North American Actuarial Journal

Vol.13 No.1

Multivariate Models of Equity Returns for Investment Guarantees Valuation

Mathieu Boudreault* and Christian-Marc Panneton†

Abstract

In this paper we investigate the valuation of investment guarantees in a multivariate (discrete-time) framework. We present how to build multivariate models in general, and we survey the most important multivariate GARCH models. A direct multivariate application of regime-switching models is also discussed, as is the estimation of these models using maximum likelihood and their comparison in a multivariate setting. The computation of the CTE provision is further presented.

We have estimated the models with a multivariate dataset (Canada, United States, United Kingdom, and Japan), and we compared the quality of their fit using multiple criteria and tests. We observe that multivariate GARCH models provide a better overall fit than regime-switching models. However, regime-switching models appropriately represent the fat tails of the returns distribution, which is where most GARCH models fail. This leads to significant differences in the value of the CTE provisions, and, in general, provisions computed with regime-switching models are higher. Thus, the results from this multivariate analysis are in line with what was obtained in the literature of univariate models.

1. Introduction

Following the report from the Task Force on Segregated Fund Investments (2002), the Canadian Institute of Actuaries (CIA) has strongly recommended the use of stochastic equity returns models. Without mandating a specific model, the actuary has the freedom to choose any model as long as it is at least as conservative than the CIA-provided calibration points on accumulation factors. The American Academy of Actuaries (AAA) followed its Canadian counterpart and also recommended stochastic models for variable annuities products (see the 2005 AAA RBC C3 Phase II report).

Since the publication of the CIA Task Force's reports, a multitude of equity returns models have been proposed. The very important contributions of Hardy (2001, 2003) with respect to the regime-switching lognormal (RSLN) model seemed to have convinced the insurance industry and the Canadian regulators (the Office of the Superintendent of Financial Institutions especially) to use this model for segregated funds valuation. Although the work from the CIA and AAA was a tremendous step in improving the future solvency of insurance companies, their focus is exclusive on single-asset (univariate) models. Insurance companies do sell segregated fund products on multiple assets, and correlation between funds can materially affect the provision.

In this paper we investigate the valuation of investment guarantees in a multivariate (discrete-time) framework. We begin by presenting how to build a multivariate time series model that supports the evidence of time-varying volatility and correlation. We then focus our investigation mainly on two

* Mathieu Boudreault, MSc, FSA, is a PhD candidate in the Department of Management Sciences, HEC Montréal, 3000 chemin de la Côte-Sainte-Catherine, Montréal (Québec), H3T 2A7, Canada, mathieu.boudreault@hec.ca. He is also an Assistant Professor in the Department of Mathematics, Université du Québec à Montréal 201 avenue du Président-Kennedy, Montréal (Québec), H2X 3Y7, Canada, boudreault.mathieu@uqam.ca.

† Christian-Marc Panneton, MSc, FSA, FCIA, CFA, CERA, is a financial risk manager at Industrial Alliance, 1080 Grande Allée Ouest, Quebec City, QC, G1K 7M3, Canada, christian-marc.panneton@inalco.com.

classes of multivariate models: the multivariate GARCH and regime-switching models.[1] To the best of our knowledge, no author in actuarial science has ever proposed multivariate time series models in the context of investment guarantees. Thus, we have surveyed in our paper the most important contributions to the literature of multivariate GARCH models that could be used in our context.[2] Consistent with current applications of multivariate regime-switching models,[3] we explain how to add a Markovian process in a multivariate setting. Modifications for local specificities are briefly mentioned. We also discuss how we can estimate these models using maximum likelihood. Moreover, we present how to compute the conditional tail expectation (CTE) provision in the context of investment guarantees, which is a simple closed-form solution in the case of multivariate regime-switching models. Different methods and approaches to select and validate a multivariate model are also discussed.

One important contribution of our paper is to analyze and compare the fit of the models to a multivariate dataset (monthly returns from Canadian, U.S., U.K., and Japanese markets). We have first fitted by maximum likelihood the models presented in this paper and have come to the conclusion that using typical model selection criteria, such as the likelihood ratio test, the AIC, or the SBC, the multivariate GARCH family of models has in general a better overall fit than the regime-switching models. We have used econometric tests to check which model better represents the dynamics of variance and covariance in the sample. We obtained that multivariate GARCH models better represent the heteroskedasticity of covariance in the chosen dataset. This is, however, at the cost of a significant lack of fit in the tails of the returns distributions because regime-switching models showed a strong capability to replicate the fat-tailness of the dataset. As a final exercise, we have computed the CTE provision for most of the models and observed that the provisions are higher with regime-switching models, reflecting the quality of the fit in the tail of the returns distribution.

In the next section we show how to build a multivariate model, we survey important multivariate GARCH models, and we present our multivariate application to the regime-switching model. Estimation of the models is discussed in Section 3, and the computation of the CTE provision with multivariate models is shown in Section 4. Section 5 is the heart of our analysis, where we display our main results in a numerical example. We present our dataset, estimate the models, analyze their fit, and compare the resulting CTE provisions in this section. Section 6 concludes.

2. MULTIVARIATE MODELS

One of the main purposes of our paper is to be able to represent the evolution of equity returns of multiple assets, all in a consistent multivariate setting. It is widely known that the volatility of financial time series data is not constant (see, e.g., Campbell, Lo, and MacKinlay 1997; Tsay 2005, and references therein), so we need to address this issue in a multivariate framework. Moreover, Tse (2000) has documented using daily data that Asian equity markets have nonconstant correlation. We strongly believe that it is very likely to be the case for North American stock markets so that we also need to model the time-varying dynamics of correlation (or any other type of dependence). We present in this section a general multivariate model construction guideline that can be used to build a multivariate model. We introduce next our different approaches to the application of regime-switching models to a multivariate setting, and finally, we survey the large family of multivariate GARCH models.

2.1 Multivariate Model Construction

Two approaches can mainly be used to build a multivariate time series model. The first technique, which we denote the direct multivariate approach, directly models the time-varying dynamics of the

[1] Because stochastic variance models can sometimes be difficult to estimate in a univariate framework, we will not discuss multivariate extensions of this model in this paper.

[2] See, e.g., Bollerslev, Engle, and Wooldridge (1988), Bollerslev (1990), Engle and Kroner (1995), and Engle (2002).

[3] See, e.g., Billio and Pelizzon (2000); Ang and Bekaert (2002a, b).

variance-covariance matrix of a multivariate distribution such as a normal or Student's distribution. The dependence relationship between the asset returns is defined by the multivariate distribution in itself. In the second method, copulas are used to link the univariate models that specify the marginal dynamics of the different markets.[4]

However, a direct application of univariate time series models with copulas will not provide time-varying correlation. Because we have noted that this feature may be important for multivariate time series modeling, the copula approach will not be investigated in this paper. Thus, the models built upon the direct multivariate approach will easily accommodate time-varying correlation. This is described below.

We begin by defining most of the notation of this paper. Suppose that $y_t^{(i)}$ denotes the continuously compounded return of the ith asset ($i = 1, 2, \ldots, N$) for the period of time between $t - 1$ and t. The column vector $\mathbf{y}_t$ represents the set of returns for each of the N assets, such that

$$\mathbf{y}_t' = [y_t^{(1)}, \ldots, y_t^{(N)}].$$

The formulation of the direct multivariate approach depends on the presence or not of a latent process. Let us begin with the case where there is no latent variable in the returns process (such as GARCH models). Assume that $\mathcal{F}_t$ contains the set of previous returns up to time t. Then the direct multivariate approach without a latent process focuses on modeling $\mathbf{y}_t|\mathcal{F}_{t-1}$ as

$$\mathbf{y}_t = f(\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \ldots) + \varepsilon_t,$$

where $f(\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \ldots)$, is some function of the past observations (such as a Vector Autoregressive model) and ε_t has a multivariate normal distribution.[5] It is supposed that

$$\varepsilon_t' = [\varepsilon_t^{(1)}, \ldots, \varepsilon_t^{(N)}]$$

is serially independent, and its variance-covariance matrix, conditional on the past information, is given by $\mathbf{H}_t = Var[\varepsilon_t|\mathcal{F}_{t-1}]$. In other words, returns are serially correlated in their squares only. Thus, the modeling of volatility and correlation is embedded in the dynamics of the conditional covariance matrix $\mathbf{H}_t$. For example, the multivariate GARCH models presented in the upcoming sections will focus on defining the dynamics of $\mathbf{H}_t$. For simplicity, we will suppose throughout this paper that

$$f(\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \ldots) = \boldsymbol{\mu}$$

is a constant, representing the mean return vector.

