the data, together with all other parameters in the model.

To define the threshold function in our model, we extend the tree-structured state space partition in Audrino and Bühlmann (2001) to a setting with multivariate thresholds in volatilities and correlations. As shown in Audrino and Trojani (2006) and Audrino (2006), the tree-structured threshold construction can incorporate a potentially large number of multivariate regimes in univariate settings in a parsimonious way. In this paper, we study a multivariate model with a potentially large number of tree-structured thresholds in volatilities and correlations, and we develop a feasible estimation strategy that can be applied to estimate the model in large dimensional applications as well. The threshold construction is obtained using a binary tree, in which each terminal node defines a local GARCH-type dynamics for volatilities and correlations over a partition cell of the multivariate state space. The estimation is performed by a simple two-step procedure that estimates the number and structure of the underlying thresholds together with the parameters of the local GARCH dynamics for volatilities and correlations. The optimal threshold structure is identified by solving a high dimensional model selection problem based on the Schwarz Bayesian Information Criterion (BIC).

We estimate our model in three distinct applications to US stock and bond market data and focus on the explanatory power for future conditional correlations, comparing the results with a set of competing models in the literature: Engle (2002) DCC-model, Ledoit et al. (2003) flexible multivariate GARCH model and Pelletier (2006) Regime Switching Dynamic Correlations (RSDC) model. Like our model, the DCC and the RSDC models can be estimated by a two step procedure that separates the estimation of the conditional volatility and correlation dynamics. In order to measure, where possible, the additional forecasting power for correlations, we es-

timate these models using a set of univariate tree structured GARCH-dynamics for volatility identical to those in our model. The flexible multivariate GARCH model cannot be estimated by a two-step estimation procedure. Therefore, the estimated volatility dynamics are different from those of our model. Our model differs from the others in its manner of specifying the correlation dynamics. The DCC and the flexible multivariate GARCH models are single-regime models for correlations, and the RSDC model specifies a very simple regime structure for conditional correlations. By contrast, our setting can account parsimoniously for GARCH-type dynamics and multivariate conditional correlation thresholds without fixing ex ante the structure and the number of thresholds.

Using our tree-structured GARCH-DCC model, we empirically study the relative importance of GARCH and threshold effects in the conditional correlation dynamics of US stock and bond returns. The conditional volatility functions of US stock returns exhibit GARCH and threshold features, but conditional correlations depend on a piecewise constant threshold function. Similar findings arise for the estimated correlation process of US government and corporate bond returns. We find both threshold and GARCH-effects for the correlation process of stock and government bond returns, with estimated thresholds that are functions of lagged stock and bond returns. In all applications, the estimated tree-structured partition of the state space improves the forecasting power for conditional correlations relative to the other multivariate GARCH models. Out-of-sample improvements are in most cases statistically and economically significant, based on different criteria recently proposed in the literature, including the test for superior predictive ability (SPA) introduced by Hansen (2005), the model confidence set approach by Hansen et al. (2003), and the economic measure of volatility and correlation timing ability for portfolio allocation in Engle and Colacito (2006) and Bandi et al. (2008). These findings highlight the importance of flexible multivariate threshold and GARCH structures for forecasting the conditional correlation of stock and bond markets.

Section 2 presents our tree-structured GARCH-DCC model and Section 3 describes the two-step estimation procedure that can be applied to estimate it. Section 4 presents our empirical study of the conditional correlation dynamics of US stock and bond returns. Section 5 summarizes the main results and concludes.

2 The model

We consider a multivariate stochastic process $(\mathbf{X}_t)_{t \in \mathbb{Z}}$ with values in $\mathbb{R}^d$:

$$\mathbf{X}_t = D_t \epsilon_t, \quad (2.1)$$

where $D_t := \text{diag}[\sigma_{1,t}, \dots, \sigma_{d,t}]$ and $\sigma_{i,t}$ is the conditional standard deviation of the i -th component of $\mathbf{X}_t$ at time $t - 1$. $(\epsilon_t)_{t \in \mathbb{Z}}$ is a zero-mean process in $\mathbb{R}^d$ with components having a unit conditional standard deviation by construction.¹ The conditional covariance matrix of ϵ_t at time $t - 1$ is denoted by R_t . Therefore, we obtain the following standard factorization of the conditional covariance matrix of $\mathbf{X}_t$:

$$\text{Cov}_{t-1}(\mathbf{X}_t) = D_t R_t D_t, \quad (2.2)$$

Our tree-structured DCC-GARCH model parameterizes the conditional volatility matrix D_t and the conditional correlation matrix R_t by means of two parametric threshold functions. Each diagonal element of D_t is modeled as a univariate tree-structured threshold GARCH(1,1)-model, as in Audrino and Bühlmann (2001) and Audrino and Trojani (2006). The conditional correlation matrix R_t is modeled according to a threshold DCC-type model described in more detail below.

2.1 Tree-structured model for D_t

Let $X_{t,j}$ be the j -th component of $\mathbf{X}_t$. In principle, the thresholds in the volatility dynamics of $X_{t,j}$ may depend on all components of $\mathbf{X}_{t-1}$. For simplicity of exposition, let us assume that they are functions of $(X_{t-1,1}, X_{t-1,j})$. Let $\mathcal{P}_j = \{\mathcal{R}_{1,j}, \dots, \mathcal{R}_{k_j,j}\}$ be a partition of the state space $G := \mathbb{R}^2 \times \mathbb{R}^+$ of $(X_{t-1,1}, X_{t-1,j}, \sigma_{t-1,j}^2)$:

$$\mathcal{P}_j = \{\mathcal{R}_{1,j}, \dots, \mathcal{R}_{k_j,j}\}, \cup_{s=1}^{k_j} \mathcal{R}_{s,j} = G, \mathcal{R}_{i,j} \cap \mathcal{R}_{s,j} = \emptyset \ (i \neq s).$$

Given a partition cell $\mathcal{R}_{i,j}$, we specify the local conditional variance dynamics of $X_{t,j}$ on $\mathcal{R}_{i,j}$ as a GARCH(1,1) model. Therefore, threshold function $\sigma_{t,j}^2$ takes the form

$$\sigma_{t,j}^2 = \sum_{i=1}^{k_j} (\alpha_{ij} + \beta_{ij} X_{t-1,j}^2 + \gamma_{ij} \sigma_{t-1,j}^2) I_{[(X_{t-1,1}, X_{t-1,j}, \sigma_{t-1,j}^2) \in \mathcal{R}_{i,j}]}, \quad (2.3)$$

where $I_{[\cdot]}$ is the indicator function and $\theta_{1,j}$ is the parameter vector:

$$\theta_{1,j} = \{\alpha_{ij}, \beta_{ij}, \gamma_{ij} ; i = 1, \dots, k_j\}.$$

To completely specify the conditional variance function (2.3), we have to define the class of partitions $\mathcal{P}_j$ that are admissible in our tree-structured model. Essentially, the only restriction we impose is that $\mathcal{P}_j$ is composed of rectangular cells $\mathcal{R}_{i,j}$, $i = 1, \dots, k_j$, delimited by a set of multivariate thresholds for $(X_{t-1,1}, X_{t-1,j}, \sigma_{t-1,j}^2)$. For example, in a model with three regimes and two thresholds, the partitioning cells $\mathcal{R}_{1,j}, \mathcal{R}_{2,j}, \mathcal{R}_{3,j}$ could be of the form:

$$\begin{aligned}\mathcal{R}_{1,j} &= \{X_{t-1,j} \leq d_1\}, \\ \mathcal{R}_{2,j} &= \{X_{t-1,j} > d_1 \text{ and } X_{t-1,1} \leq d_2\}, \\ \mathcal{R}_{3,j} &= \{X_{t-1,j} > d_1 \text{ and } X_{t-1,1} > d_2\},\end{aligned}$$

where the parameters d_1, d_2 define the two multivariate thresholds in the model. In this case, $\mathcal{R}_{1,j}$ is associated with a regime of low conditioning values $X_{t-1,j}$. $\mathcal{R}_{2,j}$ corresponds to a regime with higher conditioning values of $X_{t-1,j}$, but low values of $X_{t-1,1}$. Finally, $\mathcal{R}_{3,j}$ implies a regime in which both conditioning values $X_{t-1,1}$ and $X_{t-1,j}$ are large. The multivariate threshold function in our model is defined by means of a binary tree $\mathcal{T}_j$ in which every terminal node represents a particular cell $\mathcal{R}_{i,j}$. Details of the construction and the interpretation of binary trees for our applied examples are provided in Audrino and Trojani (2006). For each component $X_{t,j}$, estimation of (2.3) is achieved by a high dimensional model selection problem that determines the optimal number and the structure of the relevant thresholds (and hence the partition cells) in $\mathcal{P}_j$. Details of this estimation procedure for univariate tree structured GARCH(1,1) models are given in Audrino and Bühlmann (2001) and Audrino and Trojani (2006), Section 2.3.

2.2 Tree structured model for R_t

Given $\theta_1 = \{\theta_{1,j} : j = 1, \dots, d\}$, let:

$$\epsilon_t = D_t(\theta_1)^{-1} \mathbf{X}_t,$$

and:

$$R_t = \text{Corr}_{t-1}(\mathbf{X}_t) = \text{Cov}_{t-1}(\epsilon_t).$$

We model R_t by means of a tree-structured model, in which conditional correlations satisfy an Engle (2002)-type local DCC model across several multivariate thresholds. In order to keep the model tractable, we assume that thresholds in the R_t dynamics depend on ϵ_{t-1} only via the

average

$$\bar{\rho}_{t-1} = \frac{1}{d(d-1)} \sum_{u \neq v} \epsilon_{t-1,u} \epsilon_{t-1,v}$$

of the cross products of the component of ϵ_{t-1} . Intuitively, this choice allows us to account for asymmetric effects in conditional correlations, as a function of particular lagged process realizations $\mathbf{X}_{t-1}$ and specific movements in average lagged conditional correlation shocks $\bar{\rho}_{t-1}$.

