

Bayesian multivariate regime-switching models and the impact of correlation structure misspecification in variable annuity pricing

Brian Hartman^a, Chris Groendyke^b, and David Engler^c

^aDepartment of Statistics, Brigham Young University, Provo, UT, USA; ^bDepartment of Mathematics, Robert Morris University, Moon Township, PA, USA; ^cDepartment of Health Sciences Strategy, University of Utah, Salt Lake City, UT, USA.

ARTICLE HISTORY

Compiled May 7, 2019

ABSTRACT

We develop Bayesian multivariate regime-switching models for correlated assets, comparing three different ways to flexibly structure the correlation matrix. After developing the models, we examine their relative characteristics and performance, first in a straightforward asset simulation example, and later applied to a variable annuity product with guarantees. We find that the freedom allowed by the more flexible structures enables these models to more accurately reflect the actual asset dependence structure. We also show that the correlation structures inferred by most commonly used (and simplest) model will result in significantly larger estimates of the cost of the annuity guarantees.

KEYWORDS

Regime-switching; variable annuities; Guaranteed Minimum Income Benefit; Bayesian; correlation structures

1. Introduction

Accurate modeling of financial time series is essential in many areas. In addition to being crucial for option pricing and risk management, it is important to employ a sound model for financial time series when pricing or provisioning for investment guarantees such as those commonly found in modern variable and equity-indexed annuities and universal life products. Properly accounting for the uncertainty in the model and the dependence between observations can make the difference between a valuable model and a misleading one. Regime-switching models (Hamilton 1989, Hamilton and Susmel 1994) are an intuitive way to incorporate stochastic volatility and jumps into the asset pricing model. Hardy (2001) introduced these models to the actuarial literature and has since provided further model development (Hardy 2003). Hardy demonstrated that the regime-switching lognormal model fit US and Canadian index data (S&P 500 and TSE 300, respectively) relatively well.

Hardy limited investigations of regime-switching models to only single asset streams. However, when guarantees are based on multiple assets, simulating each individually may be inadequate, as doing so ignores the correlation between assets and hence exposes the writer to pricing risk. Boudreault and Panneton (2009) examined two families of multivariate models, namely regime-switching and GARCH (Generalized Auto-Regressive Conditionally Heteroskedastic) models. They found that whereas the multivariate GARCH models generally fit better throughout central elements of the distribution, regime-switching models provided better fit in the tails. Since

the majority of investment guarantees are written to protect against tail risk, further investigation of regime-switching models in multi-asset settings is warranted, paying particular attention to the modeling of the covariance structure between assets. Previous work considering a multiple asset framework when valuing guarantees embedded in variable annuity products includes Ng and Li (2011), which used a regime-switching model and considered guarantees involving multiple currencies; Ng and Li (2013), which presented models (including GARCH and regime-switching lognormal models) for pricing and hedging GMMB (Guaranteed Minimum Maturity Benefit) and GMDB (Guaranteed Minimum Death Benefit) guarantees under a two-asset framework; and Da Fonseca and Ziveyi (2017), which considers the valuation of the same two types of guarantees under two different payoff functions, using a stochastic mortality model.

When fitting regime-switching models, Boudreault and Panneton (2009) used maximum likelihood estimation and an unconstrained correlation matrix. Maximum likelihood is very quick and stable when the model is relatively simple (e.g., a few regimes and one asset), but as the model becomes more complicated, maximum likelihood estimation can become unstable. Maximum likelihood convergence is difficult to achieve with many parameters, especially when the density is multimodal, and multimodality is indeed expected in regime-switching models (Jasra et al. 2005). Additionally, the boundary constraints (non-negative-definiteness) of the covariance matrix can make convergence challenging. Thus, we will focus on Bayesian estimation, allowing for more stability and more flexibility in model specification. More importantly, Bayesian estimation makes it easier to incorporate prior beliefs about both the size and structure of the correlations, and allows for the quantification of parameter uncertainty in these regime-switching models (Hardy 2002, Hartman and Groendyke 2013).

There have been several recent developments in the utilization of regime-switching models to analyze various risks inherent to variable annuity and related products. For example, Gao et al. (2015) used a regime-switching model guided by a discrete state, continuous-time Markov process in order to allow to varying volatilities of interest and mortality rates in the pricing of a guaranteed annuity option. Costabile (2017) uses a trinomial lattice within in a regime-switching framework to find the fair price of a GMWB (Guaranteed Minimum Withdrawal Benefit) option on a variable annuity product. Kolkiewicz and Lin (2017) uses a regime-switching model of asset log-returns to evaluate the surrender risk of a ratchet equity-indexed annuity product.

Under a Bayesian framework, the constraint of positive definiteness restricts options for the prior distribution of the covariance matrix. A variety of approaches have been proposed that satisfy this constraint. A commonly used approach is that of Bernardo and Smith (1994), in which a conjugate inverse-Wishart distribution is employed as the prior. One drawback to this approach, however, is that because the degrees of freedom parameter is the only tuning parameter, it can be difficult to effectively incorporate prior information. Daniels and Kass (2001) alternatively proposed several flexible hierarchical priors. Additional prior distribution proposals include a reference prior that incorporates the eigenvalues of the covariance matrix (Yang and Berger 1994) and a log matrix prior (Leonard and Hsu 1992, Hsu et al. 2012). More recently, Wang and Pillai (2013) proposed an approach employing a scale mixture of uniform priors for the covariance matrix in high-dimensional spatial settings.

For flexibility and ease of interpretation, we adopt the separation strategy proposed by Barnard et al. (2000) in which the covariance matrix Σ is rewritten as

$$\Sigma = \mathbf{SRS},$$

where \mathbf{S} is a diagonal matrix containing the vector of element-wise standard deviations and \mathbf{R} is the correlation matrix. The approach allows for the incorporation of the positive definiteness constraint (see Section 2.4). Importantly, this approach also allows for the direct modeling of correlation structure, a primary objective in the joint modeling of multiple asset streams.

Following the approach of Liechty et al. (2004), we examined three different prior structures

for the correlation matrix, \mathbf{R} . First, we examined a common correlation model where each element of the correlation matrix follows a single prior distribution. Second, we investigated a grouped correlations model where correlations are allowed to cluster into similar groups, with each group sharing a common prior distribution. Finally, we explored use of a grouped variables model which allows the observed variables (*i.e.*, asset streams) to cluster into various groups. Under this model, correlations between asset streams in the same group share a common prior distribution. Additionally, the prior distributions of correlations between asset streams not in the same group depend on the specific group assignments.

Our contributions in this paper are three-fold. We develop new models incorporating both regime-switching and prior structure in the correlation matrix. Second, by applying these models to the stock returns of health insurance companies and banks we begin to compare the strengths and short-comings of each of the models. Finally, and most importantly for actuarial practice, we examine the impact of model choice on the pricing and reserving for products such as variable annuities. We show that the model chosen can greatly affect the estimation of asset correlations, leading to significant impacts on the pricing and provisioning for these types of products.

The remainder of this paper is organized as follows: Section 2 discusses the three models we use to describe the correlation structure between the assets, Section 3 applies these models to data from nine asset streams across two industry sectors, Section 4 gives the results of an application of these models to a variable annuity product, and Section 5 provides some conclusions and ideas for future work.