When there is a latent process, such as in a regime-switching model, modeling is slightly more complicated. Assume that $\mathcal{F}_t$ still contains the set of previous returns up to time t and that the latent process is defined by $\{\rho_t, t = 0, 1, \ldots\}$. Then the direct multivariate approach with a latent process focuses on modeling $\mathbf{y}_t|\mathcal{F}_{t-1}, \rho_t, \rho_{t-1}, \ldots$ as

$$\mathbf{y}_t = f(\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \ldots) + \varepsilon_t,$$

where f is still some function of past observations, and ε_t (given $\rho_t, \rho_{t-1}, \ldots$) has a multivariate normal distribution.

The main difference, with or without a latent process, is that $Var[\varepsilon_t|\mathcal{F}_{t-1}]$ is constant without a latent process, thus fully predictable. For example, in a multivariate GARCH model, $Var[\varepsilon_t|\mathcal{F}_{t-1}]$ is a deterministic function of past covariances and errors. With a latent process, $Var[\varepsilon_t|\mathcal{F}_{t-1}]$ is stochastic.

[4] The literature on copulas is very rich; important references are Joe (1997), Nelsen (1998), and Frees and Valdez (1998) and references therein.

[5] If the purpose of the model is to understand the dynamics of the returns, then it would also be possible to use a multivariate Student's distribution for the error term. However, since our paper requires the computation of moments of exp (y_t) (like a CTE), a Student's distribution would not be appropriate because these moments do not exist. For this reason, we will focus only on a multivariate normal distribution for the error term.

2.2 Regime-Switching Models

The class of (univariate) regime-switching models was introduced by Hamilton (1989) in the context of representing business cycles in the United States. The model has been used in numerous other settings since 1989, notably, for the valuation of segregated funds guarantees as it was initially proposed by Hardy (2001). She has suggested the use of a regime-switching lognormal model (RSLN) for the long-term dynamics of the Canadian equity returns. Details can be found in Hardy (2001, 2003) and Hardy, Freeland, and Till (2006). However, those models focus on the dynamics of a single asset, whereas most insurance companies sell many different correlated funds. We discuss in this section of multivariate extensions to the univariate regime-switching model.

In the course of financial history we have observed many times that different markets crashed simultaneously. For example, during the October 1987 crash, the 1997–1998 Asian financial crisis, the 2000 tech bubble burst, and the more recent July 2007 subprime mortgage crisis, markets all over the world suffered very important losses. This emphasizes the importance of modeling regime shifts that may influence all markets simultaneously. The multivariate regime-switching model with a global regime that we present in this subsection is an attempt to represent those types of events.

Assume there is an outside process $\{\rho_t, t = 0, 1, 2, \ldots\}$ that influences each asset in the portfolio. In other words, all assets are in the same regime every time. The returns on each asset are then given by

$$\mathbf{y}_t = \boldsymbol{\mu}_{\rho_t} + \varepsilon_{\rho_t},$$
$$\varepsilon_{\rho_t}|\rho_t \sim MVN(\mathbf{0}, \Sigma_{\rho_t}),$$

for $\rho_t = 1, 2, \ldots, K$, where $\boldsymbol{\mu}'_{\rho_t} = [\mu^{(1)}_{\rho_t}, \ldots, \mu^{(N)}_{\rho_t}]$, Σ_{ρ_t} is the covariance matrix in regime ρ_t and $\mathbf{0}$ is an N-vector of zeros. Thus,

$$(\mathbf{y}_t|\rho_t = k) \sim MVN(\boldsymbol{\mu}_k, \Sigma_k), \qquad k = 1, 2, \ldots, K,$$

which is an N-variate normal distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix Σ_k. One can see that this is a direct application of Hardy (2001) in the context of multiple assets.

For example, in a two-regime framework, this would mean that all markets are simultaneously in a low (high)-volatility regime, and if the Markov chain goes into the other state, all markets would move back to the high (low)-volatility regime, regardless of the specific dynamics of each market involved in the portfolio. This model will be known as the global regime approach throughout this paper.

In this model there is a single transition probability matrix and K correlation matrices to estimate. It is possible, however, to simplify matters by considering only a single correlation matrix for the K regimes, when there is evidence of constant correlation. Other possibilities would be to use fewer than K correlation matrices to reduce the number of parameters to estimate. We will use such an assumption in the numerical example of Section 5 with monthly data.

With a global regime, all assets are simultaneously in the same state because the latent process exerts a global influence on the returns. The specific market conditions that would drive the volatility of an asset are not accounted for in such a method. However, it is possible to add regimes in the Markov process to allow for transitions that do not affect all markets. Because they greatly increase the number of parameters to estimate, such modifications are not considered in this paper.

2.3 GARCH Models

We know that the main idea behind univariate GARCH models is to represent the current volatility as a function of past errors and volatilities.[6] The approach is similar when applied in a multivariate context: that is, we have to represent the evolution of the variance-covariance matrix of the joint errors as a

[6] For an extensive review of univariate GARCH models, the reader should check Bollerslev, Chou, and Kroner (1992), Bollerslev, Engle, and Nelson (1994), Campbell, Lo, and MacKinlay (1997), and Tsay (2005).

function of past covariances and errors. This is basically what the multivariate GARCH models surveyed in this paper do. It should be noted that we restrict the presentation of the models to those we find the most important.

2.3.1 VECH Model

Presented by Bollerslev, Engle, and Wooldridge (1988) as an approach to model the time-varying market premium in the CAPM, the VECH model, in its simplest and oldest presentation, is a direct application of a GARCH(1,1) structure, for each component of the variance-covariance matrix.[7] Because this definition will not guarantee the positive semidefiniteness (PSD) of the covariance matrix, the model has to be reformulated as

$$\mathbf{H}_t = (\mathbf{CC}') + (\mathbf{AA}') \otimes (\varepsilon_{t-1}\varepsilon'_{t-1}) + (\mathbf{BB}') \otimes \mathbf{H}_{t-1},$$

where $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are lower triangular matrices, and $\otimes$ is the element-by-element multiplication.

In an attempt to reduce the number of parameters, the model can be further simplified by assuming $\mathbf{A}$ and $\mathbf{B}$ are vectors or scalars. In the former case, the model will be referred to as the Vector VECH, and in the latter the Scalar VECH. For example, the Scalar VECH implies that there is a specific constant long-run covariance matrix that is captured by $\mathbf{CC}'$, but the way the covariance matrix is updated is the same for all variances and covariances, that is, using the constants a and b. Except for the long-run covariances, this would mean that all elements follow the same GARCH process.

2.3.3 DCC Models

Contrary to the VECH models (and many other multivariate GARCH models), the dynamic conditional correlation (DCC) class of Engle (2002) splits the modeling of $\mathbf{H}_t$ in two distinct parts: volatility and correlation. Indeed, with the DCC model, one specifies the volatility dynamics of each market and then defines the correlation process. It turns out that this two-step modeling process can also be applied to estimation, which is a tremendous advantage over all other multivariate models (including regime-switching).