To define the parametric threshold function R_t in our model, let $\mathcal{P} = \{\tilde{\mathcal{R}}_1, \dots, \tilde{\mathcal{R}}_w\}$ be a partition of the state space $\tilde{G} := \mathbb{R}^{d+1}$ of $(\mathbf{X}_{t-1}, \bar{\rho}_{t-1})$. We consider the following family of functional forms for R_t :

$$R_t = \sum_{i=1}^w c_i R_{it} I_{\{(\mathbf{X}_{t-1}, \bar{\rho}_{t-1}) \in \tilde{\mathcal{R}}_i\}} + \left(1 - \sum_{i=1}^w c_i I_{[(\mathbf{X}_{t-1}, \bar{\rho}_{t-1}) \in \tilde{\mathcal{R}}_i]}\right) Id_n \quad (2.4)$$

where $c_1, \dots, c_n \in [0, 1]$, Id_n is the d -dimensional identity matrix and the parametric processes for R_{it} , $i = 1, \dots, n$, are given by:

$$R_{it} = \text{diag}[Q_{it}]^{-1/2} Q_{it} \text{diag}[Q_{it}]^{-1/2} \quad (2.5)$$

with

$$Q_{it} = (1 - \phi_i - \lambda_i) \bar{Q} + \phi_i \epsilon_{t-1} \epsilon'_{t-1} + \lambda_i Q_{it-1}, \quad (2.6)$$

parameters $\phi_i, \lambda_i \geq 0$ such that $\phi_i + \lambda_i < 1$ for all $i = 1, \dots, w$, and $\bar{Q}$ is, as in the classical Engle (2002) DCC model, the unconditional covariance matrix of the residuals ϵ_t . Given a fixed partition $\mathcal{P}$, the parameter vector

$$\theta_2 = \{c_i, \phi_i, \lambda_i, \text{vech}(\bar{Q}) ; i = 1, \dots, w\}. \quad (2.7)$$

completely parameterizes the threshold function defining the conditional correlation function (2.4).

Since for any $i = 1, \dots, w$, the local model for Q_{it} satisfies an Engle (2002) DCC-type dynamics, positive definiteness of the resulting threshold model for R_t is easily implied by the model structure under the above conditions on the model parameters. When $\mathcal{P} = \{\tilde{G}\}$, i.e. the partition is trivial, we obtain the Engle (2002) DCC model by setting $c_1 = \dots = c_w = 1$. Therefore, this model is nested in our model. Moreover, by setting $\phi_i = \lambda_i = 0$ for $i = 1, \dots, w$, we can write R_t as:

$$R_t = \sum_{i=1}^w c_i \bar{R} I_{[(\mathbf{X}_{t-1}, \bar{\rho}_{t-1}) \in \tilde{\mathcal{R}}_i]} + \left(1 - \sum_{i=1}^w c_i I_{[(\mathbf{X}_{t-1}, \bar{\rho}_{t-1}) \in \tilde{\mathcal{R}}_i]}\right) Id_n. \quad (2.8)$$

where $\bar{R}$ is a fixed d -dimensional correlation matrix.² In this case, we obtain a piecewise constant correlation matrix defined by a multivariate threshold function over the partition $\mathcal{P}$. In contrast to the RSDC model in Pelletier (2006), this particular subcase of our model can account for a flexible description of multiple multivariate regimes in correlations, because the number and the structure of the regimes in the estimated model does not have to be fixed from the beginning. Finally, when $\phi_i > 0$ or $\lambda_i > 0$ for $i = 1, \dots, w$ and $\mathcal{P}$ is not a trivial partition, by setting $c_1 = \dots = c_w = 1$ we obtain a tree-structured DCC-model locally satisfying Engle's DCC-dynamics over the distinct partitioning cells $\tilde{\mathcal{R}}_i$.

As for the univariate tree-structured volatility dynamics of the last section, we need to define the class of admissible partitions $\mathcal{P}$ for our correlation function. Again, the only restriction we put on $\mathcal{P}$ is that it is composed of rectangular partition cells. Consistent with our assumptions, these partition cells are delimited by a set of multivariate thresholds for $(\mathbf{X}_{t-1}, \bar{\rho}_{t-1})$. In order to construct such rectangular partition cells, we make use of a binary tree, in which every terminal node represents a cell $\tilde{\mathcal{R}}_i$. Estimation of the threshold function in the correlation dynamics (2.4) is achieved by a high dimensional model selection procedure that determines the optimal number and the structure of the relevant thresholds in the underlying partition. This model selection scheme is not computationally feasible if applied directly to the multivariate time series $(\epsilon_t \epsilon'_t)_{t \in \mathbb{Z}}$. A natural way to reduce estimation complexity is to remark that the partition $\mathcal{P}$ is identical to the one implied by a corresponding tree-structured univariate model for the time series $(\bar{\rho}_t)_{t \in \mathbb{Z}}$. Indeed, since $R_t = E_{t-1}(\epsilon_t \epsilon'_t)$ it follows:

$$\begin{aligned} E_{t-1}(\bar{\rho}_t) &= \frac{1}{d(d-1)} \sum_{u \neq v} R_t^{uv} \\ &= \sum_{i=1}^w c_i \left(\frac{1}{d(d-1)} \sum_{u \neq v} R_t^{uv} \right) I_{[(\mathbf{X}_{t-1}, \bar{\rho}_{t-1}) \in \tilde{\mathcal{R}}_i]} \end{aligned} \quad (2.9)$$

where R_t^{uv} (R_{it}^{uv}) denotes the uv -th component of the matrix R_t (R_{it}). Therefore, the tree-structured model

$$\bar{\rho}_t = E_{t-1}(\bar{\rho}_t) + \eta_t, \quad (2.10)$$

where $(\eta_t)_{t \in \mathbb{Z}}$ is a martingale difference process and $E_{t-1}(\bar{\rho}_t)$ is given by equation (2.9), defines a univariate tree-structured process for $\bar{\rho}_t$ based on the same partition $\mathcal{P}$ as in the correlation dynamics (2.4). It follows, that we can exploit the univariate model (2.10) to estimate the threshold structure in equation (2.4). In particular, we can develop a model selection procedure

for selecting the optimal threshold structure in the correlation dynamics. The most simple dynamics arise in the piecewise constant case:

$$E_{t-1}(\bar{\rho}_t) = \sum_{i=1}^w c_i \left(\frac{1}{d(d-1)} \sum_{u \neq v} \bar{R}^{uv} \right) I_{[(\mathbf{X}_{t-1}, \bar{\rho}_{t-1}) \in \tilde{\mathcal{R}}_i]} \quad (2.11)$$

where $\bar{R}^{uv}$ is the uv -component of the correlation matrix $\bar{R}$ in the piecewise-constant dynamics (2.8). This piecewise constant function is the optimal one that has been estimated in our applications to the US equity market in Section 4.1 and the US bond market in Section 4.3. More generally, for $c_1 = \dots = c_n = 1$ and $\lambda_i, \phi_i > 0$, $i = 1, \dots, w$, we can also encompass the univariate dynamics of $\bar{\rho}_t$ that are consistent with a tree-structured DCC model of the form (2.4) for correlations. This threshold structure is the one we estimate in our application of Section 4.1, where we model the correlation between Treasury bond and stock returns. Model selection across this class of potential threshold functions for $E_{t-1}[\bar{\rho}_t]$ is performed using the BIC-information criterion. Once the partition $\mathcal{P}$ in (2.10) has been estimated, the parameter (2.7) of the multivariate correlation dynamics can be estimated using a multivariate conditional pseudo likelihood for ϵ_t , in which the selected partition $\mathcal{P}$ is held fixed. The next section provides additional details on the estimation procedure used to estimate our tree structured DCC model.

3 Estimation of the tree-structured DCC model

Estimation of our tree-structured model is accomplished in two steps. In the first step, an estimate of the volatility process D_t is obtained by performing d estimations of the univariate tree-structured conditional volatility dynamics $\sigma_{t,1}(\theta_{1,1}), \dots, \sigma_{t,d}(\theta_{1,d})$ implied by the specification (2.3). The resulting point estimate $\hat{D}_t := D_t(\hat{\theta}_1)$ is used to compute the estimated scaled residuals

$$\hat{\epsilon}_t := \hat{D}_t^{-1} \mathbf{X}_t. \quad (3.12)$$

The scaled residuals $\hat{\epsilon}_t$ are used in the second step of our procedure to estimate the tree-structured conditional correlation dynamics (2.4).

3.1 Estimation of tree-structured univariate GARCH-dynamics

Estimation of the d tree-structured univariate volatility functions (2.3) is achieved by a high-dimensional model selection problem, that determines the optimal structure of the relevant thresholds in any partition $\mathcal{P}_j$ of the univariate volatility dynamics (2.3), $j = 1, \dots, d$.

In a first step, a largest univariate tree-structured GARCH model is estimated for any $j = 1, \dots, d$, given a fixed maximal number M_j of possible thresholds in (2.3). This first step delivers a maximal possible partition $\mathcal{P}_j^{max}$ of the relevant state space G in the univariate volatility dynamics (2.3).

In a second step, a tree-structured model selection procedure for non-nested models is applied, that selects the optimal subpartition $\mathcal{P}_j \subset \mathcal{P}_j^{max}$ out of the maximal one. Model selection is performed according to the BIC information criterion implied by a conditionally Gaussian log likelihood for any process coordinate $X_{t,j}$, $j = 1, \dots, d$. The resulting optimal tree-structured volatility model minimizes the BIC information criterion across all tree-structured sub-partitions of $\mathcal{P}_j^{max}$. The complete algorithm used to estimate univariate tree-structured GARCH(1,1) models is given in Audrino and Bühlmann (2001) and Audrino and Trojani (2006).