2. Model

When modeling the logarithmic returns of a single asset, the normal distribution is a natural first choice because of its connection to geometric Brownian motion and the Black-Scholes-Merton model. Assume the data set contains n periodic observations, indexed by i . Now generalizing the model to C different assets across the same time period ($i \in \{1, \dots, n\}$), the multivariate normal distribution is a sensible extension. Let \mathbf{y}_i denote a $C \times 1$ vector of log returns for the i^{th} period; \mathbf{y}_i follows a multivariate normal distribution, the parameters of which depend on the regime at period i , which is given by x_i , with $x_i \in \{1, \dots, P\}$:

$$\mathbf{y}_i | x_i = p \sim N_C(\boldsymbol{\mu}_p, \boldsymbol{\Sigma}_p).$$

(Note that x_i will typically not be observed and hence must be inferred.) The regime-switching process follows a discrete-time Markov chain. The transition probability vector for regime p is $\boldsymbol{\pi}_p$, with individual elements $\pi_{p,1}, \dots, \pi_{p,P}$, and is assigned a Dirichlet prior. We define n_p as the number of data points assigned to regime p , so that $\sum_{p=1}^P n_p = n$. As mentioned above, the covariance matrix in a given regime can be decomposed into a diagonal matrix of standard deviations, \mathbf{S}_p , and a correlation matrix, \mathbf{R}_p :

$$\mathbf{y}_i | x_i = p \sim N_C(\boldsymbol{\mu}_p, \mathbf{S}_p \mathbf{R}_p \mathbf{S}_p),$$

$$\mathbf{y}_i = \begin{bmatrix} y_{i1} \\ y_{i2} \\ \vdots \\ y_{iC} \end{bmatrix}, \quad \mathbf{S}_p = \begin{bmatrix} \sigma_{p1} & 0 & \dots & 0 \\ 0 & \sigma_{p2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \sigma_{pC} \end{bmatrix}, \quad \mathbf{R}_p = \begin{bmatrix} r_{p11} & r_{p12} & \dots & r_{p1C} \\ r_{p21} & r_{p22} & \dots & r_{p2C} \\ \vdots & \vdots & \ddots & \vdots \\ r_{pC1} & r_{pC2} & \dots & r_{pCC} \end{bmatrix}, \quad \boldsymbol{\mu}_p = \begin{bmatrix} \mu_{p1} \\ \mu_{p2} \\ \vdots \\ \mu_{pC} \end{bmatrix}.$$

It is important to note that for the following models, direct sampling of r_{pab} from a normal

distribution (subject to the positive definiteness constraint) with mean μ (which is assigned a normal prior) and variance σ^2 (which is assigned an inverse gamma prior) results in some complications. Namely, the full conditionals for both μ and σ^2 will entail normalizing constants that intractably involve μ and σ^2 . To resolve this complication we employed latent variables, or shadow priors, δ_{pab} (for $a < b$ in regime p), between the likelihood of r_{pab} and the priors for μ and σ^2 . This approach drastically reduces the computational burden while providing a good approximation to the true normalizing constant. The variance of the shadow prior, ν^2 , serves as a tuning parameter: smaller values result in a better approximation to the actual model, whereas larger values improve the sampling of the posterior distribution. We experimented with a range of values for this parameter and ultimately used $\nu^2 = 0.01$ for all cases. A full discussion of the utility of shadow priors is provided by Liechty et al. (2004, 2009).

2.1. Common Correlation Model

Under the common correlation model, it is assumed that all correlations between asset streams have a common distribution for a given regime. The mean, $\boldsymbol{\mu}_p$, and the variance elements, σ_{pj}^2 , of the log returns are given multivariate normal and inverse-gamma priors, respectively. We also define

$$\bar{\mathbf{y}}_p = \frac{1}{n_p} \left[\sum_{\{x_i=p\}} y_{x_i1}, \sum_{\{x_i=p\}} y_{x_i2}, \dots, \sum_{\{x_i=p\}} y_{x_iC} \right]^T$$

as the vector of mean log returns for the points assigned to regime p . Then for this common correlation model, we have:

$$\begin{aligned} \boldsymbol{\mu}_p &\sim N_C(\mathbf{0}, \tau^2 \mathbf{I}), \\ \sigma_{pj}^2 &\sim IG(\alpha_\sigma, \beta_\sigma), \\ r_{pab} &\sim N(\delta_{pab}, \nu^2), \quad \nu^2 = 0.01, \\ \delta_{pab} &\sim N(\lambda_p, \gamma_p^2), \\ \lambda_p &\sim N(\mu_\lambda, \sigma_\lambda^2), \\ \gamma_p^2 &\sim IG(\alpha_\gamma, \beta_\gamma), \\ \boldsymbol{\pi}_p &\sim Dir(1, 1, \dots, 1), \\ \Pr(x_i = p) &= 1/P. \end{aligned}$$

These definitions and priors result in the following full conditional distributions:

$$\begin{aligned}
\boldsymbol{\mu}_p | \cdot &\sim N_C \left[(\tau^{-2} \mathbf{I} + n_p (\mathbf{S}_p \mathbf{R}_p \mathbf{S}_p)^{-1})^{-1} (n_p (\mathbf{S}_p \mathbf{R}_p \mathbf{S}_p)^{-1} \bar{\mathbf{y}}_p), (\tau^{-2} \mathbf{I} + n_p (\mathbf{S}_p \mathbf{R}_p \mathbf{S}_p)^{-1})^{-1} \right], \\
\sigma_{pj}^2 | \cdot &\sim IG \left(\alpha_\sigma + \frac{n_p}{2}, \beta_\sigma + \frac{\sum_{\{x_i=p\}} (y_{x_i,j} - \mu_{pj})^2}{2} \right), \\
f(r_{pab} | \cdot) &\propto |\mathbf{R}_p|^{-n_p/2} \exp \left\{ -\frac{(r_{pab} - \delta_{pab})^2}{2\nu^2} - \frac{1}{2} \sum_{\{x_i=p\}} (\mathbf{y}_{x_i} - \boldsymbol{\mu}_p)^T (\mathbf{S}_p \mathbf{R}_p \mathbf{S}_p)^{-1} (\mathbf{y}_{x_i} - \boldsymbol{\mu}_p) \right\}, \\
\delta_{pab} | \cdot &\sim N \left(\frac{\lambda_p / \gamma_p^2 + r_{pab} / \nu^2}{1/\gamma_p^2 + 1/\nu^2}, \frac{1}{1/\gamma_p^2 + 1/\nu^2} \right), \\
\lambda_p | \cdot &\sim N \left(\frac{\mu_\lambda / \sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \delta_{pab} / \gamma_p^2}{1/\sigma_\lambda^2 + C(C-1)/2\gamma_p^2}, \frac{1}{1/\sigma_\lambda^2 + C(C-1)/2\gamma_p^2} \right), \\
\gamma_p^2 | \cdot &\sim IG \left(\alpha_\gamma + \frac{C(C-1)}{4}, \beta_\gamma + \frac{\sum_{a=1}^{C-1} \sum_{b=a+1}^C (\delta_{pab} - \lambda_p)^2}{2} \right), \\
\Pr(x_i = p | \cdot) &\propto N_C(\boldsymbol{\mu}_p, \mathbf{S}_p \mathbf{R}_p \mathbf{S}_p) \cdot \pi_{x_{i-1},p} \cdot \pi_{p,x_{i+1}}, \\
\boldsymbol{\pi}_p | \cdot &\sim Dir(1 + m_{p1}, 1 + m_{p2}, \dots, 1 + m_{pC}), \quad \text{where } m_{jk} = \sum_{i=1}^n \mathbb{1}\{x_i = j, x_{i+1} = k\}.
\end{aligned}$$

2.2. Grouped Correlations Model

In some settings, there may be groups of correlation parameters that are similar to one another, yet differ greatly from the correlation parameters outside their group. Under the grouped correlations model, K clusters of correlations are posited. Elements \mathbf{y}_i , \mathbf{S}_p , \mathbf{R}_p , $\boldsymbol{\mu}_p$, $\boldsymbol{\pi}_p$, and σ_{pj}^2 are defined as above. In this model, however, correlations r_{pab} are drawn from a mixture of K normals. Also, for this model, θ_{pab} represents the number of the group to which the correlation r_{pab} belongs in regime p . Prior probability of membership in group k is assumed to be equal to $1/K$ for all k . Thus, for this model, we have:

$$\begin{aligned}
r_{pab} &\sim N(\delta_{pab}, \nu^2), \quad \nu^2 = 0.01, \\
\delta_{pab} | \theta_{pab} = k &\sim N(\lambda_{pk}, \gamma_{pk}^2), \\
\lambda_{pk} &\sim N(\mu_\lambda, \sigma_\lambda^2), \\
\gamma_{pk}^2 &\sim IG(\alpha_\gamma, \beta_\gamma), \\
\Pr(\theta_{pab} = k) &= 1/K.
\end{aligned}$$