Let us rewrite the covariance matrix $\mathbf{H}_t$ as

$$\mathbf{H}_t = \mathbf{D}_t\mathbf{R}_t\mathbf{D}_t, \tag{2.1}$$

where $\mathbf{R}_t$ is the time-varying correlation matrix and $\mathbf{D}_t = \text{diag}(\sigma_t^{(1)}, \ldots, \sigma_t^{(N)})$. The representation of $\sigma_t^{(i)}$ depends on the univariate model used for the volatility. For example, one can choose any univariate GARCH models for $\sigma_t^{(i)}$.[8]

The remaining task is to model the time-varying dynamics of the correlation. Engle's idea is that modeling the conditional correlation of returns is equivalent to modeling the conditional covariance of standardized errors, defined as

$$\varepsilon_t = \mathbf{D}_t^{-1}(\mathbf{y}_t - \boldsymbol{\mu}).$$

Now, define the covariance matrix of ε_t as $\mathbf{Q}_t$ (we will specify its dynamics later), so that

$$\mathbf{R}_t = \boldsymbol{\Delta}_t^{-1}\mathbf{Q}_t\boldsymbol{\Delta}_t^{-1}$$

with $\boldsymbol{\Delta}_t = \text{diag}(\sqrt{Q_t^{(1,1)}}, \ldots, \sqrt{Q_t^{(N,N)}})$. We see that the dynamics of correlation will depend on the choice of dynamics for $\mathbf{Q}_t$.

[7] In the most complex specification of the model, each covariance element is a function of the lagged volatilities and errors of each market. In the literature the diagonal VECH (or simply VECH in most cases) is a GARCH(1,1) process for each element of the covariance matrix, without accounting for cross-sectional effects.

[8] The model chosen for $\sigma_t^{(i)}$ does not have to be the same as for $\sigma_t^{(j)}$ ($i \neq j$). An exponential GARCH(1,1) and a threshold GARCH(2,1) could be selected for specific markets if it is believed to be more appropriate. Unless stated otherwise, GARCH(1,1) models are used.

Engle (2002) proposes the following two models for the correlation. The mean-reverting (MR) model is a scalar GARCH(1,1) model such that

$$\mathbf{Q}_t = (1 - \alpha - \beta)\mathbf{S} + \alpha(\varepsilon_{t-1}\varepsilon'_{t-1}) + \beta\mathbf{Q}_{t-1}, \tag{2.2}$$

where $\mathbf{S}$ is the sample correlation matrix. The integrated (INT) model is such that $\alpha = 1 - \lambda$ and $\beta = \lambda$.

The advantages of the model are obvious: it allows a very parsimonious representation of the correlation dynamics and the possibility of estimating the model in two steps. According to Engle (2002), it will provide consistent estimates as long as the volatility parameters are consistent. However, they will not be efficient. To estimate a DCC model with N assets, one would perform N separate univariate GARCH estimations. Given the residuals of the first step, an $N + 1$–th optimization is necessary to get the parameters of the matrices $\mathbf{Q}_t$.

Finally, we adopt the following convention. When a model is estimated in two steps using Engle's approach, the prefix DCC- will be added to the name of the model. Thus, the DCC-MR model corresponds to equation (2.2), while its integrated version will be denoted by DCC-INT.

2.3.3 Constant Correlation (CCORR) Model

One very important model can be found using the DCC class of Engle (2002): the constant correlation (CCORR) model of Bollerslev (1990). In the original specification of the model, a GARCH(1,1) model is used for each volatility, and a constant (over time) correlation matrix $\mathbf{R}_t = \mathbf{R}$ is assumed. When both correlation and volatility parameters are estimated at the same time, the model is referred to as the CCORR model because it matches the earliest definition presented by Bollerslev (1990).

However, one can also view the CCORR model as in equation (2.2) using $\alpha = \beta = 0$. In this case it provides theoretical justification to estimate the model in two steps. When it is estimated in this fashion, the model will be designated as DCC-CCORR.

One might wonder if a constant correlation assumption is appropriate, especially with regard to the results of Tse (2000). In fact, some financial time series do feature constant correlation (futures and exchange rates). Moreover, monthly data, which are frequently used with investment guarantees modeling, will not necessarily have the same behavior as daily data because of the aggregation of daily returns. Consequently this model will act as a benchmark to check the presence of time-varying correlation in monthly data.

2.3.4 Other Models

Another model that has also been considered in the numerical analysis is the BEKK model of Baba et al. (published in Engle and Kroner 1995). Its most important features are (1) guaranteed PSD of the covariance matrix, (2) covariances account for cross-sectional effects, and (3) more parsimonious than the full VECH model. However, in our numerical analysis the second feature is not worth the added complexity, in terms of both parsimony and adequation, so the BEKK model will not be further discussed.

Other important contributions to the literature of multivariate GARCH models are the factor ARCH (FARCH) model of Engle, Ng, and Rothschild (1990) and the asymmetric version of the VECH, BEKK, FARCH, and CCORR models presented in Kroner and Ng (1998). Because of its implicit dependence upon beta pricing models, we will not use the factor ARCH model in our context. Moreover, asymmetric models add many parameters that are very likely to be rejected on the basis of parsimony. For the sake of brevity, we will thus limit our analysis to the models presented.

3. Estimation

In this section we briefly discuss the estimation of the models presented earlier using the maximum likelihood estimation (MLE) technique.

3.1 Regime-Switching Models

Let us first discuss the estimation of multivariate regime-switching models given that it has been written as a global regime model with K states.[9] In this case one can directly use the loglikelihood function presented in Hamilton (1994) and Hardy (2001, 2003), where the univariate p.d.f. of $y_t|y_{t-1}, \ldots, y_0$ is replaced by the multivariate p.d.f. of $\mathbf{y}_t|\mathbf{y}_{t-1}, \ldots, \mathbf{y}_0$.

3.2 Multivariate GARCH Models

The estimation of multivariate GARCH models depends on how the model has been built. For the VECH (including its vector and scalar versions) and CCORR models, it has been argued that all parameters have to be found simultaneously. In this case one uses the fact that $\mathbf{y}_t|\mathcal{F}_{t-1}$ is normally distributed with mean μ and (conditional) variance $\mathbf{H}_t$.

The DCC class of models is built and estimated in two steps. Thus, one finds the parameters of the volatility processes with usual techniques and then, given the previous estimation, the likelihood function is maximized with respect to the correlation parameters. For more details about the estimation procedure of DCC models, one should refer to Engle (2002).

4. CTE PROVISIONS

We now turn our attention to the computation of the conditional tail expectation (CTE) in a framework with multiple assets. We will first derive this quantity under a regime-switching setting, and then we will discuss simulation and approximation approaches for GARCH models.

4.1 Regime-Switching Models

One major advantage of regime-switching models over GARCH models is the existence of a closed-form solution for the provision. Indeed, Hardy (2001, 2003) provides an analytic formula for the CTE with a univariate RSLN with two regimes. Fortunately this formula will also apply in a multivariate framework.

Denote by $\mathbf{w}$ the allocation to each of the N assets, with $\Sigma_{i=1}^{N} w_i = 1$ and $\mathbf{w}' = [w_1, \ldots, w_N]$. Given the current regime, or equivalently, the number of visits to each state, the scalar $y_t^{(P)}$, which is the portfolio return with

$$y_t^{(P)} \equiv \mathbf{w}'\mathbf{y}_t, \qquad t = 0, 1, 2, \ldots \tag{4.1}$$

has a normal distribution. This is equivalent to rebalancing the portfolio at each period.[10] Thus, in a global regime with two states, the CTE is given by equation (9.18) of Hardy (2003), with the modification that

$$\begin{aligned} \mu^*(k) &= k\mathbf{w}'\boldsymbol{\mu}_1 + (n - k)\mathbf{w}'\boldsymbol{\mu}_2, \\ \sigma^*(k) &= \sqrt{k\mathbf{w}'\Sigma_1\mathbf{w} + (n - k)\mathbf{w}'\Sigma_2\mathbf{w}}, \end{aligned} \tag{4.2}$$

where n is the total number of periods considered (or the horizon of the CTE).[11]

---

[9] This formulation includes the possible addition of one or two regimes to account for local dynamics that do not affect other markets.

[10] If the rebalancing of the portfolio is done frequently, so that the periodic return on each asset $y_t^{(i)}$ is small enough, then using the expansion of the exponential, we have

$$\sum_{i=1}^{N} w_i \exp(y_t^{(i)}) \approx \exp\left(\sum_{i=1}^{N} w_i y_t^{(i)}\right).$$

[11] The formula for three regimes can be derived in a similar fashion, conditioning on the number of visits to each regime. For the sake of brevity, the derivation of the CTE for three regimes will not be presented in this paper. If one wants to see the details of the formula, it is possible to contact the authors of this paper; a document especially designed for this purpose is readily available.