The construction of the largest partition $\mathcal{P}_j^{max}$ proceeds as follows. We first fix a maximal number $M_j + 1$ of partition cells in the tree. Because of the tree-structured construction of $\mathcal{P}_j^{max}$, this first step implies a maximal number $M_j + 1$ of conditional volatility regimes (i.e., the number of terminal nodes in the binary tree). A parsimonious specification of the maximal number M_j of thresholds ensures a statistically and computationally tractable model dimension. Moreover, it avoids (over) fitting an overly flexible model dynamics, which would result in a poor out-of-sample forecasting power. For any coordinate axis of the multivariate state space that has to be split we search for multivariate thresholds over grid points that are empirical α -quantiles of the data along the relevant coordinate axis. We fix the empirical quantiles as $\alpha = i/\text{mesh}$, $i = 1, \dots, \text{mesh} - 1$, where mesh determines the fineness of the grid on which we search for multivariate thresholds. Typically, we choose mesh = 8. The partition of the state space $G = \mathbb{R}^d \times \mathbb{R}^+$ into a maximal number of $M_j + 1$ cells is performed as follows. A first threshold $d_1 \in \mathbb{R}$ or $\mathbb{R}^+$ in one coordinate indexed by a component index $\iota_1 \in \{1, 2, \dots, d + 1\}$ partitions G as

$$G = \mathcal{R}_{left} \cup \mathcal{R}_{right},$$

where $\mathcal{R}_{left} = \{(\mathbf{X}_{t-1}, \sigma_{t-1}^2) \in \mathbb{R}^d \times \mathbb{R}^+; (\mathbf{X}_{t-1}, \sigma_{t-1}^2)_{\iota_1} \leq d_1\}$ and $(\mathbf{X}_{t-1}, \sigma_{t-1}^2)_{\iota_1}$ denotes the ι_1 -component of the tuple $(\mathbf{X}_{t-1}, \sigma_{t-1}^2)$. $\mathcal{R}_{right}$ is defined analogously using the relation ' $>$ ' instead of ' $\leq$ '. In a second step, one of the partition cells $\mathcal{R}_{left}$, $\mathcal{R}_{right}$ is further partitioned with a second threshold d_2 and a second component index ι_2 , in the same way as above. We then iterate this procedure. For the m -th iteration step, we specify a new pair (d_m, ι_m) , which defines a new threshold d_m for the coordinate indexed by ι_m , and an existing partition cell that

is going to be split into two sub-cells. For a new pair $(d, \iota) \in \mathbb{R} \times \{1, \dots, d+1\}$ refinement of an existing partition $\mathcal{P}^{(old)}$ is obtained by picking $\mathcal{R}_{j^*} \in \mathcal{P}^{(old)}$ and splitting it as

$$\mathcal{R}_{j^*} = \mathcal{R}_{j^*,left} \cup \mathcal{R}_{j^*,right} . \quad (3.13)$$

This procedure produces a new (finer) partition of G , given by

$$\mathcal{P}^{(new)} = \{\mathcal{R}_j, \mathcal{R}_{j^*,left}, \mathcal{R}_{j^*,right}, j \neq j^*\}. \quad (3.14)$$

In this partition, the tuple (d, ι) describes a threshold d and a component index ι such that $\mathcal{R}_{j^*,left} = \{(\mathbf{X}_{t-1}, \sigma_{t-1}^2) \in \mathcal{R}_{j^*}; (\mathbf{X}_{t-1}, \sigma_{t-1}^2)_\iota \leq d\}$. $\mathcal{R}_{j^*,right}$ is defined analogously, with the relation ' $>$ ' instead of ' $\leq$ '. The whole procedure finally determines a partition $\mathcal{P}_j^{max} = \{\mathcal{R}_{1,j}, \dots, \mathcal{R}_{M_j+1,j}\}$. This partition can be represented and summarized by a binary tree in which every terminal node represents a partition cell of $\mathcal{P}_j^{max}$. To select the specific threshold and component index (d, ι) in each iteration step of the above procedure we optimize the corresponding conditional negative (pseudo) log-likelihood in the model.

3.2 Estimation of tree-structured DCC-dynamics

In the first step, we estimate the optimal partition $\mathcal{P}$ using the tree-structured model (2.10) for $\bar{\rho}_t$ and the scaled estimated residuals $\hat{\epsilon}_t$. In the second step, we fix the partition $\hat{\mathcal{P}}$ - say - estimated for the univariate model (2.10), and estimate the parameter θ_2 in (2.7) by a multivariate pseudo maximum likelihood estimator.

(i) *Estimation of the univariate tree structured model (2.10).* Let

$$\hat{\bar{\rho}}_t = \sum_{u \neq v} \hat{\epsilon}_{t-1,u} \hat{\epsilon}_{t-1,v} / [d(d-1)]. \quad (3.15)$$

The following tree-structured model for $\hat{\bar{\rho}}_t$ is estimated; compare with equation (2.10):

$$\hat{\bar{\rho}}_t = E_{t-1}(\hat{\bar{\rho}}_t) + \eta_t, \quad (3.16)$$

where $(\eta_t)_{t \in \mathbb{Z}}$ is a martingale difference sequence and

$$E_{t-1}(\hat{\bar{\rho}}_t) = \sum_{i=1}^w c_i \left(\frac{1}{d(d-1)} \sum_{u \neq v} \hat{R}_{it}^{uv} \right) I_{[(\mathbf{X}_{t-1}, \hat{\bar{\rho}}_{t-1}) \in \tilde{\mathcal{R}}_i]} . \quad (3.17)$$

In this equation, $\hat{R}_{it}$ denotes for $i = 1, \dots, n$ a constant correlation matrix when the tree-structured model for correlations implies a piecewise constant correlation matrix. It then follows

in this case that the conditional mean of $\widehat{\rho}_t$ is simply a piecewise constant threshold function. More generally, for $c_1 = \dots = c_n = 1$ and $\phi_i > 0$ or $\lambda_i > 0$, where $i = 1, \dots, w$, the local conditional correlation matrix $\hat{R}_{it}$ is simply defined in the same way as R_{it} in equation (2.5) and (2.6), but with $\hat{\epsilon}_{t-1}$ replacing ϵ_{t-1} in equation (2.6).

We apply to the series $\widehat{\rho}_t$ the same estimation procedure given in the last section for individual conditional variances. First, we estimate a largest univariate tree-structured model for $\widehat{\rho}_t$, given a fixed maximal number M of possible thresholds in (2.10). In all our empirical applications, we fix the maximal number of candidate thresholds in model (2.10) at $M = 4$. A tree-structured model selection procedure for non nested models is then applied, which selects the optimal subpartition $\mathcal{P}$ out of the maximal one. Model selection is performed according to the BIC criterion implied by a conditionally Gaussian pseudo log likelihood for $\widehat{\rho}_t$.³ In our empirical study, we found that this procedure offers a simple and effective way to reduce the computational costs implied by the estimation of our multivariate tree-structured model. In particular, in the applications of Section 4 a piecewise constant conditional correlation function is estimated for equity returns and different types of bond returns. However, local DCC-type structures are found to better model the conditional correlations between equity and bond returns.

(ii) *Estimation of the tree-structured conditional correlation function R_t .* In the second step of our estimation procedure, we fix the partition $\widehat{\mathcal{P}}$ estimated in step (i) and we estimate the parameter vector θ_2 in (2.7) by a pseudo maximum likelihood estimator $\widehat{\theta}_2$ for θ_2 , under a Gaussian multivariate conditional pseudo likelihood for $\widehat{\epsilon}_t$. If in step (i) the optimal threshold function does not imply piecewise constant correlations, we estimate the matrix $\overline{Q}$ in the dynamics (2.6) by doing correlation targeting, as proposed by Engle and Sheppard (2001) and Pelletier (2006). If in step (i) a piecewise constant correlation structure has been selected, in the second step we estimate a piecewise constant correlation process of the form (2.8). In such a case, we estimate the constant matrix $\overline{R}$ by doing correlation targeting in a rolling window of one year of data. The piecewise constant correlation structure significantly reduces the number of parameters over which the likelihood function has to be maximized.

3.3 Consistency

Proofs of the consistency of our model selection procedure are very difficult to obtain for the case where the true model is in the class of tree-structured GARCH-DCC models. Analogously to CART, it is possible to prove theorems that study the behavior of the prevailing parameter

estimators when growing the tree. However, such results do not imply model selection consistency. Furthermore, it is quite unlikely that the “correct” generating process in our example and other similar ones with real data is indeed exactly a tree-structured model for volatilities and correlations. For this reason, it is more important to prove consistency of the estimators for the parameters of a tree-structured model under a model misspecification than it is to prove consistency of the model selection strategy under the assumption of a correctly-specified, tree-structured model. Such consistency results can be found in Audrino and Bühlmann (2001). Based on these results, consistency of the two-step estimates $(\hat{\theta}_1, \hat{\theta}_2)$ in the tree structured DCC-GARCH model under a possible model misspecification can be derived in the standard way under mild regularity conditions; see, for instance, Newey and McFadden (1994). Moreover, efficient estimates can be obtained by performing a further one step Newton-Raphson estimation of the full likelihood, using as starting values the parameter estimates obtained from the two-step procedure (see, e.g., Pagan, 1986).

4 Results

In this section, we test the in-sample and out-of-sample explanatory power of our tree-structured GARCH-DCC (TreeDCC) model in three different applications involving the econometric analysis of US stock and bond returns. We compare our model with several multivariate GARCH models that have been recently proposed in the literature. Some of these models are nested in ours and can be estimated by a two-step estimation procedure:

- The CCC-GARCH model, as proposed by Bollerslev (1990); this model is nested in our model.
- The DCC-GARCH model, as proposed by Engle (2002); this model is nested in our model.
- The RSDC-GARCH model with switching regimes in conditional correlations, as proposed in Pelletier (2006). This model is not formally nested in ours.

Since the individual volatility processes are estimated separately from the correlation dynamics in these models, we can easily focus in our empirical study on the additional explanatory for conditional correlations, which is the main topic of this paper. To achieve this goal, we estimate identical volatility processes as in our tree-structured DCC-GARCH setting. The volatility processes are all specified as univariate tree-structured GARCH(1,1) processes. We also study

the performance of our model relative to the flexible multivariate GARCH setting in Ledoit et al. (2003), which does not include thresholds in volatilities or correlations and is based on a more general correlation dynamics than the one implied by the Engle (2002) DCC model. Therefore, this model is not nested in our setting. Flexible multivariate GARCH models have been shown by Ledoit et al. (2003) to describe the dynamics of stock returns quite accurately. Therefore, they are further natural competitors of our approach, especially in applications that study the multivariate dynamics of stock markets, as in our first empirical example.