These definitions and priors result in the following full conditional distributions:

$$\begin{aligned}
f(r_{pab}|\cdot) &\propto |\mathbf{R}_p|^{-n_p/2} \exp \left\{ -\frac{(r_{pab} - \delta_{pab})^2}{2\nu^2} - \frac{1}{2} \sum_{\{x_i=p\}} (\mathbf{y}_{x_i} - \boldsymbol{\mu}_p)^T (\mathbf{S}_p \mathbf{R}_p \mathbf{S}_p)^{-1} (\mathbf{y}_{x_i} - \boldsymbol{\mu}_p) \right\}, \\
\delta_{pab} | \theta_{pab} = k, \cdot &\sim N \left(\frac{\lambda_{pk}/\gamma_{pk}^2 + r_{pab}/\nu^2}{1/\gamma_{pk}^2 + 1/\nu^2}, \frac{1}{1/\gamma_{pk}^2 + 1/\nu^2} \right), \\
\lambda_{pk} | \cdot &\sim N \left(\frac{\mu_\lambda/\sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\theta_{pab}=k\}} \delta_{pab}/\gamma_{pk}^2}{1/\sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\theta_{pab}=k\}}/\gamma_{pk}^2}, \frac{1}{1/\sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\theta_{pab}=k\}}/\gamma_{pk}^2} \right), \\
\gamma_{pk}^2 | \cdot &\sim IG \left(\alpha_\gamma + \frac{\sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\theta_{pab}=k\}}}{2}, \beta_\gamma + \frac{\sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\theta_{pab}=k\}} (\delta_{pab} - \lambda_{pk})^2}{2} \right), \\
\Pr(\theta_{pab} = k | \cdot) &\propto \exp \left\{ \frac{-(\delta_{pab} - \lambda_{pk})^2}{2\gamma_{pk}^2} \right\}.
\end{aligned}$$

2.3. Grouped Variables Model

Rather than allowing the correlation values to cluster, it may be more natural to allow for groupings between the actual variables. Within each asset group, correlations between variables have a common prior distribution. Cross-group correlations are specified separately for each pairwise group combination, where there are L defined groups. Let C_w denote the total number of intra-group asset pairs, and C_v denote the total number of cross-group asset pairs, so that $C_w + C_v = C(C-1)/2$. For this model, θ_{pi} denotes the group number to which asset i belongs in regime p . Prior probability of membership in group l is assumed to be equal to $1/L$ for all l . Elements \mathbf{y}_i , \mathbf{S}_p , \mathbf{R}_p , $\boldsymbol{\mu}_p$, $\boldsymbol{\pi}_p$, and σ_{pj}^2 are defined as in the previous models, so that we have:

$$\begin{aligned}
r_{pab} &\sim N(\delta_{pab}, \nu^2), \quad \nu^2 = 0.01, \\
\delta_{pab} &\sim N(\lambda_{pgh}, \gamma_{pgh}^2), \\
\lambda_{pgh} &\sim N(\mu_\lambda, \sigma_\lambda^2) \quad \text{for } g \leq h, \\
\gamma_{pgh}^2 &\sim IG(\alpha_\gamma, \beta_\gamma) \quad \text{for } g \leq h, \\
\Pr(\theta_{pi} = l) &= 1/L,
\end{aligned}$$

where $g = \min(\theta_{pa}, \theta_{pb})$ and $h = \max(\theta_{pa}, \theta_{pb})$.

These definitions and priors result in the following full conditional distributions:

$$\begin{aligned}
f(r_{pab}|\cdot) &\propto |\mathbf{R}_p|^{-n_p/2} \exp \left\{ -\frac{(r_{pab} - \delta_{pab})^2}{2\nu^2} - \frac{1}{2} \sum_{\{x_i=p\}} (\mathbf{y}_{x_i} - \boldsymbol{\mu}_p)^T (\mathbf{S}_p \mathbf{R}_p \mathbf{S}_p)^{-1} (\mathbf{y}_{x_i} - \boldsymbol{\mu}_p) \right\}, \\
\delta_{pab} | \{\theta_{pa}, \theta_{pb}\} = \{g, h\}, \cdot &\sim N \left(\frac{\lambda_{pgh}/\gamma_{pgh}^2 + r_{pab}/\nu^2}{1/\gamma_{pgh}^2 + 1/\nu^2}, \frac{1}{1/\gamma_{pgh}^2 + 1/\nu^2} \right),
\end{aligned}$$

$$\lambda_{pgh} | \cdot \sim N \left(\frac{\mu_\lambda / \sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\{\theta_a, \theta_b\}=\{g,h\}\}} \delta_{pab} / \gamma_{pgh}^2}{1 / \sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\{\theta_a, \theta_b\}=\{g,h\}\}} / \gamma_{pgh}^2}, \frac{1}{1 / \sigma_\lambda^2 + \sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\{\theta_a, \theta_b\}=\{g,h\}\}} / \gamma_{pgh}^2} \right),$$

$$\gamma_{pgh}^2 | \cdot \sim IG \left(\alpha_\gamma + \frac{\sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\{\theta_{pa}, \theta_{pb}\}=\{g,h\}\}}}{2}, \beta_\gamma + \frac{\sum_{a=1}^{C-1} \sum_{b=a+1}^C \mathbb{1}_{\{\{\theta_a, \theta_b\}=\{g,h\}\}} (\delta_{pab} - \lambda_{pgh})^2}{2} \right),$$

$$\Pr(\theta_{pi} = l | \cdot) \propto \exp \left\{ - \sum_{j=1}^{i-1} \frac{(\delta_{pji} - \lambda_{pcd})^2}{2\gamma_{pcd}^2} - \sum_{j=i+1}^C \frac{(\delta_{pij} - \lambda_{pcd})^2}{2\gamma_{pcd}^2} \right\},$$

where $c = \min(\theta_{pj}, l)$ and $d = \max(\theta_{pj}, l)$.

2.4. Gibbs Sampling

For most model elements, straightforward Gibbs or Metropolis-Hastings sampling is effective and sufficient. With regard to sampling of the covariance elements, however, special care must be taken to ensure the positive definiteness of \mathbf{R}_p . A straightforward approach for such settings is outlined by Barnard et al. (2000) and Ritter and Tanner (1992) in which candidate values are drawn from a grid of values that preserve this condition (conditional on the other parameter values). This griddy Gibbs approach was employed for sampling of covariance elements in the current setting.

3. Stock Return Data Analysis

As a first example of our method in an actual setting, we fit the three models to weekly total return data from nine companies across two sectors from February 2014 through May 2016 (Yahoo! Inc. 2017). We used data from four health insurers [United Healthcare (UNH), Aetna (AET), Cigna (CI), and Humana (HUM)] and five banks [Bank of America (BAC), Key Bank (KEY), Wells Fargo (WFC), Citigroup (C), JP Morgan (JPM)]. We chose weekly (as opposed to daily) data to provide a sufficient number of observations over the two-year period. Previous authors have shown that data that is more frequent than weekly is best fit by a regime-switching model with many regimes; for example, see Hartman and Heaton (2011), which finds that 6-7 regimes provide the best fit for daily data. With weekly data, we fit a model with fewer regimes, and which is consequently much easier to interpret. Additionally, most conditions, options, or guarantees on variable annuities, including those we use in this paper as examples, trigger on a monthly, quarterly, or annual basis. Having data measured more frequently would not change the resulting payment patterns. For each model, the MCMC chain was run for 100,000 iterations, discarding the first 10,000 as burn-in. Hyperparameter settings are available in Appendix A.