4.2 GARCH Models

The computation of the CTE generally relies on the tractability of the r.v. $\Sigma_{j=1}^{n} y_{T+j}^{(P)}|\mathcal{F}_T$. Even though $y_{T+1}^{(P)}$ given $\mathcal{F}_T$ is normal, $y_{T+2}^{(P)}, \ldots, y_{T+n}^{(P)}$ clearly are not so that we cannot compute the CTE with a closed-form solution. In this case we have to rely on simulation or an approximation.

Let us begin by presenting the approximation to the true value of the CTE. We can assume that

$$(\sigma_{t+k}^{(P)})^2 \equiv \mathbf{w}'E[\mathbf{H}_{t+k}|\mathcal{F}_t]\mathbf{w}, \qquad k = 1, 2, \ldots,$$

which is equivalent to using a deterministic volatility pattern, equal to the predicted volatility given by the model. The CTE would then be computed using the CTE of a lognormal distribution with independent but not identically distributed returns. Using Hardy (2003)'s equation (9.17), we get that the CTE after n periods is

$$\widehat{CTE}_\alpha[L] \approx G - \frac{\exp(n\mu^{(P)} + \frac{1}{2}\Sigma_{k=1}^{n}(\sigma_{t+k}^{(P)})^2)}{1 - \alpha}\,\Phi\left(-z_\alpha - \sqrt{\sum_{k=1}^{n}(\sigma_{t+k}^{(P)})^2}\right),$$

where $\mu^{(P)} \equiv \mathbf{w}'\boldsymbol{\mu}$. Note that we did not discount the CTE and we excluded expenses.

This approximation should be lower than the real CTE value because extreme returns, which would be provided by scenarios with high variances, will not occur as often because the variances are smoothed by their expectation. However, the volatilities generated by GARCH models tend to stabilize very quickly, so this approximation should be quite reliable. Because the latter approximation is a lower bound on the value of the provision, simulation remains a good solution in this case, as long as the number of paths used is large enough, say, at least 10,000.

Even though the size of the error made with the approximation is hard to define, this approximation can be of great value to actuaries because of its ease and quickness of computation. A comparison of the approximation with the true simulated value done on a very frequent basis may help the actuary determine the importance of the error and adjust the approximated value accordingly. Moreover, the approximation can be used in combination with simulation in a control variate setting to reduce the number of simulations necessary or to increase precision in Monte Carlo estimates.

5. NUMERICAL EXAMPLE

We now turn to the application of the models that we have presented earlier. As a first step, we will present the dataset we will use throughout this example. Given the dataset, we estimate various regime-switching models along with the GARCH models presented. Their fit will be analyzed with selection criteria and validation tests. Then, for a few selected models, we will compute the CTE provision and further analyze the differences obtained between models.

It is important to note that given the constraints of a paper, it is impossible to present the results of every test applied to each possible model and, similarly, compute the CTE for each model and portfolio considered. This is why we adopt a parsimonious approach in this paper in which models are dropped after each step. Thus, this example is more an illustration of how to use the tests than a formal and exhaustive model selection process. Moreover, nothing replaces the actuary's judgment when it comes to determining the required capital, and the results presented in this paper are no exception.

5.1 Data

We have chosen a dataset for which the markets considered would be important from the point of view of a Canadian insurance company. We used monthly returns data from January 1956 to September 2005 for the following markets: (1) Canada (S&P TSX total return index), (2) United States (S&P 500 total return index), (3) United Kingdom, (4) Japan (TOPIX index).

The data for Canadian and Japanese markets come from Bloomberg, while U.S. data come from both Bloomberg and Ibbotson. The dataset for the United Kingdom has been built from the FTSE All-Shares

Table 1
Annualized Descriptive Statistics for Monthly Historical Returns

	Canada	U.S.	U.K.	Japan
Mean	9.37%	10.02%	7.22%	7.68%
Volatility	15.60%	14.56%	17.39%	18.61%
Asymmetry	−0.90	−0.60	−0.39	0.07
Kurtosis	3.61	2.54	1.37	8.68

index for the period from April 1962 to September 2005 (Bloomberg). Prior to 1962 we have used the Actuaries Investment Index from the Institute of Actuaries.

Historical data cover the same period for all four indices. No adjustments are made for currency effect because the index performances are measured in their local currencies. We show in Table 1 descriptive statistics on the monthly continuous returns. Note that because Canadian and U.S. market data include dividends, we expect the returns to be higher on these markets. Because dependence plays an important role in a multivariate model, we show in Table 2 the correlation between the different pairs of markets.

For the computation of the CTE provisions, we will consider nine different portfolios, comprised of five mixes between the markets presented. Apart from a 100% allocation in each market, the portfolios analyzed are the following: (1) 50% U.S.–50% Canada, (2) 50% U.S.–50% U.K., (3) 50% U.S.–50% Japan, (4) 33.3% U.S.–33.3% Canada–33.3% U.K., (5) 25% U.S.–25% Canada–25% U.K.–25% Japan. We assume throughout the example a monthly rebalancing as formalized by equation (4.1).

5.2 Estimation and Selection

5.2.1 Regime-Switching Models

Let us estimate as a first step a set of regime-switching models that will differ on the number of regimes and the number of correlation matrices. The models are denoted as MRSLN(K, n_R), where K is the number of regimes and n_R is the number of correlation matrices used in the estimation process. To differentiate models on the basis of parsimony, we use the likelihood ratio test (LRT), in addition to the Akaike (AIC) and Schwartz-Bayes (SBC) Information Criteria.[12] The number of parameters in the model is denoted by $n^{(P)}$.

The estimation results are shown in Table 3. According to the AIC parsimony criterion, the addition of a third regime to the basic MRSLN(2,1) model is significant. This is also verified with the likelihood ratio test (LRT), when the MRSLN(2,1) model is picked as the base model. Note that all models are

Table 2
Empirical Correlation Matrix for Four Markets Presented

	Canada	U.S.	U.K.	Japan
Canada	1	0.7773	0.5231	0.2995
U.S.	0.7773	1	0.5515	0.2929
U.K.	0.5231	0.5515	1	0.2761
Japan	0.2995	0.2929	0.2761	1

[12] More details about these criteria can be found in Klugman, Panjer, and Willmot (2004) or Hardy (2003).

Table 3
Estimation Results for Regime-Switching Models

Model	$n^{(P)}$	Log-likelihood	AIC	SBC	LRT
MRSLN(2,1)	24	4,423.72	4,399.72	**4,347.04**	n.a.
MRSLN(2,2)	30	4,426.64	4,396.64	4,330.79	0.4414
MRSLN(3,1)	36	4,457.81	**4,421.81**	4,342.79	$<\mathbf{10}^{-7}$
MRSLN(3,3)	48	4,463.14	4,415.14	4,309.77	$<\mathbf{10}^{-7}$

specific cases of the MRSLN(3,3), making it possible to use the LRT. When the SBC is computed, none of the improvements are important.

In light of Table 3, there is no sign of time-varying correlation in the monthly data because no parsimony criterion picked a model with more than one correlation matrix.

In conclusion, on the basis of Table 3, the MRSLN(2,1) and MRSLN(3,1) models are the most parsimonious. We choose those models to be further analyzed.

5.2.2 GARCH Models

We have also performed a maximum likelihood estimation on numerous multivariate GARCH models. The results are presented in Table 4. We have split the table in two distinct parts. The first portion of the table refers to models that have been estimated in one step, that is, a full optimization, while the second part of the table are models estimated in two steps, consistent with Engle's DCC class of models. In the latter class of models, a GARCH(1,1) model was assumed for each of the four markets.