To quantify and compare the in-sample and out-of-sample fit of the different models, we compute several goodness-of-fit statistics for conditional covariances. Since the individual volatility processes are identical for all but the flexible multivariate GARCH model, this comparison allows us to investigate the additional explanatory power of our model for explaining the correlation dynamics. We consider the following goodness of fit measures:

- The multivariate negative log-likelihood statistic (NL),
- The multivariate version of the classical mean absolute error statistic (MAE),
- The multivariate version of the classical mean squared error statistic (MSE),

The last two performance measures require the specification of sensible values for the unknown true conditional covariance matrix. A powerful way of computing good proxies for this matrix is by means of the the so-called realized covariance approach, which is the natural multivariate version of the realized volatility approach proposed, among others, in Andersen et al. (2001, 2003) and Barndorff-Nielsen and Shephard (2001, 2002a). We follow this approach in our first two real data applications, in which we collect tick-by-tick return data to compute the realized covariance between returns using the methodology proposed in Corsi and Audrino (2007). In the third and last application, we do not have tick-by-tick data at our disposal to compute the realized covariances between returns. Therefore, we use products of centered returns as (noisy) proxies for the unknown covariances between returns.⁴

The different statistics used to quantify the in-sample and out-of-sample goodness of fit in our empirical analysis are defined as follows: (IS denotes in-sample and OS denotes out-of-sample):

$$\text{IS-NL: } -\log\text{-likelihood}(\mathbf{X}_1^n; \hat{\phi})$$

$$\text{OS-NL: } -\log\text{-likelihood}(\mathbf{Y}_1^{n_{\text{out}}}; \hat{\phi})$$

$$\text{IS-MAE: } \frac{1}{d^2} \sum_{i,j=1}^d \frac{1}{n} \sum_{t=1}^n |v_{t,ij} - \hat{v}_{t,ij}|$$

$$\text{OS-MAE: } \frac{1}{d^2} \sum_{i,j=1}^d \frac{1}{n_{\text{out}}} \sum_{t=1}^{n_{\text{out}}} |v_{t,ij} - \hat{v}_{t,ij}(\mathbf{Y}_1^{t-1})|$$

$$\text{IS-MSE: } \frac{1}{d^2} \sum_{i,j=1}^d \frac{1}{n} \sum_{t=1}^n |v_{t,ij} - \hat{v}_{t,ij}|^2$$

$$\text{OS-MSE: } \frac{1}{d^2} \sum_{i,j=1}^d \frac{1}{n_{\text{out}}} \sum_{t=1}^{n_{\text{out}}} |v_{t,ij} - \hat{v}_{t,ij}(\mathbf{Y}_1^{t-1})|^2$$

where in the OS performance measures the expression $\hat{v}_{t,ij}(\mathbf{Y}_1^{t-1})$ is the ij -th covariance prediction implied by our out-of-sample data $\mathbf{Y}_1^{n_{\text{out}}} = \{\mathbf{Y}_1, \dots, \mathbf{Y}_{n_{\text{out}}}\}$ at time t under the parameter estimates obtained from the in-sample data $\mathbf{X}_1^n = \{\mathbf{X}_1, \dots, \mathbf{X}_n\}$. $v_{t,ij}$ is the realized covariance between the return series i and j at time t , in our first two applications, or the product of the centered returns of series i and j at time t , in our last application. In all cases, a lower goodness of fit measure indicates a higher forecasting power of a model for conditional correlations.

To evaluate whether differences in performance among the models considered are statistically and/or economically significant, we perform a series of tests recently proposed in the literature. We investigate the statistical relevance of the improvements in forecasting accuracy of the TreeDCC model relative to the other approaches:

- We perform the superior predictive ability (SPA) test introduced by Hansen (2005). This allows us to verify whether each of the models considered is significantly outperformed by one (or more) of the alternatives.
- We construct a model confidence set (MCS) at the 5% and 10% confidence levels as proposed by Hansen et al. (2003). We introduce the MCS approach with the goal of characterizing the best subset of models out of a set of competing ones; see Appendix A for more details about the MCS construction.

We also evaluate the alternative model specifications using an economic criterion. We apply the methodology suggested by West et al. (1993) and Fleming et al. (2001, 2003) to quantify the

economic benefit of different correlation forecasts in the context of a portfolio strategy based on volatility timing. Similarly to Engle and Colacito (2006) and Bandi et al. (2008), we employ the variance component of an investor's long-run mean-variance utility as a metric to quantify the economic differences between the alternative correlation (covariance) forecasts:

$$AU = \frac{\lambda}{2} \frac{1}{n_{\text{out}}} \sum_{t=1}^{n_{\text{out}}} (R_t^p - \bar{R}^p)^2,$$

where the portfolio return at time t is given by

$$R_t^p = R^f + \mathbf{w}_{t-1}'(\mathbf{Y}_t - R^f \mathbf{I}_d), \quad t = 1, \dots, n_{\text{out}}.$$

$\bar{R}^p$ is the sample mean of the portfolio returns across the out-of-sample period and λ is a coefficient of risk aversion. As in Bandi et al. (2008), we use three values of $\lambda = 2, 7, 10$. In the computation of the portfolio returns, R^f is the risk-free rate, $\mathbf{I}_d$ is a $d \times 1$ unit vector, and $\mathbf{w}_{t-1}$ is a d -vector of portfolio weights obtained by solving the classical mean-variance optimization problem at time $t - 1$, given a fixed target return on the portfolio. In our real data applications, we set R^f equal to the average value of the US three-month rate in the out-of-sample period. In the mean-variance optimization, we use the one-step-ahead conditional covariance forecasts $\hat{v}_{t,ij}(\mathbf{Y}_1^{t-1})$ obtained from the alternative multivariate GARCH models under investigation. We can interpret the difference between the quantity AU implied by the TreeDCC model and each other model as the fee that an investor would be willing to pay to switch from correlation forecasts generated by each alternative model to those of the TreeDCC model. We also use the generalization of Diebold-Mariano test for pairwise equal predictive ability (EPA) and the joint test introduced in Engle and Colacito (2006) to study the statistical significance of the differences in estimated economic gains. For details, see again Engle and Colacito (2006) and Bandi et al. (2008).

4.1 First real data application: US equity returns

We consider a multivariate time series of (annualized) daily log-returns for ten US stocks: Alcoa, Citigroup, Hasbro, Harley Davidson, Intel, Microsoft, Nike, Pfizer, Tektronix and Exxon. Data are for the sample period between January 2, 2001 and December 30, 2005, amounting to 1256 trading days. The source of the data is *Tick Data*, a division of Nexa Technologies, Inc. (see the webpage <http://www.tickdata.com>). Using these tick-by-tick data, we construct realized covariances with the method in Corsi and Audrino (2007) and obtain the quantities $v_{t,ij}$ needed to compute our goodness-of-fit measures.

We split the sample into two subperiods. The first one consists of $n = 752$ trading days, from January 2, 2001 to December 31, 2003. Data from this subperiod are used for in-sample estimation and performance evaluation. The second subperiod consists of the remaining $n_{\text{out}} = 504$ observations, up to December 30, 2005, and is used for out-of-sample performance evaluation.

We focus on differences in goodness-of-fit implied by the conditional correlation matrix dynamics under the different model settings. We estimate our model in two steps as follows. First, we separately estimate the univariate conditional volatility dynamics for each single return series and include as possible conditioning variables in the threshold definition (i) its estimated conditional volatility and (ii) the first lag of all components in the multivariate return series. This threshold volatility structure has proven to produce good empirical results in applications of tree-structured GARCH models to financial data; see, for example, Audrino and Trojani (2006). This first step of the estimation procedure is kept identical for all models in which volatilities can be estimated separately from correlations: the CCC, the DCC, the RSDC and our tree-structured DCC model. In this way, we ensure that differences in the goodness of fit of these models with respect to the estimated conditional covariance matrix dynamics are exclusively due to differences in the explanatory power with respect to conditional correlations. In the second step of our estimation procedure, we estimate possible tree-structured thresholds and GARCH-type dynamics in conditional correlations. We include as possible conditioning variables for the definition of the threshold structure of conditional correlations (i) the first lag of the average conditional correlation shocks across returns and (ii) the first lag of all components of our multivariate return series; see again Section 2.2 for details.

4.1.1 Estimation results

The estimation results of our TreeDCC-GARCH model for the ten-dimensional time series of US stock returns introduced above are summarized in Table 1.

TABLE 1 ABOUT HERE.

Table 1, Panel A, highlights that at most two regimes are necessary to model the individual conditional variance dynamics accurately. The most important predictor variables impacting on the corresponding threshold structures are the lagged returns of Microsoft and Harley Davidson. The structure of the estimated conditional correlation dynamics in our model is summarized in Panel B of Table 1. Similar to volatilities, the most important and statistically significant predictor variable impacting on the threshold structure of conditional correlations is the lagged

return of Harley Davidson. Moreover, the complete threshold structure of conditional correlations is characterized using only two further lagged stock returns, the returns of Alcoa and Intel (in descending order of statistical significance), implying four correlation regimes overall. The first regime is associated with simultaneously low lagged Harley Davidson and Alcoa returns.⁵ The second one arises for lagged, low Harley Davidson returns and for large Alcoa returns. The third regime is obtained for lagged low Intel returns and large Harley Davidson returns. Finally, the fourth regime is caused by contemporarily not too low returns of Harley Davidson and Intel. An important difference between the estimated volatility and correlation dynamics is that the local correlation dynamics never exhibit GARCH-type effects across the different correlation regimes. In other words, conditional correlations are regime dependent, but piecewise constant. The local average correlation levels in the different regimes are similar and vary from 0.897 to 0.962. However, we find that the BIC criterion increases significantly in all cases when incorporating additional correlation regimes into the model. This finding is also supported by our out-of-sample tests on the model's forecasting power for correlations, which indicate an additional clear superiority of our model relative to a CCC model with constant correlations.