We experimented with fitting models containing $P = 2, 3$, and 4 regimes, but all of the results we present here correspond to the models with 2 regimes, as our results indicated that additional regimes were unnecessary for this data. We refer to the regime containing the greater number of observations as the primary regime and the other regime as the secondary regime. We found that when we added additional regimes to the model, the algorithm for partitioning the data into regimes tended to place the same data points into the primary regime, while subdividing the

secondary regime into two or more partitions. Furthermore, these new smaller secondary regimes had very similar inferred parameter values (and hence similar interpretations) to one another, implying that there was little or no additional benefit from their inclusion. As a result, we believe that the models with two regimes were the best and most parsimonious representations of the data. In later sections, we discuss the regime interpretation for these models.

For this data we were primarily interested in exploring the magnitude and structure of the correlations of pairs of assets both within and between sectors, how the assets and correlations are grouped by the models, and in particular how these results vary among the three models of correlation structure.

We also explored the assignment of points to the different regimes. Often, when regime-switching models with two regimes are applied to asset return streams, one regime has a high mean and low standard deviation, representing a normal state of the economy, whereas the other has a lower mean and higher standard deviation, representing a recession or turbulent state (Hardy 2003, Hartman and Heaton 2011).

3.1. Common Correlation Model

Under the common correlation model (a special case of the grouped correlations and grouped variables models with $K = 1$ or $L = 1$), in the primary regime we see that all the posterior means of the correlation elements are between about 0.65 and 0.9, with the cross-sector correlations being the smallest (see Figure 1). The common prior distribution acts to pull all of the individual correlation elements together. There is a similar effect in the secondary regime, but the magnitude of all the elements is much smaller. Neither regime shows strong patterns of grouping by asset sector in the posterior distributions of the correlation elements. There are fewer data in the secondary regime (only about 20%), allowing the prior to have more impact.

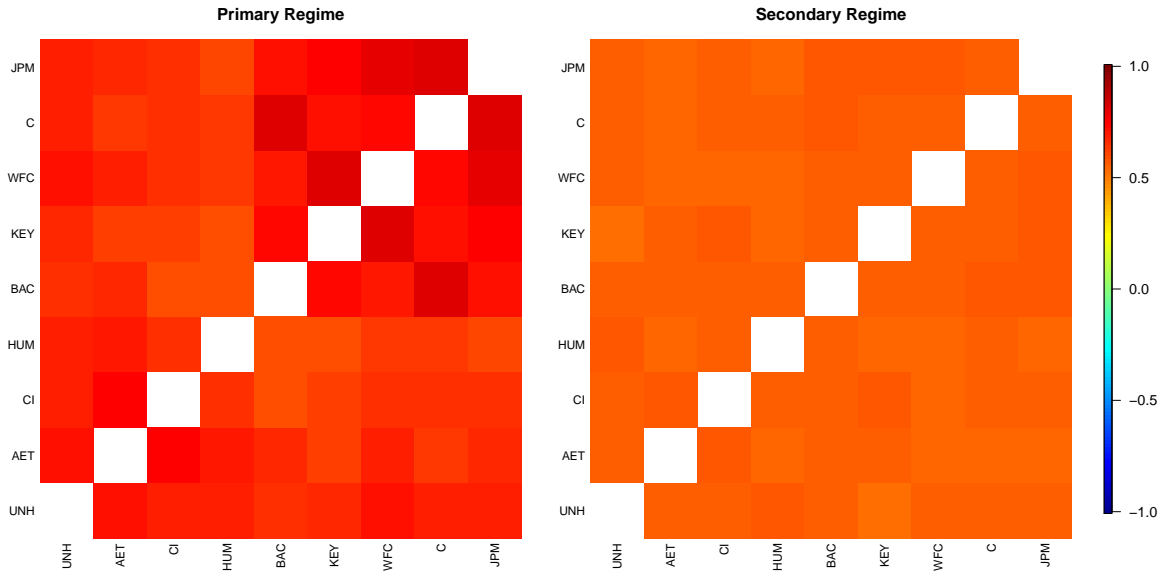


Figure 1.: Posterior means of health and bank correlation parameters under the common correlation model

When we examine the assignment of points to regimes, we notice significant differences in the μ parameters between regimes; in the primary regime, the bank stocks all have positive mean returns and the health insurance stocks all have a negative mean returns, while the opposite is true in the secondary regime. Figure 2 shows the estimated means of the posterior distributions

of the μ_p parameters for the nine asset streams for the common correlation model. Detailed values and standard errors can be seen in Table 4 in the appendix. Thus, for this model, we can see that there are clear differences between the behavior of the parameters in the two regimes, especially for the mean parameters, and to a lesser extent for the correlation parameters; further, there is evidence that the model has placed the data into the two regimes largely based on the values of the mean returns for the asset streams.

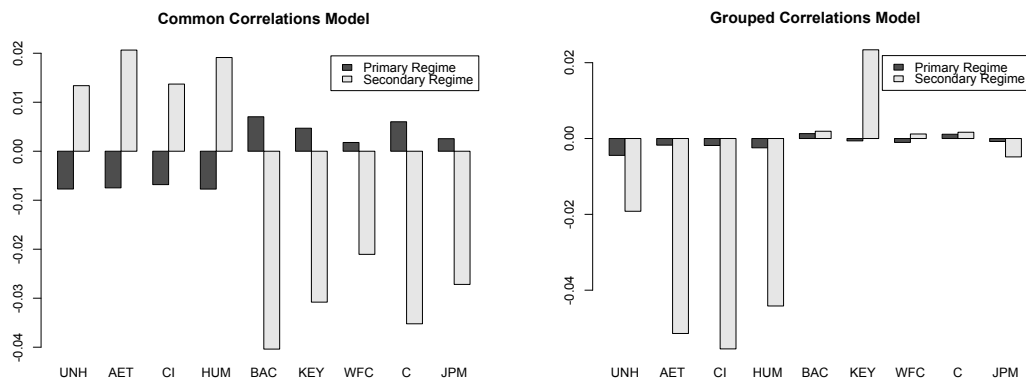


Figure 2.: Posterior means of μ_p parameters under the common correlation model (left) and the grouped correlations model (right).

3.2. Grouped Correlations Model

The structure of the grouped correlations model allows the individual correlation elements to cluster with other elements more freely. This allows us to test if the common correlation model is too restrictive. We first set K , the number of correlation groups, to two.




Figure 3.: Posterior means of bank and health correlation parameters under the under the grouped correlations model with $K = 2$

We immediately see strong evidence that the common correlation model is overly restrictive. When the correlation elements are allowed to cluster more freely, in the primary regime all of the within-sector correlations (both for banks and health insurers) cluster around 0.8, while the cross-sector correlations are much smaller, around 0.4 (see Figure 3). We can see from Figure 4 that the correlations cluster strongly into one cluster of within-sector correlations and one cluster for cross-sector correlations. (We used a hierarchical clustering algorithm, using a Euclidean distance metric and complete linkage method, (R Core Team 2016) to examine the clustering pattern of these correlations.) There is little evidence of sub-clustering within either of these two clusters. In the secondary regime, the correlation elements are generally much weaker (most are close to zero), and exhibit much less pattern, with no apparent clustering by sector.