Let us begin by analyzing the first part of Table 4. According to the AIC criterion, the most parsimonious model is the Vector VECH model. A notable runner-up is the CCORR model. Penalizing further for the additional parameters of the model, we obtain that the Scalar VECH model is also interesting. It seems to be a very strong assumption, but when we compare the estimated parameters of the VECH and Scalar VECH, such an assumption looks plausible.

We have chosen to present the results of the DCC class of models separately because of the two-step estimation process. We see that the DCC-INT model does a very good job in terms of parsimony with both criteria. Note that even though the model has 23 parameters, those are estimated in two steps, which considerably reduces the workload. The DCC-INT model will clearly be chosen because it is simple to estimate and provides results advantageously comparable to the models estimated in one step.

In conclusion, on the basis of Table 4, the Vector VECH, Scalar VECH, CCORR, and DCC-INT models are the most parsimonious GARCH models, and we choose those to be further analyzed.

Moreover, when we compare the results of Table 3 with Table 4, using an equivalent number of parameters or penalized likelihood criteria, we see that when we consider both families of multivariate models, the GARCH models have in general a better overall fit than the regime-switching models.

Table 4
Estimation Results for GARCH Models

Model	$n^{(P)}$	Log-likelihood	AIC	SBC
VECH	34	4,423.367	4,389.36	4,314.73
Vector VECH	22	4,424.746	**4,402.74**	4,354.45
Scalar VECH	16	4,415.596	4,399.59	**4,364.47**
BEKK	46	4,444.34	4,398.34	4,297.36
CCORR	22	4,424.38	*4,402.38*	4,354.08
DCC-INT	23	4,433.468	**4,410.46**	**4,359.98**
DCC-MR	24	4,434.135	*4,410.13*	4,357.45
DCC-CCORR	22	4,413.76	4,391.76	4,343.46

5.3 Validation

A model selection process based upon criteria such as the LRT, AIC, and SBC is a good initial step when it comes to determining whether the added complexity is worth the gain in fit. However, it does not tell whether the fit in itself is good or not. The second set of tests presented next will focus on how good the model approximates the data. With a multivariate model, what matters is the adequation of the heteroskedasticity in variance and covariance as well as the tails of the returns distributions. This last feature is very important for the determination of conservative and realistic provisions for investment guarantees. Before we begin, we need to present how to compute the standardized residuals in both families of models. This is a crucial input for tests presented next.

Assume that the parameters of the model have been previously estimated. Let us denote by $\tilde{\varepsilon}_t' = [\tilde{\varepsilon}_t^{(1)}, \ldots, \tilde{\varepsilon}_t^{(N)}]$ the vector of standardized residuals at time t. The standardized residuals for a multivariate GARCH model is obtained by computing

$$\tilde{\varepsilon}_t = \text{chol}(\hat{\mathbf{H}}_t)(\mathbf{y}_t - \hat{\boldsymbol{\mu}}),$$

where chol is the lower triangular matrix obtained from a Cholesky decomposition of $\hat{\mathbf{H}}_t$, and $\hat{\mathbf{H}}_t$ is the conditional covariance matrix computed with the estimated parameters $\hat{\boldsymbol{\theta}}_t$.

For a regime-switching model, the approach is similar, but we have to weigh the result using the conditional probability of being in each regime. This is described in Hardy, Freeland, and Till (2006), and in the context of a multivariate model, we get[13]

$$\tilde{\varepsilon}_t = \sum_{k=1}^{K} \text{chol}(\hat{\boldsymbol{\Sigma}}_k)(\mathbf{y}_t - \hat{\boldsymbol{\mu}}_k)\Pr(\rho_t = k|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0),$$

where $\hat{\boldsymbol{\mu}}_k$ and $\hat{\boldsymbol{\Sigma}}_k$ are the estimated mean and covariance matrix in regime k and $\Pr(\rho_t = k| \ldots)$ is the conditional probability of being in regime k, given the returns observed up to time t.

It should be noted that the standardized residuals computed in this fashion smooth the effects of regime uncertainty. When $\Pr(\rho_t = k| \ldots)$ is close to 1 or 0, standardized residuals will behave very closely to a normal distribution, and failure or acceptance of the following tests will be reliable. Otherwise, the failure of normality or heteroskedasticity tests in the case of the regime-switching model may be due in part to regime uncertainty.

5.3.1 Normality Tests

According to statistical theory, the standardized residuals of a well-fitted model should be i.i.d. normally distributed random variables for Gaussian innovations. As a first step, we have computed the Jarque-Bera and the Shapiro-Wilk tests, on the standardized residuals of each asset, for each of the previously selected models. Moreover, we have also added the MSRLN(3,3) and BEKK models in the analysis to check if these complex models rejected earlier would perform better in validation tests.

Table 5 shows the p-values of the Jarque-Bera test. P-values lower than 5% are highlighted in bold, whereas values between 5% and 10% are shown in italic. The results are obvious: practically all GARCH models presented fail the Jarque-Bera test at a 5% significance level, while regime-switching models perform better, especially with the North American markets. The results from the Shapiro-Wilk are somewhat different but still support the evidence shown by the Jarque-Bera test. Thus, the results of the Shapiro-Wilk test are not presented for the sake of brevity.

A final word with respect to Table 5 is that no regime-switching model does remarkably better than the others, in terms of the normality tests. In this case we may choose the most parsimonious model, the MRSLN(2,1). Moreover, all GARCH models do poorly with the Canadian market and only slightly better with the American market.

---

[13] Hardy, Freeland, and Till (2006) also present the zero-one standardized residuals, but since they give very similar results, we will use only the one presented.

Table 5
***P*-Values of the Jarque-Bera Test for Normality Applied on Each Market**

Jarque-Bera	S&P TSX	S&P 500	Topix	FTSE
MRSLN(2,1)	0.4864	0.1547	$<\mathbf{10}^{-4}$	$<\mathbf{10}^{-4}$
MRSLN(3,1)	0.5168	0.1345	**0.0002**	0.1274
MRSLN(3,3)	0.7357	*0.0541*	**0.0045**	0.1069
CCORR	$<\mathbf{10}^{-4}$	*0.0874*	$<\mathbf{10}^{-4}$	$<\mathbf{10}^{-4}$
Vector VECH	$<\mathbf{10}^{-4}$	**0.0049**	$<\mathbf{10}^{-4}$	$<\mathbf{10}^{-4}$
Scalar VECH	$<\mathbf{10}^{-4}$	**0.0051**	$<\mathbf{10}^{-4}$	$<\mathbf{10}^{-4}$
DCC-INT	$<\mathbf{10}^{-4}$	**0.0443**	$<\mathbf{10}^{-4}$	$<\mathbf{10}^{-4}$
BEKK	$<\mathbf{10}^{-4}$	*0.0654*	$<\mathbf{10}^{-4}$	$<\mathbf{10}^{-4}$

A model that poorly fits the tails of a distribution of returns will most likely fail those normality tests. To check this, we have plotted the univariate Q-Q plots for the MRSLN(2,1) and the CCORR models in Figures 1 and 2 to help understand where GARCH models fail. Note that the Q-Q plots are representative of their own class of models. Take, for example, the case of the S&P TSX market. There are a few deviations from the dotted line with the regime-switching model, while deviations from the normal quantiles are much more significant with the multivariate GARCH model. Thus, the failure to

Figure 1
Q-Q Plot of Standardized Residuals for the MRSLN(2,1) Model for Each Market