4.1.2 Multivariate performance results

We now compare the accuracy of the conditional correlation predictions implied by our model with those implied by the CCC, the DCC, the RSDC and the flexible multivariate GARCH model.⁶ In Table 2, we present the goodness-of-fit measures defined in Section 4 for the real data application to our ten-dimensional stock returns series.

TABLE 2 ABOUT HERE.

Our multivariate tree structured model achieves the highest goodness-of-fit overall across the models with identical volatility dynamics (TreeDCC, CCC, DCC and RSDC models). The RSDC model achieves a better in-sample log Likelihood criterion, but it implies worse MAE and MSE criteria both in sample and out-of-sample. Since the RSDC model has more than twice the number of parameters of the other models, we interpret this finding as an evidence of overfitting for this particular model. The TreeDCC model's better performance relative to CCC, DCC and RSDC ranges between 0.5% and to 17% depending on the goodness-of-fit measure applied. In comparison with the flexible multivariate GARCH model, we find that the out-of-sample forecasting performance of the latter specification is slightly better in two out of three

cases, despite this model’s having more than double the number of parameters of ours. The heavy parametrization of the flexible multivariate GARCH model can make it unpracticable for settings that include dozens to hundreds of individual time series. Our TreeDCC model can produce flexible conditional covariance specifications combined with a sufficient degree of parsimony. This is why the TreeDCC model, like the CCC, DCC and RSDC models, can be used to estimate the conditional variance-covariance dynamics of very high-dimensional time series settings.

4.1.3 Statistical and economic significance of the improvements.

In this section, we provide additional evidence of the statistical and economic significance of the goodness-of-fit improvements provided by our model. We focus first on the statistical significance of improvements in conditional variance-covariance forecasts. To this end, we first compute Hansen (2005) tests of superior predictive ability. p -values of these tests are reported in Table 2 for each model and each out-of-sample performance measure under investigation.

TABLE 2 ABOUT HERE.

Consistent with our previous results, the TreeDCC and flexible multivariate GARCH models are the only ones not significantly outperformed by any other model at the 5% significance level based on the MAE and MSE statistics. With respect to the negative log Likelihood statistic, the RDSC model is also not significantly outperformed by any other model at the 20% significance level. In this case, however, the p -values implied by TreeDCC and flexible GARCH models are even higher (above 85% and 60%, respectively). To formally characterize the set of models not significantly outperformed by other ones, we also computed Hansen et al. (2003) model confidence sets (MCS). Consistent with the previous findings for SPA tests, we find that the 5%-MCS based on the OS-MAE and OS-MSE criteria using the range and semi-quadratic statistics consist solely of the TreeDCC and flexible MGARCH models. 10%-MCS based on the OS-NL criterion include the TreeDCC, flexible MGARCH and RSDC models.

We now economically quantify the improvements implied by the different variance-covariance forecasts and compute the difference between the variance component of the average utility AU implied by TreeDCC model and the one implied by each of the other models. Table 3 presents results for three levels of risk aversion parameters $\lambda = 2, 7, 10$ and three target expected returns (6%, 10% and 15%) for the mean variance optimal portfolio. The resulting difference between average utilities is shown as an annualized fee and the (annualized) risk-free rate is 3%.

TABLE 3 ABOUT HERE.

The economic gains relative to CCC and DCC models range from about 1 to 180 basis points. The economically most important gains arise for highly risk-averse investors ($\lambda = 10$) and high target expected returns of 15%. Economically, very small gains arise for low risk aversion ($\lambda = 2$) and a target expected return of 6%. Consistent with our previous findings, economic gains are clearly smaller relative to RSDC and flexible MGARCH models. In this case, for a high risk aversion $\lambda = 10$ and a target expected return of 15% they are 52 and 81 basis points, respectively. We can use a Diebold-Mariano-type test and the joint test proposed in Engle and Colacito (2006) to test the statistical significance of these differences. Similarly to Bandi et al. (2008), we find that these differences are not significant at the 5% significance level.

Overall, these results provide evidence that economically and statistically significant differences in variance-covariance forecasts of our TreeDCC model are likely to arise with respect to the CCC and DCC models, whereas with respect to the RSDC and flexible GARCH models, such differences are less likely.

4.1.4 Sensitivity analysis

To investigate the robustness of our findings, it is useful to study the sensitivity of results to moderate changes in the structure of the estimated TreeDCC model. We perform this task along several dimensions.

First, we investigate whether moderate changes in the estimated location and threshold parameters of the conditional correlation process imply significantly different out-of-sample performances, given a fixed tree-structure. Panel A of Table 4 presents results for a TreeDCC model in which the location parameter estimates $\hat{c}_k$ have been altered by plus/minus one standard error. Similar findings are produced by modifying parameters $\hat{c}_k$ by a larger amount.

TABLE 4 ABOUT HERE.

Overall, the resulting effect on the out-of-sample performance relative to the estimated TreeDCC model are economically small. The largest impact arises by modifying the location parameter c_4 in the fourth regime. Using Diebold-Mariano tests for forecasting accuracy, we find that the only statistically significant change in the MAE and MSE out-of-sample criteria arises when changing the location parameter c_4 . However, these changes are typically smaller than the changes in MAE and MSE observed for the CCC, DCC, RDSC and flexible GARCH models.

Second, we analyze whether changes in the estimated TreeDCC threshold parameters imply significantly different out-of-sample performances, given fixed location parameters c_k and the tree-structure. Panel B of Table 4 presents results for a TreeDCC model in which the threshold parameter estimates are moved to the nearest possible threshold value, in positive and negative directions, on the discrete grid of admissible threshold values in our estimation procedure. Overall, differences in out-of-sample performance relative to the estimated TreeDCC model are economically small, even if they are larger than in the case where the location parameters c_k were changed. Using Diebold-Mariano tests for forecasting accuracy, we find that the only statistically significant change in the MAE and MSE out-of-sample criteria arises when changing the first threshold parameter d_1 . This is a natural finding, since d_1 is the threshold parameter associated with the first binary split in the tree-structure of our estimated TreeDCC model and it is therefore the statistically most significant one. However, it is interesting to note that if we also re-estimate the location parameters after having changed the threshold values, no significant difference in MSAE and MSE out-of-sample performance measures persists.

Third, we investigate whether moderate changes in the tree-structure of the estimated TreeDCC model imply significantly different out-of-sample performances. We consider the second, third and fourth best (in-sample) TreeDCC models estimated according to the procedure described in Section 3. The second-best estimated TreeDCC model implies four regimes, with the lagged returns of Harley-Davidson, Intel, and Tektronix as relevant predictor variables. The third- (fourth-) best estimated TreeDCC model implies three regimes, with the lagged returns of Harley-Davidson and Tektronix (Harley-Davidson and Alcoa) as relevant predictor variables. In all these models, the estimated first binary split in the tree is the one with the lagged return of Harley-Davidson as predictor variable, implying a quite remarkable stability for the estimation of the first-, most important, threshold. Panel B of Table 4 summarizes the out-of-sample performance results of the estimated second-, third-, and fourth-best TreeDCC models. Differences in performance relative to the first-best TreeDCC model are tiny and Diebold-Mariano tests for differences in forecasting power are never significant. According to these findings, the good forecasting performance of our TreeDCC modeling approach relative to other models seems to be quite robust across various similar choices of the parameters and threshold structure in the model. We obtain the same basic findings in a sensitivity analysis of results for the next two empirical applications below, but we do not report them here for brevity.

4.2 Second real data application: US stock index and bond returns

We consider a two-dimensional time series of (annualized) daily log-returns for the US S&P500 stock index and the US 30-year Treasury bond. The time period under investigation goes from January 3, 1996 to October 30, 2003 and contains 1899 trading days. The data is provided by *Tick Data*. As in the previous section, we exploit the tick-by-tick data to construct the series of realized volatilities and covariances between stock index and bond returns. As before, for forecasting evaluation purposes we split the sample in two sub-periods. The first sub-period consists of $n = 1219$ trading days and goes from January 3, 1996 to December 29, 2000. The second sub-period consists of the last three years of data ($n_{\text{out}} = 680$ observations).

4.2.1 Estimation results

The estimation procedure follows the same steps as the one described in Section 4.1. The individual variance structures and correlation threshold functions estimated for our TreeDCC model are summarized in Table 5.

TABLE 5 ABOUT HERE.

The estimated threshold functions for volatility each depend only on one lagged return of each time series of returns. Similarly to previous findings in the literature, e.g., Audrino and Trojani (2006), we find that the conditional variance of stock index returns implies more than two regimes. The estimated conditional variance of Treasury bond returns implies only two regimes.

We estimate only two regimes as well for the estimated conditional correlation function of stock index and Treasury bond returns. These regimes depend on the lagged return of the S&P500 index only. Each regime features local GARCH-type DCC effects as in Engle (2002), in which the regime-dependent parameters ϕ_k and λ_k imply a different persistence of correlation shocks: Here the lagged S&P500 index return is below or above the threshold value $d_1 = -3.847$. This threshold corresponds to the 37.5% quantile of the distribution of S&P500 index returns. The differences in the estimated parameters of the local DCC dynamics are statistically significant. Interestingly, we find that correlation shocks are quite substantially more persistent, conditional on a sufficiently negative past stock index return ($X_{t-1,S\&P500} \leq -3.847$). This might be interpreted as an indication that correlation shocks between bond and stock returns are likely to last for a longer time conditional on flight-to-quality effects caused by a drop in the stock market. It is also interesting to note that the correlation dynamics estimated for

stock index and bond returns are quite different from those estimated earlier in our application to a ten-dimensional stock returns time series. The flexibility of our TreeDCC setting is crucial in order to take these different dynamic correlation features of some asset returns adequately into account.

4.2.2 Multivariate performance results

Table 6 presents results for the goodness-of-fit measures in Section 4, estimated using stock index and Treasury bond return data.

TABLE 6 ABOUT HERE.