When we increased the number of groups K to three, we observed a similar effect. For the primary regime, the within-sector correlations are still larger than the cross-sector correlations. The health and bank within-sector correlations stay relatively close to each other, even with the added flexibility of a third group. Figure 5 again shows that the correlations cluster strongly into one cluster of within-sector correlations and one cluster for cross-sector correlations. However, unlike the $K = 2$ case, there is some evidence of intra-sector sub-clustering inside the within-sector cluster. The within-bank correlations and within-health correlations tend to mostly form their own sub-clusterings, though these groupings are not perfect. In the secondary regime, the results are largely similar to the case where $K = 2$, though the magnitude of the correlation elements is a bit greater than in this previous case. In this case, we also see little or no evidence of clustering by sector in the secondary regime.

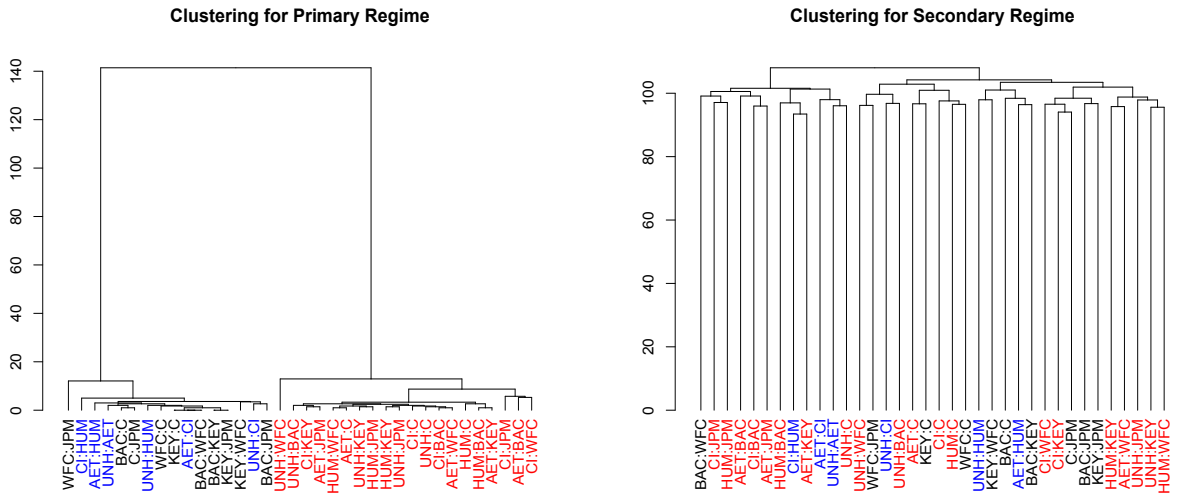


Figure 4.: Dendrograms for bank and health correlations in the grouped correlations model ($K = 2$). Blue labels correspond to intra-bank correlations, black labels correspond to intra-health correlations, and red labels represent inter-sector correlations.

We can see from Figure 2 that for some assets, the means are greater in the primary regime, while for others, the means have greater values in the secondary regime. This implies that a high return / low return interpretation of the two regimes is unlikely to be applicable here. The corresponding parameter means for the $K = 3$ case are not shown here; the results for these μ_p parameters are very similar to the $K = 2$ case.

Indeed, upon closer inspection, it appears more likely that the points being assigned to the secondary regime by this model were not distinguished by their means, but rather the degree to which the assets within a sector moved or clustered with each other. To quantify this, we used




Figure 5.: Dendrograms for bank and health correlations in the grouped correlations model ($K = 3$). Blue labels correspond to intra-bank correlations, black labels correspond to intra-health correlations, and red labels represent inter-sector correlations.

the intra-sector ranges to measure the return divergence at each time point. That is, for each time point i , we have computed

$$\Delta y_i^\tau = \max(\Delta y_i^B, \Delta y_i^H), \text{ where}$$

$$\Delta y_i^B = \max_{j \in \text{Bank}} (y_{ij}) - \min_{j \in \text{Bank}} (y_{ij}) \quad \text{and} \quad \Delta y_i^H = \max_{j \in \text{Health}} (y_{ij}) - \min_{j \in \text{Health}} (y_{ij}).$$

Figure 6 plots the maximum of the two intra-sector ranges (Δy_i^τ) for each point in the dataset. We can see clearly that the assignment of points to the secondary regime corresponds very strongly to the larger values of the intra-sector ranges.

3.3. Grouped Variables Model

In the grouped variables model, we allow the individual assets, as opposed to the correlation elements, to cluster. We first set L , the number of asset groups, to two.

Examining the posterior means of the correlation (r_{pab}) elements, we can see from Figure 7 that in the primary regime, the highest correlations belong to the intra-bank pairs of assets, followed closely by the correlations corresponding to the intra-health sector, while the health/bank cross-sector correlations are significantly smaller in magnitude. Not surprisingly, we can see that the correlations are grouped strongly by sector, though not quite as strongly as in the grouped correlations model. In the secondary regime, the correlation means are smaller in magnitude, roughly centered around zero, and show no discernible pattern by asset category; these are very similar to the results found for the grouped correlations model.

For this grouped variables model, it is also interesting to explore the clustering of assets, based on the inferred clusters assigned by the algorithm. As with the grouped correlations model, we used a hierarchical clustering algorithm, using a Euclidean distance metric and complete linkage method to examine the clustering pattern of these assets. We can see from Figure 8 that in the primary regime, the assets are strongly clustered into their sectors, whereas in the

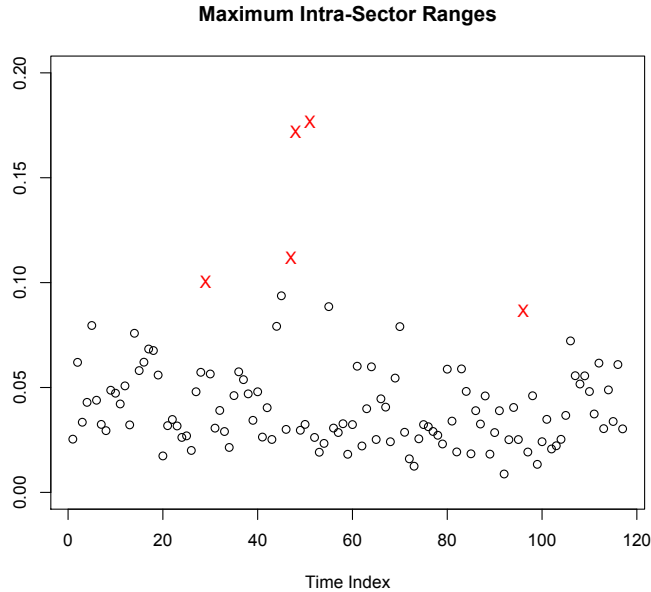


Figure 6.: Maximum intra-sector ranges. Points most commonly assigned to the secondary regime are marked by a red X.

secondary regime, there is little to no apparent clustering. All of these results reinforce the findings above, namely that the assignment of points to regimes in this model is primarily driven by the relationships between the assets, i.e., the correlation elements, rather than their individual means or variances.

When we allowed for three groups in this model, the conclusions did not materially change. The correlation posterior distributions look very similar to those produced in the previous case with $L = 2$. That is, in the primary regime, the highest correlations belong to the intra-bank pairs of assets, followed closely by the correlations corresponding to the intra-health sector, while the health/bank cross-sector correlations are significantly smaller in magnitude; in the secondary regime, the correlation densities are smaller in magnitude, roughly centered around zero, and show no discernible pattern by asset category.

Regarding the assignment of data points to regimes, we found that the grouped variables model, for both the cases of $L = 2$ and $L = 3$ groups, assigned points to the primary and secondary regimes in a very similar manner to the grouped correlations model. As a result, the resulting regimes have the same interpretations as they did under the grouped correlations model, and most of the other model parameters such as the μ_p parameters were estimated similarly to the corresponding parameters in the grouped correlations model.