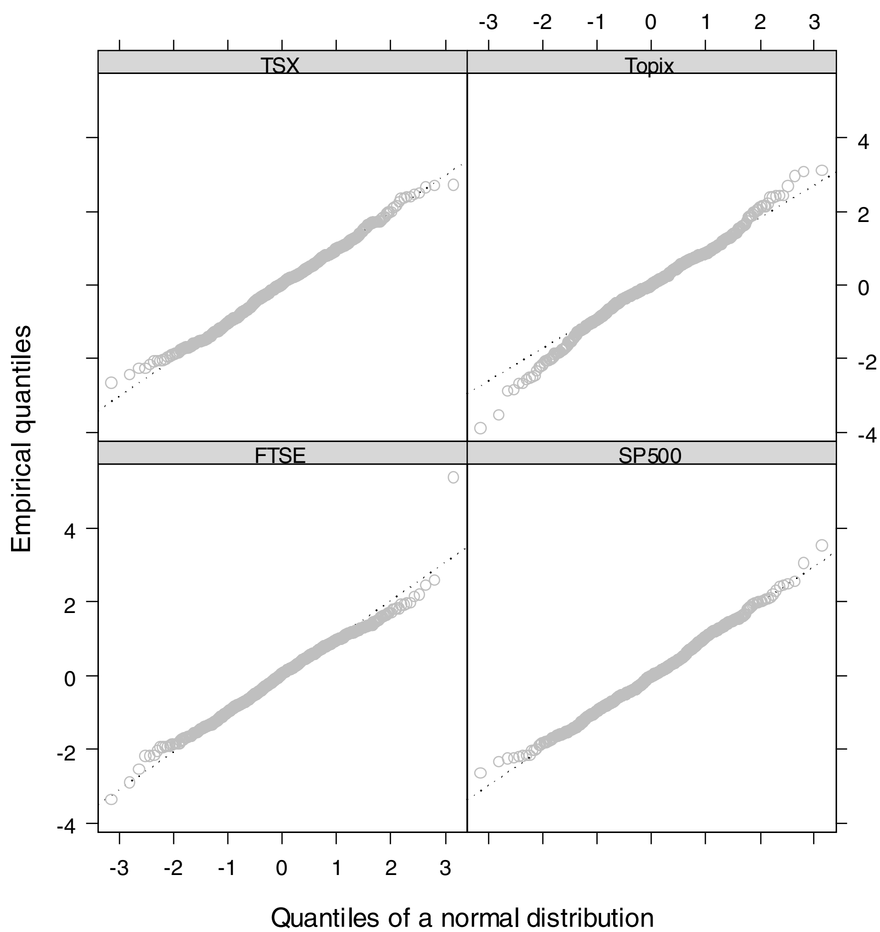

Figure 2
Q-Q Plot of Standardized Residuals for the Constant Correlation GARCH (CCORR) Model for Each Market

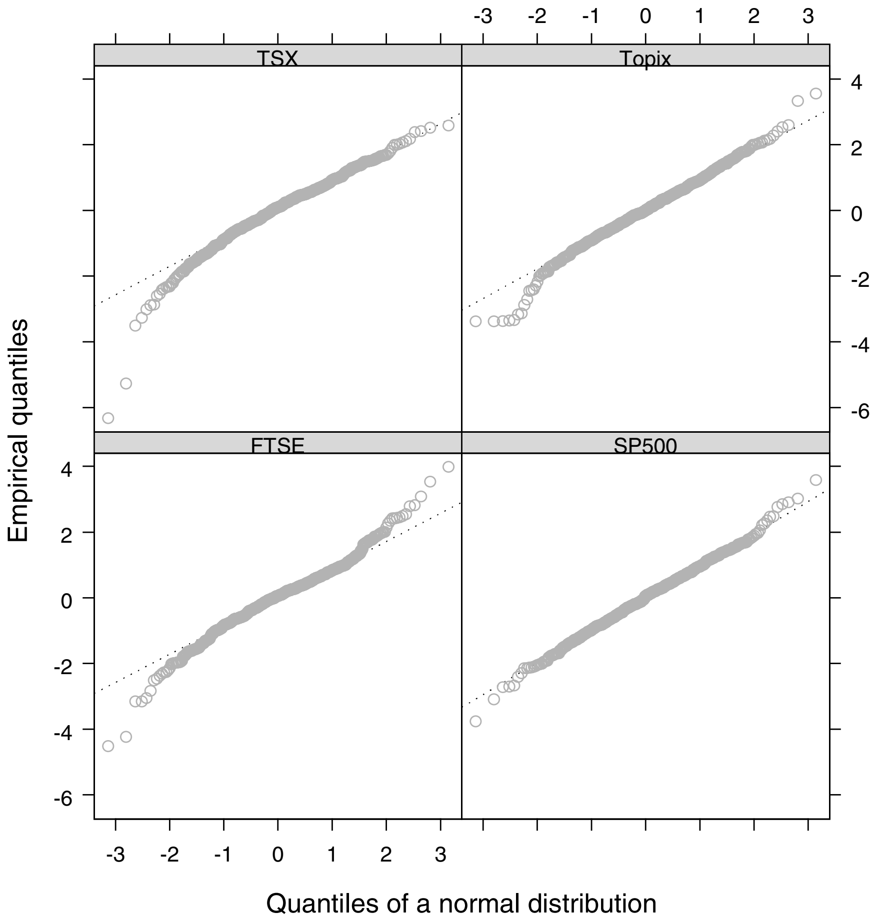

meet normality for GARCH models is mainly due to a lack of fit in the left tail of the returns distributions. This is consistent with Hardy, Freeland, and Till (2006).

A multivariate extension of a univariate Q-Q plot also exists that takes into account the fit of all markets. To compute this type of Q-Q plot, we rely on the fact that if $\tilde{\varepsilon}_t$ is indeed a vector of i.i.d. N (0, 1) r.v., then $\tilde{\varepsilon}_t'\tilde{\varepsilon}_t \sim \chi^2_N$. Thus, it suffices to compute the T values of $\tilde{\varepsilon}_t'\tilde{\varepsilon}_t$ and compare them to their corresponding χ^2_N quantile. This has been plotted in Figure 3 for the constant correlation model. The fact that the tails of the returns distribution were poorly fitted in each market clearly shows in this graph as we have a significant deviation over the 45° line. This is a sign of a misestimation of the tails.

5.3.2 Tests for Heteroskedasticity

We now want to determine if the selected multivariate models are capable to represent the time-varying dynamics of variance and covariance. In this case the standardized residuals should no longer show signs of heteroskedasticity. To do so, we will use specific tests that check for heteroskedasticity in the data. We will use three distinct tests: an ARCH test applied on each market, Engle's multivariate extension to the basic univariate ARCH test, and Hosking's multivariate portmanteau test.

Figure 3

Multivariate Q-Q Plot of Standardized Residuals for the Constant Correlation (CCORR) GARCH Model

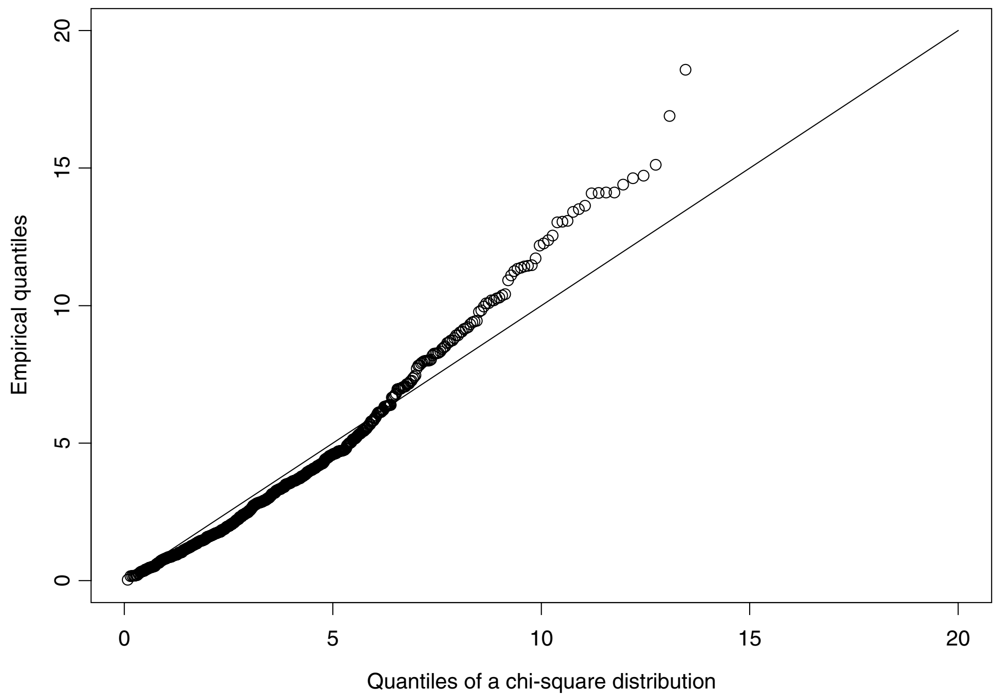

We begin with an application of the ARCH test to the standardized residuals.[14] This will allow us to determine whether there is remaining heteroskedasticity in the variances of each market. Heteroskedasticity in the covariances is not taken into account with this test. A univariate ARCH test is a check for zero coefficients in the regression of $(\tilde{\varepsilon}_t^{(i)})^2$ on $(\tilde{\varepsilon}_{t-1}^{(i)})^2, \ldots, (\tilde{\varepsilon}_{t-m}^{(i)})^2$.