The TreeDCC model clearly has the best goodness-of-fit results across all in-sample and out-of-sample measures used. The improvements in out-of-sample performance range from 1% to 25%, depending on the goodness-of-fit measure applied. Constant or piecewise-constant conditional correlation dynamics (like the ones implied by the CCC and RSDC models) are largely rejected by the data and leads to very inaccurate correlation forecasts. In contrast to the previous application to single stock returns, GARCH-type DCC effects are now crucial in order to improve the model’s out-of-sample forecasting power for correlations.

4.2.3 Statistical and economic significance of the improvements.

p -values of Hansen (2005) SPA tests are reported in Table 6. They show that the TreeDCC model is the only one not significantly outperformed by any other model at the 10% confidence level and for all performance criteria used.

TABLE 6 ABOUT HERE.

Engle DCC model yields quite good results and is not outperformed by any model at the 5% confidence level. All other models are clearly dominated at standard significance levels. MCS results consistently support these findings. 10%-MCS using the range and the semi-quadratic statistics consist only of the TreeDCC model for all performance criteria used. 5%-MCS additionally include the DCC model

We conclude this section by economically quantifying the forecast improvements of our TreeDCC model relative to the competing ones. Table 7 presents the annualized fees (in basis points) that a conditional mean-variance investor would be ready to pay in order to forecast

future variance-covariance matrices with the TreeDCC instead of the other models considered. The (annualized) risk free rate is equal to 3%.

TABLE 7 ABOUT HERE.

The largest economic gains arise with respect to the CCC and RSDC models: For risk aversion parameters $\lambda \geq 7$ and target expected return larger than 10%, they range from approximately 100 to 440 basis points per year. Sizable economic gains between 70 and 294 basis points for risk aversion parameters $\lambda \geq 7$ and target expected return larger than 10% arise also with respect to the flexible MGARCH model. Economic gains relative to the DCC model are small in almost all cases and never larger than 43 basis points per year. Using Diebold-Mariano-type tests we find that estimated economic gains of the TreeDCC are significant at the 1% significance level relative to the CCC and RSDC model. Estimated economic gains relative to the flexible MGARCH model are significant at the 5% significance level. Improvements relative to the classical DCC model are not statistically significant. Overall, these results provide evidence that economically and statistically significant differences in variance-covariance forecasts of our TreeDCC model are likely to arise with respect to the CCC, RSDC and flexible GARCH models, and less so relative to the DCC model, which was clearly dominated by the TreeDCC model in the previous empirical application.

4.3 Third real data application: US government and corporate AAA bond returns

We conclude our empirical analysis by studying a three-dimensional time series of daily bond index log-returns for the following bond types: Corporate AAA intermediate bonds (5 years maturity), corporate AAA long bonds (10 years maturity) and Treasury long bonds (10 years maturity). The data under investigation are for the time period between April 23, 1996 to December 31, 2002, for a total of 1745 observations. We use the first 1223 observations (until the end of 2000) to estimate the different models. The remaining ones are used for forecasting evaluation purposes. The data is provided by *Lehmann Brothers*.

Since no tick-by-tick data are available in this setting, we rely on products of (centered) returns as proxies for the unobserved conditional covariances of returns. The main disadvantage of this approach is that the implied MAE and MSE measures are only noisily estimated and small differences of the unknown goodness-of-fit across models can be easily obscured by a low signal-to-noise ratio. Thus, estimated goodness-of-fit measures are likely to have discrimination power

only between models with large differences in forecasting ability. Moreover, as recently shown in Patton (2006), the MAE criterion leads to optimal forecasts that are significantly biased when conditional covariances are noisily approximated. Therefore, we report goodness-of-fit results only for the NL and MSE measures.

4.3.1 Estimation results

The individual variance structures and correlation threshold functions estimated for our TreeDCC model are summarized in Table 8.

TABLE 8 ABOUT HERE.

The conditional variance dynamics estimated for the long Treasury and corporate bond return series features a single regime. The one estimated for the intermediate corporate bond return series implies a single threshold, which depends on the lagged return of intermediate corporate bonds returns only. The conditional correlation dynamics feature a piecewise constant threshold structure with three regimes that depends exclusively on the lagged return of intermediate AAA bonds.

4.3.2 Multivariate performance results

Table 9 presents estimated goodness-of-fit measures defined in Section 4.

TABLE 9 ABOUT HERE.

Our TreeDCC model achieves the best goodness-of-fit across all criteria with the only exception being the IS-NL statistic. Estimated differences in the OS-NL measure amount to approximately 10%. Differences in the OS-MSE statistic are clearly smaller than in the previous applications, which is likely a consequence of the noisily estimated MSE statistic without using variance-covariance proxies computed with tick-by-tick data.

4.3.3 Statistical and economic significance of the improvements.

p -values of Hansen (2005) SPA tests are reported in Table 9 for the OS-NL and OS-MSE performance measures.

TABLE 9 ABOUT HERE.

The only model not outperformed by any other one is the TreeDCC. With respect to the OS-NL statistic, all other models are dominated by another one. Results for the OS-MSE statistic are more mixed, probably because of the noisily proxied variance-covariance matrix or returns. While the RSDC and flexible MGARCH models are clearly outperformed, the CCC and DCC models are not outperformed at the 5% significance level. 5% MCS implied by the OS-NL performance measure all consist of the TreeDCC model only. 10% MCS implied by the OS-MSE performance measure additionally include the CCC model. 5% MCS additionally include the CCC and DCC models.

We conclude by economically quantifying the forecast improvements of our TreeDCC model relative to the competing ones. Table 10 presents annualized fees that a conditional mean-variance investor would be ready to pay in order to switch from any of the alternative model forecasts to those implied by the TreeDCC model. The risk-free rate is equal to 2.5%.

TABLE 10 ABOUT HERE.

As expected, estimated economic gains are smaller than in the previous two applications. In some cases, e.g., relative to the flexible MGARCH model and for high risk aversions and target expected returns, they are also relevant. Diebold-Mariano-type test results imply a highly statistically significant difference of the economic gains of the TreeDCC relative to the flexible MGARCH model. The results of this test relative to the other models are not significant.

5 Conclusions

We propose a new multivariate DCC-GARCH model that extends previous models by admitting multivariate thresholds in conditional volatilities and correlations. The thresholds are modeled by a tree-structured partition of the multivariate state space and are estimated with all other model parameters. Three real data applications support the overall higher forecasting power of the TreeDCC model for return correlations, relative to Bollerslev’s CCC model, Engle’s DCC model, Pelletier’s RSDC model, and the flexible MGARCH model. We find that the conditional correlations of financial data are often characterized by multivariate thresholds and local GARCH-type structures and that the forecast improvements of our TreeDCC model are often economically relevant. Our model can cope in a parsimonious way with such features of the data even in applications with large cross-sections of financial assets. An interesting venue for future research is the joint empirical modeling of the dynamic correlation of the returns of several

asset classes, which are likely to exhibit rich threshold and GARCH-type effects that could be parsimoniously taken into account by our model.

Acknowledgements

We thank the editors, Torben Andersen and Serena Ng, the associate editor and two anonymous referees for many valuable comments. Financial support by the Swiss National Science Foundation (grants 100012-103781, 100012-105745 and NCCR FINRISK) and by the Foundation for Research and Development of the University of Lugano is gratefully acknowledged.

A Testing statistical relevance: The Model Confidence Set

We test formally for differences in forecasting power of the competing models in order to select, if possible, a best one (or a best subset of models) that significantly dominates the others in our real data application. To this end, we apply the Model Confidence Set (MCS) method proposed by Hansen et al. (2003).

Without loss of generality, let us denote by $\hat{D}_{t,wk}$ the differences of each term in the OS-MSE statistic:

$$\hat{D}_{t,wk} = \tilde{U}_{t;\text{model}_w} - \tilde{U}_{t;\text{model}_k}, \quad t = 1, \dots, n_{\text{out}}, \quad w, k = 1, \dots, 9, w < k,$$

where

$$\sum_{t=1}^{n_{\text{out}}} \tilde{U}_{t;\text{model}} = \text{OS-MSE}.$$

Statistics based on time averages $\bar{D}_{wk}$ of $\hat{D}_{t,wk}$ allow us to investigate whether there is a systematic difference in out-of-sample forecasting power between the different models. Tests based on $\hat{D}_{t,wk}$ are t-type tests. In a similar way, one can proceed by using a different out-of-sample goodness-of-fit statistic. In our application, we compute tests based on OS-MSE, OS-MAE and OS-NL.

The MCS is defined as the smallest set of models which, at a given confidence level α , cannot be significantly distinguished based on forecasting power. The MCS is determined after sequentially trimming the set of candidate models, which in our application consists of the five multivariate GARCH specifications introduced above. At each step of such a trimming procedure, the null-hypothesis of equal predictive ability (EPA) $\mathcal{H}_0 : \mathbb{E}[D_{t,wk}] = 0, \forall w, k \in \mathcal{M}$ is tested for the relevant set of models $\mathcal{M}$ at a confidence level α . In a first step, $\mathcal{M}$ consists of all models under investigation. If, in the first step, $\mathcal{H}_0$ is rejected, then the worst-performing model according to the relevant criterion is eliminated. The test procedure is then repeated for the new set $\mathcal{M}$ of surviving models, and it is iterated until the first non-rejection of the EPA hypothesis occurs. The set of resulting models is called the model confidence set $\widehat{\mathcal{M}}_\alpha$ at the given confidence level α . In our application, we work with $\alpha = 0.05, 0.10$.

Our tests of EPA are based on the range statistic T_R and the less conservative semi-quadratic statistic T_{SQ} :

$$T_R = \max_{k,w \in \mathcal{M}} \frac{|\bar{D}_{kw}|}{\sqrt{\widehat{\text{var}}(\bar{D}_{kw})}} \quad \text{and} \quad T_{SQ} = \sum_{k < w} \frac{\bar{D}_{kw}^2}{\widehat{\text{var}}(\bar{D}_{kw})},$$

where the sum in T_{SQ} is taken over the models in $\mathcal{M}$, $\bar{D}_{kw} = n_{\text{out}}^{-1} \sum_{t=1}^{n_{\text{out}}} \hat{D}_{t,kw}$, and $\widehat{\text{var}}(\bar{D}_{kw})$

is an estimate of $\text{var}(\overline{D}_{kw})$ obtained from a block-bootstrap of the series $\widehat{D}_{t,kw}$, $t = 1, \dots, n_{\text{out}}$. Using statistics T_R or T_{SQ} , we test the null hypothesis EPA at confidence level α for model set $\mathcal{M}$. If hypothesis EPA is rejected for model set $\mathcal{M}$, we compute a worst-performing index, in order to trim the worst-performing model from $\mathcal{M}$.