4. Variable Annuity Application

Here we apply the three correlation models to a situation involving a variable annuity product with a guarantee. We fit the models to monthly returns from some of the funds offered as investment choices in a popular variable annuity product (Allianz 2018) from January 2005 through June 2017. Again, we used two regimes in all cases, and allowed for two, three, and four groups in both the grouped correlations and grouped variables models. For each model, the MCMC chain was run for 5,000,000 iterations (thinning every 100 iterations) with 50,000 burn-in iterations. We used the same values for all of the hyperparameters as were used in the

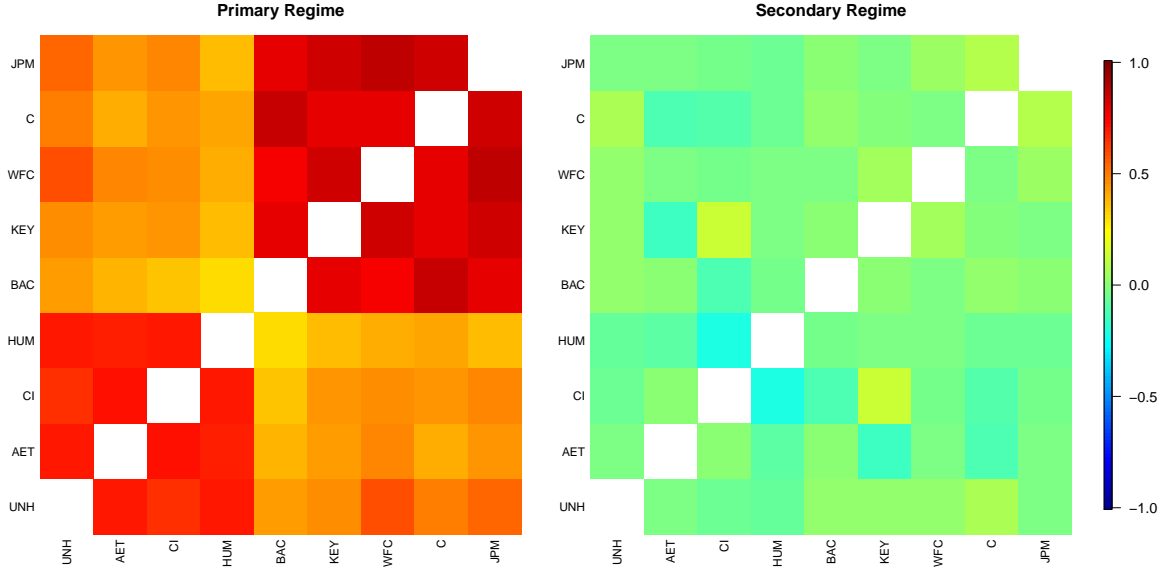


Figure 7.: Posterior means of bank and health correlation parameters under the grouped variables model ($L = 2$).

previous analysis. The posterior samples produced for each model were used in the subsequent inference and simulation described later in this section.

The variable annuity product considered here is a standard variable annuity design with a guaranteed minimum income benefit. We assume that the policyholder makes a single deposit of 10,000 at the onset of the contract. This money is invested in various investment funds, in proportions chosen by the policyholder. After some time period, the policyholder can choose to begin to receive benefit payments, payable for the remainder of their lifetime. For simplicity, we assume that the benefit is taken in the form of a single life annuity, though typically other payout options — such as single life with a period certain or joint and last survivor — would often be available.

During the accumulation phase (i.e., prior to the first benefit payment), the nominal account value grows with the performance of the underlying funds. The nineteen fund choices used for this application are listed in the appendix. In general, the policyholder could allocate their investment in any manner among these fund choices. For simplicity, we assumed that the investment was allocated equally among all nineteen funds; we also assumed a quarterly rebalancing of the allocation. This hypothetical annuity product offers downside investment protection to the policyholder in the form of a quarterly ratchet (or 0% quarterly roll-up): if the nominal account value at the end of a quarter is less than its value at the end of the previous quarter, the account value is reset to its previous quarterly value. The nominal account is charged an expense charge to cover the product’s maintenance expenses and the cost of the guarantee; this expense charge is 45 bp per quarter, or 190 bp annually, during the accumulation phase. We assume no withdrawals or deaths during the accumulation phase.

The payout phase begins when the policyholder elects to begin to take annuity payments, commonly at time of retirement. The amount of the annual benefit payment, payable for as long as the annuitant is alive, is given in this case by:

$$BA = BR \times I \times \prod_{i=1}^{QBR} \max(1 + r_i, 1) \times (1 - ME_i), \quad (1)$$

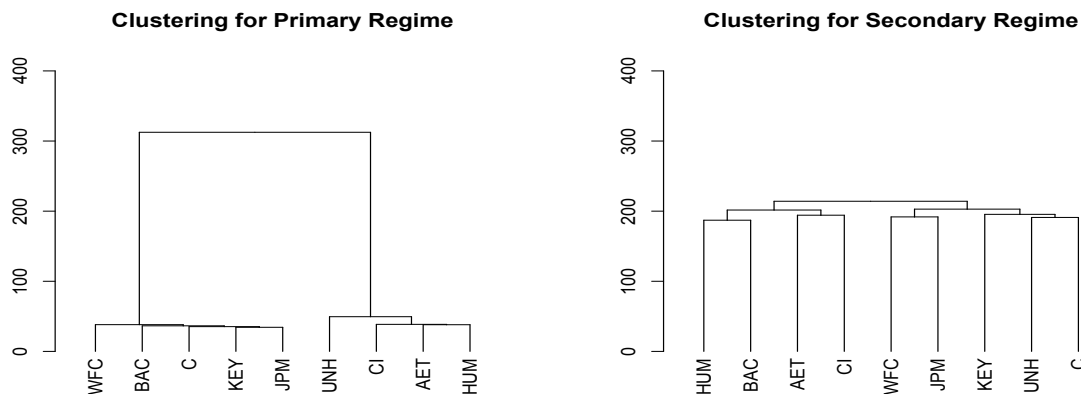


Figure 8.: Dendrograms for bank and health assets under the grouped variables model ($L = 2$).

where BA is the annual benefit payment, I is the amount of the single deposit, BR is the benefit rate (assumed to be 4% here), QBR is the number of quarterly time periods from the deposit to retirement, ME_i is the quarterly maintenance expense rate (0.0045 here), and r_i represents the quarterly asset portfolio return in period i .

In the remainder of this section we first describe the results of the inference for the model parameters for this data under the three correlation models; we then perform a simulation study where we use the inference results to simulate asset valuations for this product under the guarantees described above for these models. Finally, we consider how the results change as some of the underlying assumptions vary.

4.1. Inference Results

We consider here the inferred posterior distributions of the model parameters under the three different models of correlation structure for these 19 assets. For the grouped correlations and grouped variables models, we focus our attention on the case where we assume that the assets or correlations cluster into two groups, that is, $K = 2$ or $L = 2$. Both the common correlation and grouped correlations model typically assigned about 90-95% of the data to the primary regime, whereas the grouped variables model usually assigned 60-70% of the points to the primary regime. With respect to the μ and σ parameters, the grouped correlations model showed some evidence of producing a primary regime with high mean / low variance and a secondary regime with low mean / high variance; however, this pattern was not present in the other two correlation models. The inferred overall mean log returns under the different regimes are shown in Table 1.

Table 1.: Overall mean log returns by model.