Table 6 shows the p-value of the ARCH test on each market, for each model, for six lags only. The results for 12 lags have been computed but are very similar to those with six lags and, thus, are not presented. In general, we see that multivariate GARCH models better reflect the time-varying dynamics of variance than do the regime-switching models. The latter models do reasonably well at reflecting the variance dynamics of the American market. However, it seems to impose an inadequate variance

Table 6

ARCH Test on Each Market for Selected Models

ARCH test (6 lags)	S&P TSX	S&P 500	Topix	FTSE
MRSLN(2,1)	**0.0009**	0.1824	**0.0002**	**0.0005**
MRSLN(3,1)	**0.0130**	0.1739	**0.0008**	0.4040
MRSLN(3,3)	**0.0068**	0.3300	**0.0009**	0.3095
CCORR	0.9979	0.8420	0.6192	*0.0901*
Vector VECH	0.9871	0.7537	*0.0955*	**0.0336**
Scalar VECH	0.9970	0.6297	0.3149	**0.0039**
DCC-INT	0.9955	0.8684	0.6975	0.1052
BEKK	0.9985	0.4130	**0.0388**	*0.0550*

[14] More details about ARCH tests can be found in Bollerslev, Chou, and Kroner (1992), Bollerslev, Engle, and Nelson (1994), Campbell, Lo, and MacKinlay (1997), and Tsay (2005).

Table 7
Hosking's Multivariate Portmanteau Test Applied to Squared Standardized Residuals

Hosking	6 Lags	12 Lags
MRSLN(2,1)	$<\mathbf{10^{-4}}$	$<\mathbf{10^{-4}}$
MRSLN(3,1)	**0.0088**	0.1022
MRSLN(3,3)	**0.0059**	**0.0474**
CCORR	0.7178	0.9433
Vector VECH	**0.0038**	*0.0862*
Scalar VECH	**0.0001**	**0.0054**
DCC-INT	0.2203	0.7637
BEKK	**0.0075**	**0.0435**

process for the Canadian and foreign markets, so that regime-switching models fail the ARCH test with those markets. Multivariate GARCH models do very well with North American and Japanese markets, while the FTSE generates troubles for all GARCH models.

Although the ARCH test is simple to use, its application in a multivariate setting might not detect if the time changes of covariance are well represented. We have to turn to especially designed multivariate tests. Hosking (1980) and Engle (2002) both introduced multivariate tests that can be applied to assess the quality of the fit of a multivariate model, in terms of the covariance dynamics. Hosking has proposed a multivariate extension of the Ljung-Box test that we will apply on the squared standardized residuals, in order to check if there is remaining heteroskedasticity. Engle discusses an extension of the univariate ARCH test that accounts for cross-products.

Applying Hosking (1980) to the squared standardized residuals, we obtain Table 7. We see that when we account for the time-varying dynamics of covariance, at least two of the three regime-switching models presented fail Hosking's test, showing that there remains heteroskedasticity in the variance and covariance. This might be because of its poor performance with the ARCH test on foreign markets and its lower p-values for the North American markets. The CCORR and DCC-INT GARCH models have a high significance level, showing that these models adequately represent the variance and covariance dynamics. The analysis here would certainly benefit from another test that has a similar purpose. For this reason we will perform the ARCH test of Engle (2002), applied to a multivariate context.

Engle (2002)'s idea to extend the univariate ARCH test to a multivariate framework is to add the cross-products of standardized residuals in the regression. We have used exactly the same regression as in Engle; the results are shown in Table 8. We see that adding the cross-products of standardized residuals in (2002) the ARCH test significantly increases the number of rejections, showing that many models do not adequately represent the covariance dynamics. Consistent with Table 6, most regime-switching models still have heteroskedasticity in the covariances. Multivariate GARCH models do better

Table 8
Multivariate ARCH Test on Each Market for Selected Models

Engle (5 Lags)	S&P TSX	S&P 500	Topix	FTSE
MRSLN(2,1)	**0.0103**	0.6388	**0.0020**	$<\mathbf{10^{-4}}$
MRSLN(3,1)	**0.0134**	0.8472	**0.0079**	0.2188
MRSLN(3,3)	**0.0051**	0.9016	**0.0086**	0.1794
CCORR	0.6346	0.4702	0.5034	**0.0280**
Vector VECH	0.7130	*0.0747*	**0.0317**	**0.0078**
Scalar VECH	0.7367	**0.0201**	0.1030	**0.0001**
DCC-INT	0.6929	0.4705	0.4570	0.1360
BEKK	0.5239	**0.0339**	**0.0193**	**0.0130**

than regime-switching models in this test. Engle's test confirms the results obtained with Hosking's test, showing which market results in an overall failure of the latter.

One might argue that regime uncertainty might have played a role in the failure of regime-switching models to heteroskedasticity tests. We believe it is not likely to be the case for several reasons. First, the probabilities of being in each regime are either very small or very large so that the test should be reasonably good. As a result, normality tests were not rejected for most regime-switching models. Moreover, the fact that we could not reject heteroskedasticity in the U.S. subsample illustrates the relative reliability of the test for this specific dataset. Finally, dynamics of variance account for an important part of the total loglikelihood, and when we compare likelihood-based tests (AIC, SBC, or direct comparison of log-likelihoods), GARCH models do have a better overall fit.

In light of the results obtained with normality and heteroskedasticity tests, the MRSLN(2,1), MRSLN(3,1), CCORR, and DCC-INT models are excellent candidates for modeling purposes. More complex models such as the MRSLN(3,3) and BEKK do not provide an important improvement in the quality of the fit to be further considered. However, those two models will also be used in conjunction with the computation of the CTE provision to check whether they provide important differences.

5.4 CTE Provisions

In this section we will compute the CTE provision for some selected models in order to measure the impact of the multivariate model on the provision level. Moreover, we will be able to assess the effects of using data from other markets on univariate provisions, that is, portfolios with a 100% allocation in one market. Throughout this example we assume that the guarantee is 100% of the initial investment and funded using a 3% annual management expense ratio (MER). In computing the provision, we have used the actuarial approach, that is, the reserve is set up by projecting future losses without accounting for the assets backing the liability. For illustration purposes, we used a 95% level 10-year CTE without mortality and discounting. Finally, unless stated otherwise, the CTEs for regime-switching models have been computed using the closed-form solution of equation (9.18) of Hardy (2003) with the modification of (4.2). The CTEs for GARCH models have been calculated using 10,000 simulations.

Table 9 shows the 95%, 10-year CTE for allocations of 100% in one of the markets. It shows that there is a lot of variation between the CTE provisions. For an allocation of 100% in the Canadian (U.S.) market, the lowest provision is 30% (22%) while the largest is 49% (36%), which is significantly more. Variations are fortunately much less important for the other two markets, as the difference between the minimum and maximum (range) provisions is either 4.5% or 7% (in absolute values). A large difference between classes of models could have been anticipated, but even within each class of models, provisions can change a lot. Differences of approximately 5% (in absolute values) between very similar models are not rare.