The worst-performing index for Model_k is computed as the mean across models $w \neq k$ of statistic $\overline{D}_{wk}$. More specifically, it is defined as $\overline{D}_k / \sqrt{\text{var}(\overline{D}_k)}$, where $\overline{D}_k = \text{mean}_{w \neq k \in \mathcal{M}} \overline{D}_{kw}$. As above, our estimate of $\text{var}(\overline{D}_k)$ is based on a block-bootstrap. The model with the highest worst-performing index is finally trimmed from $\mathcal{M}$. Consistency of estimates of the (asymptotic) distributions of T_R and T_{SQ} can be proved under mild regularity conditions on the bootstrap. For details, see Hansen et al. (2003).

Notes

¹To simplify the notation, conditional means of $\mathbf{X}_t$ have been set to zero in (2.1).

²For this case, the trivial parametrization $\bar{R} = \text{diag}[\bar{Q}]^{-1/2} \bar{Q} \text{diag}[\bar{Q}]^{-1/2}$ is superfluous.

³See again Audrino and Trojani (2006), Section 2.3, for details on this estimation procedure.

⁴All estimated models also include a linear autoregressive conditional mean function modeled by a simple diagonal VAR(1) process.

⁵Note that with the terminology “low” (“large”) returns we mean in fact returns that are below (above) the thresholds in the estimated threshold functions.

⁶Since all models, with the only exception of the flexible multivariate GARCH model, have the same individual volatility dynamics, any difference in the goodness-of-fit for the forecasts of the return covariance matrix is due to a difference in the quality of the forecasts for conditional correlations.

References

- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. 2001. The distribution of exchange rate volatility. *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. 2003. Modeling and forecasting realized volatility. *Econometrica* **71**(2), 579–625.
- Audrino F. (2006). Tree-structured multiple regimes in interest rates. *Journal of Business & Economic Statistics* **24**, No. 3, 338-353.
- Audrino, F. and Bühlmann, P. (2001). Tree-Structured GARCH Models. *Journal of the Royal Statistical Society, Series B*, **63**, No. 4, 727-744.
- Audrino, F. and Trojani, F. (2006). Estimating and predicting multivariate volatility regimes in global stock markets. *Journal of Applied Econometrics* **21**, No. 3, 345-369.
- Bandi, F., Russell, J.R. and Zhu, Y. (2008). Using high-frequency data in dynamic portfolio choice. *Econometric Reviews* **27**, 163-198.
- Barndorff-Nielsen, O. E. and Shephard, N. 2001. Non-gaussian ornstein-uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society, Series B*, **63**, 167–241.
- Barndorff-Nielsen, O. E. and Shephard, N. 2002a. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B*, **64**, 253–280.
- Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics* **72**, 498-505.
- Corsi, F. and Audrino, F. (2007). Realized correlation tick-by-tick. Discussion Paper, Department of Economics, University of St. Gallen.
- Engle, R.F. (2002). Dynamic conditional correlation - a simple class of multivariate GARCH models. *Journal of Business and Economic Statistics* **20**, 339-350.
- Engle, R. and Colacito, R. (2006). Testing and valuing dynamic correlations for asset allocation. *Journal of Business and Economic Statistics* **24**, 238-253.

- Engle, R.F. and Sheppard, K. (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. Mimeo, University of California, San Diego.
- Fleming, J., Kirby, C. and Ostdiek, B. (2001). The economic value of volatility timing. *Journal of Finance* **56**, 329-352.
- Fleming, J., Kirby, C. and Ostdiek, B. (2001). The economic value of volatility timing using realized volatility. *Journal of Financial Economics* **67**, 473-509.
- Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics* **23**, 365-380.
- Hansen, P.R., Lunde, A. and Nason, J.M. (2003). Choosing the best volatility models: The model confidence set approach. *Oxford Bulletin of Economics and Statistics* **65**, 839-861.
- Ledoit, O., Santa-Clara, P. and Wolf, M (2003). Flexible multivariate GARCH modeling with an application to international stock markets. *The Review of Economics and Statistics*, **85**, Issue 3, 735-747.
- Newey, W.K. and McFadden, D. (1994). Large sample estimation and hypothesis testing. *Handbook of Econometrics*, Vol. **4**, 2111-2245. North-Holland, Amsterdam.
- Pagan, A. (1986). Two stage and related estimators and their applications. *The Review of Economic Studies* **53**, No. 4, 517-538.
- Patton, A.J. (2006). Volatility forecast evaluation and comparison using imperfect volatility proxies. Working paper, London School of Economics.
- Pelletier, D. (2006). Regime switching for dynamic correlations. *Journal of Econometrics* **131**, 445-473.
- West, K.D., Edison, H.J. and Cho, D. (1993). A utility-based comparison of some models of exchange rate volatility. *Journal of International Economics* **35**, 23-45.

Panel A: Individual conditional variance structures.

Series	Regimes	Optimal predictors
Alcoa	1	—
Citigroup	2	Microsoft
Hasbro	2	Harley Davidson
Harley Davidson	2	Harley Davidson
Intel	1	—
Microsoft	1	—
Nike	2	Exxon
Pfizer	2	Microsoft
Tektronix	2	Harley Davidson
Exxon	1	—

Panel B: Conditional correlation structure and parameters.

Cond. corr. structure	Cond. corr. parameters
$\mathcal{R}_k$	c_k
$X_{t-1, \text{Harley Davidson}} \leq -19.983$ and	0.918
$X_{t-1, \text{Alcoa}} \leq -3.553$	(0.029)
$X_{t-1, \text{Harley Davidson}} \leq -19.983$ and	0.897
$X_{t-1, \text{Alcoa}} > -3.553$	(0.056)
$X_{t-1, \text{Harley Davidson}} > -19.983$ and	0.962
$X_{t-1, \text{Intel}} \leq -15.016$	(0.016)
$X_{t-1, \text{Harley Davidson}} > -19.983$ and	0.935
$X_{t-1, \text{Intel}} > -15.016$	(0.023)

Table 1: Estimation results for a multivariate time series of ten daily (annualized) US stock returns (in %). Data are for the in-sample time period between January 2, 2001 and December 31, 2003, consisting of 752 observations. Estimated individual conditional variance structures (Panel A) and estimated conditional correlation structure and parameters (Panel B) are for the tree-structured GARCH-DCC model fit. Standard errors computed using 1000 model-based bootstrap replications are given in parentheses.

US equity returns: Goodness-of-fit results.							
Model	# par.	IS-			OS-		
		NL	MAE	MSE	NL	MAE	MSE
CCC-GARCH	80	34959	269.413	176147	21694 (0.0002)	105.470 (0)	73205.9 (0)
DCC-GARCH	82	34942	269.525	177848	21652 (0.008)	103.259 (0)	72731.0 (0)
RSDC-GARCH	173	34684	276.604	182999	21634 (0.2203)	105.470 (0)	74520.1 (0)
TreeDCC-GARCH	84	34927	241.830	152813	21594 (0.857)	86.0280 (0.124)	68831.5 (0.453)
F-MGARCH	185	34378	188.850	135608	21588 (0.603)	79.6050 (0.509)	69832.7 (0.0826)

Table 2: Goodness-of-fit of different models for a multivariate time series of ten daily (annualized) US stock returns (in %). Data are for the time period between January 2, 2001 and December 30, 2005, for a total of 1256 observations. The in-sample estimation period goes from the beginning of the sample to the end of 2003 (752 observations). NL, MAE and MSE are multivariate versions of the standard univariate negative log-likelihood, the mean absolute error, and the mean squared error statistics. # par. reports the number of parameters estimated by the different models. For the out-of-sample performance measures, p -values of superior predictive ability (SPA) tests are reported in parentheses.

US equity returns: Annualized fees (in basis points).				
Target	Alternative model	$\lambda = 2$	$\lambda = 7$	$\lambda = 10$
6%	CCC-GARCH	2.24	7.87	11.24
	DCC-GARCH	1.97	6.88	9.83
	RSDC-GARCH	0.66	2.31	3.31
	F-MGARCH	0.94	3.29	4.73
10%	CCC-GARCH	12.24	42.85	61.22
	DCC-GARCH	10.70	37.46	53.52
	RSDC-GARCH	3.60	12.59	17.99
	F-MGARCH	4.37	15.29	21.85
15%	CCC-GARCH	35.98	125.92	179.88
	DCC-GARCH	31.46	110.09	157.28
	RSDC-GARCH	10.58	37.02	52.88
	F-MGARCH	16.17	56.60	80.85

Table 3: The table contains the annualized fees (in basis points) that a conditional mean-variance investor with absolute risk-aversion parameter $\lambda = 2, 7$, and 10 would be willing to pay to perform volatility timing using the one-step-ahead conditional correlation and covariance forecasts from the TreeDCC model (benchmark model) versus those obtained using the CCC, DCC, RSDC, and the flexible MGARCH models. The portfolio weights are obtained by minimizing the one-step-ahead conditional variance forecast of a portfolio containing ten US stocks and the risk-free asset, for a given target expected return on the portfolio. The annual risk-free rate is set to be equal to 3%. Data are for the out-of-sample time period beginning in January, 2004 and ending in December, 2005, for a total of 504 daily observations.