Model	Primary Regime Mean	Secondary Regime Mean
Common Correlation	0.0058	0.0074
Grouped Correlations	0.0073	-0.0286
Grouped Variables	0.0019	0.0105

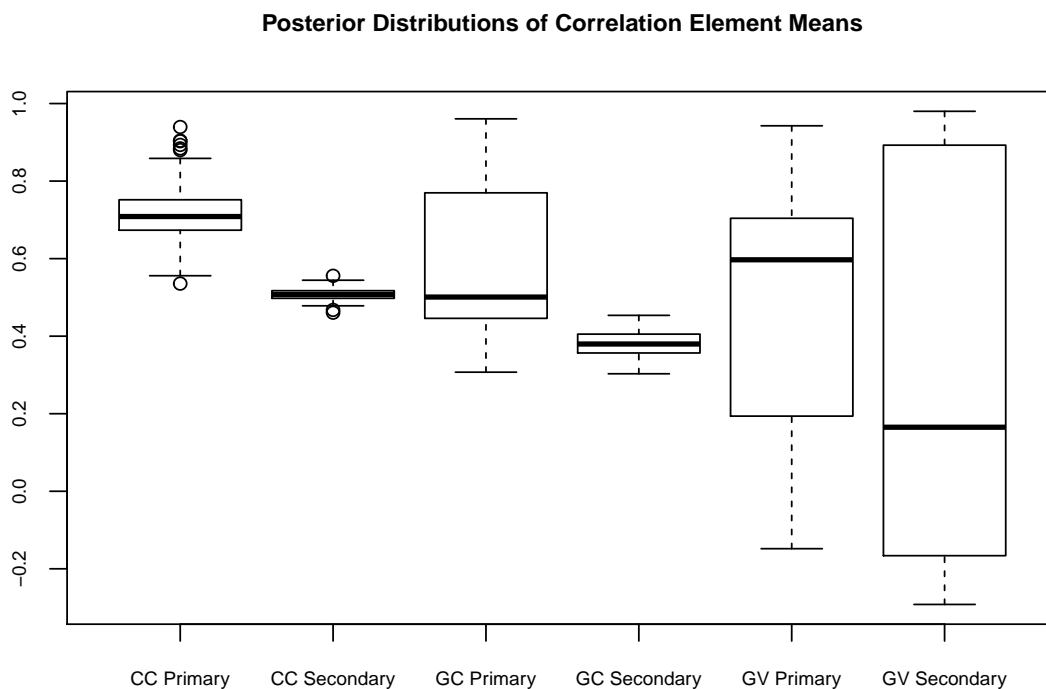
Next we consider the correlation elements. As was the case in the previous example, we find that the three models produced widely varying results in terms of the posterior distributions of these correlation parameters. In the common correlation model, under each regime, the correlation elements were all fairly similar in magnitude, i.e., all of the elements clustered in a

relatively tight range, with little or no discernible structural patterns among assets or correlation elements. In general, the correlation was significantly stronger in the primary regime than in the secondary regime.

As was the case in the previous data set, we see here that the more complex correlation structures allows more freedom for the correlation elements. The biggest difference in the results of the grouped correlations model (as compared with the common correlation model) for this data was that in the primary regime, there was a bigger spread in the mean values of the correlation elements. In addition, the results showed evidence of significant clustering among correlation elements in the primary regime.

By allowing the assets to cluster together, the grouped variables model shows the biggest spread among correlation means in both the primary and secondary regimes. Like the grouped correlations model, the primary regime showed evidence of clustering; however, unlike the grouped correlations model, this clustering was also present in the secondary regime. These results again indicate that the common correlation model is likely too restrictive, and the more complex models are better able to describe the correlation structure in the data. Boxplots of the means of the posterior distributions of the correlation elements are given in Figure 9.

Figure 9.: Boxplots of the means of the posterior distributions of inferred correlation elements for primary and secondary regimes under the common correlation (CC), grouped correlation (GC), and grouped variables (GV) models.



Analyzing the results for the cases where 3 or 4 groups were allowed for the grouped correlations and grouped variables models, the results did not change materially from the cases with 2 groups. There was little or no evidence that the models with more groups provided better fit to the data or that a larger number of groups was needed for this data.

4.2. Simulation Results

Here we discuss the results from a simulation study which uses the results of the preceding model parameter inference to simulate the asset valuations of an insurer issuing the VA product. We consider the results for the three correlation models as a function of the length of time after retirement that the annuitant lives. In all cases, we assume that the accumulation phase lasts for 25 years; we simulate a payout phase lasting up to 25 years. Again, we focus on the cases with two groups ($K = 2, L = 2$), with the results being similar in the cases where we allowed for 3 or 4 asset or correlation groupings. The hyperparameter settings are the same as those used in the previous example (see appendix for more details).

We used the joint posterior parameter density from the previous section to simulate 50,000 series of quarterly asset portfolio returns. The estimated densities of these simulated returns are shown in Figure 10. The common correlation model has the thickest tails overall because of the high estimated correlations between the assets. The grouped correlations model is much more left-skewed than the other two. This is because the grouped correlations model is assigning regimes largely based on the return means, while the common correlation and grouped variables models assigned regimes largely based on correlation differences. The grouped correlations model is then able to have a low mean secondary regime, causing the left-skewness.

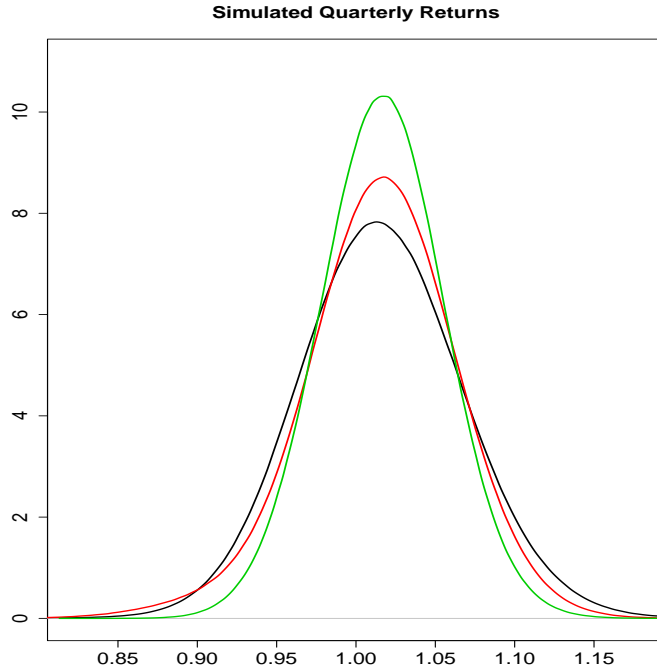


Figure 10.: Simulated quarterly returns for three different correlation models. The black curve represents the common correlation model, the red curve represents the grouped correlations model, and the green curve represents the grouped variables model.

Using Equation (1) and the simulated returns, we then examined the distributions of benefit amounts generated by the three models. As we can see in Table 2, the three correlation models produce significantly different simulated benefit amounts for our hypothetical variable annuity product. Due to the quarterly 0% minimum return (rollup) feature, it is primarily the right tail of returns that drives the patterns of benefit amounts. Thus, the common correlation model, which has the thickest right tail, results in the largest benefit level; the grouped variables model, having the thinnest right tail, produces the smallest benefit amounts, while the grouped correlations

model falls in between. For the same reason, the common correlation model produces the largest spread of benefit amounts, while the grouped variables model produces the smallest. Due to the right tail of all three quarterly returns distributions, the resulting distribution of benefit amounts is right-skewed for all three correlation models.

Table 2.: Simulated benefit amounts produced by hypothetical variable annuity product

Model	Q1	Median	Mean	Q3
Common Correlation	711	1,030	1,264	1,531
Grouped Correlations ($K = 2$)	621	868	1,021	1,243
Grouped Variables ($L = 2$)	528	709	798	960

Using the simulated asset returns and benefit amounts, we next analyzed the accumulation of insurer assets for this VA product, both in the accumulation and payout phases. Figure 11 shows the simulated accumulation of assets through time for all three models. The grouped variables model produces the largest median simulated asset values throughout the accumulation phase, with the grouped correlations model producing the smallest asset values, and the common correlation model falling in between. This is consistent with the asset returns simulated from these models, with the grouped correlations model having a somewhat thicker left tail.