When we compare the provisions generated by both families of models, regime-switching models provide much higher CTE values for North American markets while both families are equivalent in

Table 9

CTE Provision for Different Selected Models, with Allocation of 100% in Each Market

CTE	S&P TSX	S&P 500	Topix	FTSE
MRSLN(2,1)	43.90%	32.00%	54.00%	56.20%
MRSLN(3,1)	47.70	34.70	54.50	58.80
MRSLN(3,3)	48.80	35.90	54.80	60.40
CCORR	32.88	24.67	51.51	55.70
DCC-INT	35.07	23.55	53.54	53.82
BEKK	29.71	22.33	55.91	60.18
Minimum	29.71	22.33	51.51	53.82
Maximum	48.80	35.90	55.91	60.40

Table 10
CTE Provision for Different Selected Models and Portfolios

CTE	Ptf 1	Ptf 2	Ptf 3	Ptf 4	Ptf 5
MRSLN(2,1)	35.20%	38.60%	33.00%	36.80%	33.80%
MRSLN(3,1)	38.70	41.80	34.30	40.50	36.80
MRSLN(3,3)	37.10	42.30	32.10	40.10	35.10
CCORR	24.70	33.00	27.28	27.61	25.18
DCC-INT	24.97	32.53	25.80	26.77	25.97
BEKK	21.84	34.00	26.69	26.91	24.94
Minimum	21.84	32.53	25.80	26.77	24.94
Maximum	38.70	42.30	34.30	40.50	36.80

terms of CTE levels for the Japanese and British markets. The global regime approach seems to be perfectly suited for North American markets, capturing most of the common dynamics reflected on the provisions.

Table 10 shows the 95% 10-year CTE provisions for the five main portfolios considered in the paper. Although there are still differences between the models, the provisions are much more stable among models, with the notable exception that GARCH provisions are generally lower. This is mainly because all portfolios have an important share in either two North American markets, driving the provisions up for regime-switching models.

6. CONCLUSION

We have presented in this paper multivariate time series models that can be used in the context of valuation of investment guarantees. Two families of models have been considered: multivariate GARCH models (VECH, CCORR, and DCC) and regime-switching models. We have estimated those models to a four-markets dataset (Canada, U.S., U.K., and Japan) using typical estimation techniques (maximum likelihood), and we have computed the CTE provisions for these models.

In general, we have found that multivariate GARCH models had a better overall fit than regime-switching models. For some models it was true even if we did not account for some parsimony criterion. The better fit of the GARCH models was also reflected with tests of heteroskedasticity of variance and covariance. Indeed, it showed that, in general, the dynamics of variance and covariance observed on the markets was better reflected by multivariate GARCH models. However, the latter models failed to adequately represent the fat tails of the distribution of returns, so that GARCH models failed most normality tests. We have seen that this was not the case for regime-switching models because they showed very interesting Q-Q plots, and we could not reject the normality assumption of the residuals for most markets and models.

Looking at the models more specifically, we have found that the most parsimonious models also do reasonably well in terms of the quality of the fit. They are the two (or three) regime multivariate regime-switching model with one correlation matrix (MRSLN(2,1), MRSLN(3,1)), the constant correlation GARCH (CCORR), and the DCC integrated model (DCC-INT).

The good performance of GARCH models with respect to likelihood criteria comes from the fact that covariance dynamics are adequately represented, at the cost of sacrificing the left tail. Since the latter accounts for a very small portion of the sample, the regime-switching models had a disappointing overall fit. Model selection should depend on what matters: overall fit versus the fit of the tails. Multivariate GARCH models clearly have a better overall fit, but multivariate regime-switching models better represent the fat tails of the returns distribution. This is reflected in the levels of the CTE provisions because, in general, provisions computed with multivariate regime-switching models are higher than with multivariate GARCH models.

Considering the results that we have obtained with this specific dataset, we cannot recommend a specific model for all purposes because it mostly depends on how the model will be used. There is a tradeoff between capturing fat tails and heteroskedasticity in the covariances (overall or global fit), and none of the models presented do good in both areas. If the overall fit is what matters, then it appears from our work that a multivariate GARCH model is preferred. If the fit in the tails is more critical, as with CTE provisions, then according to our analysis, a multivariate regime-switching model is preferred.

REFERENCES

AMERICAN ACADEMY OF ACTUARIES (AAA). 2005. Recommended Approach for Setting Regulatory Risk-Based Capital Requirements for Variable Annuities and Similar Products. City: AAA.

ANG, A., AND G. BEKAERT. 2002a. Regime Switches in Interest Rates. *Journal of Business & Economic Statistics* 20: 163–82.

———. 2002b. International Asset Allocation with Regime Shifts. *Review of Financial Studies* 15: 1137–87.

BILLIO, M., AND L. PELIZZON. 2000. Value-at-Risk: A Multivariate Switching Regime Approach. *Journal of Empirical Finance* 7: 531–54.

BOLLERSLEV, T. 1990. Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model. *Review of Economics and Statistics* 72: 498–505.

BOLLERSLEV, T., R. Y. CHOU, AND K. F. KRONER. 1992. ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence. *Journal of Econometrics* 52: 5–59.

BOLLERSLEV, T., R. F. ENGLE, AND D. B. NELSON. 1994. ARCH Models. In *Handbook of Econometrics* 4, ed. R. F. Engle and D. L. McFadden, pp. 2961–3038. Amsterdam: North Holland.

BOLLERSLEV, T., R. F. ENGLE, AND J. M. WOOLDRIDGE. 1988. A Capital Asset Pricing Model with Time-Varying Covariances. *Journal of Political Economy* 96: 116–31.

CAMPBELL, J. Y., A. W. LO, AND A. C. MACKINLAY. 1997. *The Econometrics of Financial Markets*. Princeton: Princeton University Press.

CANADIAN INSTITUTE OF ACTUARIES (CIA). 2002. Report of the Task Force on Segregated Fund Investment Guarantees. Ontario: CIA.

ENGLE, R. F. 2002. Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business and Economic Statistics* 20: 339–50.

ENGLE, R. F., AND K. F. KRONER. 1995. Multivariate Simultaneous Generalized ARCH. *Econometric Theory* 11: 122–50.

ENGLE, R. F., V. K. NG, AND M. ROTHSCHILD. 1990. Asset Pricing with a Factor-ARCH Covariance Structure: Empirical Estimates for Treasury Bills. *Journal of Econometrics* 45: 213–37.

FREES, E. W., AND E. A. VALDEZ. 1998. Understanding Relationships Using Copulas. *North American Actuarial Journal* 2: 1–25.

HAMILTON, J. D. 1989. A New Approach to the Economic Analysis of Non-Stationary Time Series and the Business Cycle. *Econometrica* 50: 987–1007.

———. 1994. *Time Series Analysis*. Princeton: Princeton University Press.

HARDY, M. R. 2001. A Regime-Switching Model of Long Term Stock Returns. *North American Actuarial Journal* 5: 41–53.

———. 2003. *Investment Guarantees* (*Modeling and Risk Management for Equity-Linked Life Insurance*). New York: Wiley.

HARDY, M. R., R. K. FREELAND, AND M. C. TILL. 2006. Validation of Long-Term Equity Return Models for Equity-Linked Guarantees. *North American Actuarial Journal* 10: 28–47.

HOSKING, J. R. M. 1980. The Multivariate Portmanteau Statistic. *Journal of the American Statistical Association* 75: 602–8.

Joe, H. 1997. *Multivariate Models and Dependence Concepts*. Monographs on Statistics and Applied Probability 73. London: Chapman & Hall.

KLUGMAN, S. A., H. H. PANJER, AND G. E. WILLMOT. 2004. *Loss Models: From Data to Decisions*. 2nd ed. New York: Wiley-Interscience.

KRONER, K. E., AND V. K. NG. 1998. Modeling Asymmetric Comovements of Asset Returns. *Review of Financial Studies* 11: 817–44.

NELSEN, R. B. 1998. *An Introduction to Copulas*. Lecture Notes in Statistics 139. New York: Springer.

TSAY, R. S. 2005. *Analysis of Financial Time Series*. 2nd ed. Wiley Series in Probability and Statistics. New York: John Wiley.

TSE, Y. K. 2000. A Test for Constant Correlations in a Multivariate GARCH Model. *Journal of Econometrics*, 98: 107–27.

Discussions on this paper can be submitted until July 1, 2009. The authors reserve the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.