Panel A: TreeDCC sensitivity analysis: location parameters.								
Parameter	+1 standard error				−1 standard error			
changed	NL	MAE	MSE	Δ AU	NL	MAE	MSE	Δ AU
c_1	21595	86.2212	68882.6	< 0.12	21593	85.8391	68783.9	< 0.11
c_2	21594	86.1893	68869.9	< 0.61	21594	85.7198	68765.2	< 0.50
c_3	21595	86.3414	68911.8	< 0.60	21592	85.7175	68753.5	< 0.45
c_4	21597	87.4117	69148.0	< 1.88	21592	84.6671	68528.3	< 1.04
all c_k	21599	88.0797	69317.7	< 3.19	21589	84.0113	68367.8	< 1.88

Panel B: TreeDCC sensitivity analysis: threshold parameters.						
Threshold	Direction	NL	MAE	MSE	Δ AU	
d_1 in Harley Davidson	positive	21592	85.5696	68724.9	< 4.56	
d_1 in Harley Davidson	negative	21595	86.2449	68885.8	< 1.92	
d_2 in Alcoa	positive	21594	86.0447	68835.2	< 0.21	
d_2 in Alcoa	negative	21594	86.0019	68825.7	< 0.03	
d_3 in Intel	positive	21595	86.4877	68934.2	< 1.89	
d_3 in Intel	negative	21593	85.7289	68759.3	< 2.74	

Panel C: TreeDCC sensitivity analysis: tree-structure.				
TreeDCC model	NL	MAE	MSE	Δ AU
second-best	21595	85.9964	68825.6	< 0.10
third-best	21594	86.1244	68846.2	< 0.01
fourth-best	21693	86.1569	68852.4	< 0.10

Table 4: TreeDCC forecasting stability analysis. Multivariate out-of-sample performance results when one or more parameters in the TreeDCC model are changed: location parameters (Panel A) are changed adding (subtracting) one standard error, threshold parameters (Panel B) are changed according to the nearest grid points (empirical quantiles) used in the estimation procedure, and the tree-structure (Panel C) is changed considering the second, third and fourth best (in-sample) performing TreeDCC models. As before, NL, MAE and MSE are multivariate versions of the standard univariate negative log-likelihood, the mean absolute error, and the mean squared error statistics. | Δ AU| is the annualized fee (in absolute values and in basis points) an investor has to pay to switch from the estimated TreeDCC model and the modified version of it.

Panel A. Individual conditional variance structures.		
Series	Regimes	Optimal predictors
S&P500	3	S&P500
30-years Treasury bond	2	30-years Treasury bond

Panel B: Conditional correlation structure and parameters.		
Cond. corr. structure	Cond. corr. parameters	
$\mathcal{R}_k$	ϕ_k	λ_k
$X_{t-1, \text{S\&P500}} \leq -3.847098$	0.0490	0.9129
	(0.001)	(0.024)
$X_{t-1, \text{S\&P500}} > -3.847098$	0.0222	0.9724
	(0.002)	(0.019)

Table 5: Estimation results for a two-dimensional time series of daily (annualized) returns (in %) for the US S&P500 index and the US 30-years Treasury bond. Data are for the in-sample time period between January 3, 1996 and December 29, 2000, consisting of 1219 observations. Estimated individual conditional variance structures (Panel A) and estimated conditional correlation structure and parameters (Panel B) are for the tree-structured GARCH-DCC model fit. Standard errors computed using 1000 model-based bootstrap replications are given in parentheses.

US index and bond returns: Goodness-of-fit results.							
Model	# par.	IS-			OS-		
		NL	MAE	MSE	NL	MAE	MSE
CCC-GARCH	25	9334.7	63.6310	21197.8	5511.5 (0.0003)	97.2933 (0)	27592.1 (0)
DCC-GARCH	27	9299.2	59.4418	20458.1	5451.7 (0.4469)	74.8853 (0.0868)	22357.5 (0.3319)
RSDC-GARCH	30	9334.7	63.6310	21197.8	5511.5 (0.0014)	97.2933 (0)	27592.0 (0)
TreeDCC-GARCH	29	9290.2	58.7296	20413.5	5440.9 (0.8435)	70.7387 (0.2711)	20544.6 (0.2246)
F-MGARCH	13	9389.2	65.0761	22752.7	5483.4 (0.0012)	86.1455 (0.0007)	28271.6 (0.0652)

Table 6: Goodness-of-fit of different models for a two-dimensional time series of daily (annualized) returns (in %) on the US S&P500 index and the 30-years Treasury bond. Data are for the time period between January 3, 1996 and October 30, 2003, for a total of 1899 observations. The in-sample estimation period goes from the beginning of the sample to the end of 2000 (1219 observations). NL, MAE and MSE are the multivariate versions of the standard univariate negative log-likelihood, the mean absolute error, and the mean squared error statistics. # par. reports the number of parameters estimated in the different models. For the out-of-sample performance measures, p -values of superior predictive ability (SPA) tests are reported in parentheses.

US index and bond returns: Annualized fees (in basis points).

Target	Alternative model	$\lambda = 2$	$\lambda = 7$	$\lambda = 10$
6%	CCC-GARCH	5.54	19.38	27.68
	DCC-GARCH	0.54	1.87	2.68
	RSDC-GARCH	5.54	19.38	27.68
	F-MGARCH	3.68	12.89	18.41
10%	CCC-GARCH	30.14	105.50	150.72
	DCC-GARCH	2.92	10.21	14.59
	RSDC-GARCH	30.14	105.50	150.72
	F-MGARCH	20.05	70.18	100.26
15%	CCC-GARCH	88.58	310.05	442.92
	DCC-GARCH	8.57	29.99	42.85
	RSDC-GARCH	88.58	310.05	442.92
	F-MGARCH	58.93	206.25	294.64

Table 7: The table contains the annualized fees (in basis points) that a conditional mean-variance investor with absolute risk-aversion parameter $\lambda = 2, 7$, and 10 would be willing to pay to perform volatility timing using the one-step-ahead conditional correlation and covariance forecasts from the TreeDCC model (benchmark model) versus those obtained using the CCC, DCC, RSDC, and the flexible MGARCH models. The portfolio weights are obtained by minimizing the one-step-ahead conditional variance forecast of a portfolio containing the US S&P500 index, the 30-years US Treasury bond, and the risk-free asset for a given target expected return on the portfolio. The annual risk-free rate is set to be equal to 3%. Data are for the out-of-sample time period beginning in January, 2001 and ending in October, 2003, for a total of 680 daily observations.

Panel A. Individual conditional variance structures.		
Series	Regimes	Optimal predictors
Corporate AAA intermediate	2	Corporate AAA intermediate
Corporate AAA long	1	—
Government long	1	—

Panel B: Conditional correlation structure and parameters.	
Cond. corr. structure	Cond. corr. parameters
$\mathcal{R}_k$	c_k
$X_{t-1, \text{Corporate AAA intermediate}} \leq -0.0583$	0.965 (0.059)
$-0.0583 < X_{t-1, \text{Corporate AAA intermediate}} \leq 0$	0.985 (0.039)
$X_{t-1, \text{Corporate AAA intermediate}} > 0$	0.947 (0.038)

Table 8: Estimation results for a three-dimensional time series of daily returns (in %) of the following US bonds: 5-years corporate AAA (intermediate), 10-years corporate AAA (long) and 10-years government (long). Data are for the in-sample time period between April 23, 1996 and December 29, 2000, consisting of 1223 observations. Estimated individual conditional variance structures (Panel A) and estimated conditional correlation structure and parameters (Panel B) are for the tree-structured GARCH-DCC model fit. Standard errors computed using 1000 model-based bootstrap replications are given in parentheses.

US corporate and government bond returns: Goodness-of-fit results.					
Model	# par.	IS-		OS-	
		NL	MSE	NL	MSE
CCC-GARCH	20	13.4818	0.1716	433.01 (0)	0.2135 (0.4221)
DCC-GARCH	22	−297.56	0.1717	454.41 (0)	0.2140 (0.0903)
RSDC-GARCH	29	−336.12	0.1730	418.16 (0.0013)	0.2164 (0)
TreeDCC-GARCH	23	−100.02	0.1716	377.91 (0.4721)	0.2131 (0.5932)
F-MGARCH	24	−41.331	0.1844	425.32 (0)	0.2344 (0)

Table 9: Goodness-of-fit of different models for a three-dimensional time series of daily returns (in %) of US 5-years and 10-years corporate AAA bonds and 10-years Treasury bonds. Data are for the time period between April 23, 1996 and December 31, 2002, for a total of 1745 observations. The in-sample estimation period goes from the beginning of the sample to the end of 2000 (1223 observations). NL and MSE are multivariate versions of the standard univariate negative log-likelihood and the mean squared error statistics. # par. reports the number of parameters estimated in the different models. For the out-of-sample performance measures, p -values of superior predictive ability (SPA) tests are reported in parentheses.

US corporate and government bond returns: Annualized fees (in basis points).

Target	Alternative model	$\lambda = 2$	$\lambda = 7$	$\lambda = 10$
6%	CCC-GARCH	0.93	3.24	4.63
	DCC-GARCH	0.53	1.86	2.66
	RSDC-GARCH	0.22	0.78	1.11
	F-MGARCH	1.34	4.69	6.69
10%	CCC-GARCH	4.25	14.88	21.25
	DCC-GARCH	2.44	8.55	12.22
	RSDC-GARCH	1.02	3.58	5.12
	F-MGARCH	6.14	21.48	30.69
15%	CCC-GARCH	11.81	41.33	59.04
	DCC-GARCH	6.79	23.75	33.93
	RSDC-GARCH	2.84	9.95	14.22
	F-MGARCH	17.08	59.77	85.39

Table 10: The table contains the annualized fees (in basis points) that a conditional mean-variance investor with absolute risk-aversion parameter $\lambda = 2, 7$, and 10 would be willing to pay to perform volatility timing using the one-step-ahead conditional correlation and covariance forecasts from the TreeDCC model (benchmark model) versus those obtained using the CCC, DCC, RSDC, and the flexible MGARCH models. The portfolio weights are obtained by minimizing the one-step-ahead conditional variance forecast of a portfolio containing the US 5-years and 10-years corporate AAA bonds, the 10-years Treasury bonds, and the risk-free asset for a given target expected return on the portfolio. The annual risk-free rate is set to be equal to 2.5%. Data are for the out-of-sample time period beginning in January, 2001 and ending in December, 2002, for a total of 522 daily observations.