Once the annuity moves from the accumulation phase to the payout phase, the effect of the larger simulated benefit amounts produced by the common correlation model begins to take over, causing the asset accumulations simulated by this model to drop more rapidly than the other models. Obviously, this effect continues to grow as the annuitant lives longer, i.e., more payments are made from the insurer's assets; in the event that the payout phase lasts longer than roughly seven or eight years, the higher benefit amount produced by the common correlation model causes the simulated asset amounts for this model to drop below those produced by the grouped correlations model.

Of great importance to writers of such annuities will be the performance of the products and associated guarantees in some of the more unfavorable scenarios. Table 3 gives the VaR (Value at Risk) and CTE (Conditional Tail Expectation) for the worst 5% and 10% of simulated asset accumulations at year 45 (i.e., after 20 years in the payout phase) under the three correlation models. According to these metrics, the simulated results for the tail are worse for the common correlation model than for either of the grouped models.

Table 3.: VaR and CTE for insurer asset accumulation amounts produced by simulated returns

Model	90% VaR	CTE 90	95% VaR	CTE 95
Common Correlation	-41,992	-55,692	-50,760	-65,541
Grouped Correlations ($K = 2$)	-37,867	-51,674	-46,859	-61,532
Grouped Variables ($L = 2$)	-9,232	-14,417	-13,522	-17,496

We also examined the effects that varying the levels of some of the other assumptions and product features had on the above simulation results. Within reasonable ranges, we found that the results were not greatly impacted by different levels of the expense charge or minimum quarterly guaranteed rate of return (rollup). Clearly, decreases in the expense charge or increases in the minimum quarterly rate will both lead to greater benefit amounts, but these effects do not change the relative results under the three models. Similarly, increasing the annual benefit rate would cause the benefit amount to increase in all cases, and a corresponding decrease in insurer assets.

The results of this section have demonstrated that for this variable annuity example, the three correlation models produce significantly differing results. In particular, the common correlation




Figure 11.: Simulated insurer asset accumulations through time. Solid lines show median values for the correlation models, while dashed lines indicate the upper and lower quartiles. Black lines represent the common correlation model, red lines represent the grouped correlations model, and green lines represent the grouped variables model.

model results in parameter inference that yields conservative results, relative to the other two models. Since it is likely that the actual correlation structure present in the data is more complex than can be represented by the common correlation model, this simpler model will serve to overprice investment guarantees such as the one seen in this example.

5. Conclusion

In this paper, we have explored the application of regime-switching models to modeling asset return streams — and specifically the correlation between pairs of asset returns — using three different correlation structures in a Bayesian framework. The common correlation model, the grouped correlations model, and the grouped variables model offer varying levels of flexibility and interpretability when modeling the correlations between asset streams. There are many applications for which accurate modeling of these between-asset correlations is important, including the pricing of investment guarantees; hence, this is a valuable area of exploration. We found that in many circumstances, the added flexibility of the more complex structures was indeed valuable and led to both better modeling of the correlations and valuable insights into asset groupings.

In particular, when we applied these models to the asset returns of nine assets (comprised of four health insurance stocks and five bank stocks), we found several interesting results. First and most importantly, there was strong evidence that the common correlation model was too restrictive and the added flexibility afforded by the grouped correlations and grouped variables model appeared to be valuable in describing the correlation structure of these assets.

Not surprisingly, the grouped correlations model appears to cluster the correlations according to their actual sector groupings, while the grouped variables model was successful in clustering the actual assets according to sector. It appeared that two groups were sufficient in both the grouped correlations and grouped variables model; allowing the possibility of a third group did not significantly improve the models.

We also found that the assignment of time points to regimes varied by model. For the common correlation model, the regimes were marked by divergent mean asset returns by sector. However, for the grouped correlations and grouped variables models, regime assignment was driven more by the clustering structure of the assets than by the means or variances of the assets. In all cases, this differs from many regime switching applications that have regimes which can be interpreted as being a high-mean/low-variance regime and a low-mean/high-variance regime.

When applied to a generic variable annuity product with a guaranteed minimum income benefit, we found that the three correlation models resulted in significantly different levels of benefits and widely varying patterns of simulated asset accumulations for the insurer. These differences were due to the variations in the distribution of simulated portfolio asset returns, which themselves resulted from the discrepancies in the inference for the underlying parameter values, in particular, the correlation elements. The simpler correlation model lacked the flexibility to allow the correlation elements to find their natural groupings; this rigidity led to the appearance of more positively correlated assets, tending to produce thicker-tailed return distributions. These differences are important to insurers, as they may drive the pricing and/or reserving of such annuity guarantees.

There are many avenues for future work with regard to these models. First, we may consider how these models perform under different types of asset return data. For example, it might be interesting to consider how the models classify assets or market indices from different countries, and which types of clustering and regime classification patterns are observed in these cases. Another line of inquiry would be to further consider the impact of using overly simplistic correlation models on the pricing and reserving of various types of investment guarantees offered by insurers in products such as variable or equity-indexed annuities or variable universal life products. While we have considered here a hypothetical annuity product with a relatively generic GMIB (Guaranteed Minimum Income Benefit) rider, the volume and complexity of guarantees offered in the marketplace is large and ever-expanding. A thorough exploration of the impact of the differences of these correlation models on the pricing and reserving of such guarantees would be of great benefit to researchers and practicing actuaries alike. Finally, it would be interesting to explore whether the results we found regarding the assignment of points to regimes would generalize to other situations involving the modeling of correlations between multiple assets and to what extent the correlation structure would inform the interpretation of the regimes as it has done here.

The presented model has been shown to work well in settings with moderate dimensionality. If, however, joint modeling of a very high number of asset streams is of interest, investigation of alternative covariance structures may be warranted. For example, implementation of an approach similar to Wang and Pillai (2013), an approach shown to be computationally efficient in high-dimensional spatial settings, might be useful. However, it may be more difficult to impose clustering structure under such an approach.

Acknowledgements

The authors would like to thank the anonymous reviewers and the editorial staff, whose comments and suggestions greatly improved the quality of this paper.

References

- Allianz (2018), ‘Allianz vision variable annuity’.
URL: <https://www.allianzlife.com/annuities/variable-annuities/vision>
- Barnard, J., McCulloch, R. and Meng, X.-L. (2000), ‘Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage’, *Statistica Sinica* **10**(4), 1281–1312.
- Bernardo, J. and Smith, A. (1994), ‘Bayesian theory wiley’, *New York*.
- Boudreault, M. and Panneton, C. (2009), ‘Multivariate models of equity returns for investment guarantees valuation’, *North American Actuarial Journal* **13**(1), 36–53.
- Costabile, M. (2017), ‘A lattice-based model to evaluate variable annuities with guaranteed minimum withdrawal benefits under a regime-switching model’, *Scandinavian Actuarial Journal* **2017**(3), 231–244.
- Da Fonseca, J. and Ziveyi, J. (2017), ‘Valuing variable annuity guarantees on multiple assets’, *Scandinavian Actuarial Journal* **2017**(3), 209–230.
- Daniels, M. J. and Kass, R. E. (2001), ‘Shrinkage estimators for covariance matrices’, *Biometrics* **57**(4), 1173–1184.
- Gao, H., Mamon, R. and Liu, X. (2015), ‘Pricing a guaranteed annuity option under correlated and regime-switching risk factors’, *European Actuarial Journal* **5**(2), 309–326.
- Hamilton, J. (1989), ‘A new approach to the economic analysis of nonstationary time series and the business cycle’, *Econometrica: Journal of the Econometric Society* **57**(2), 357–384.
- Hamilton, J. and Susmel, R. (1994), ‘Autoregressive conditional heteroskedasticity and changes in regime’, *Journal of Econometrics* **64**(1-2), 307–333.
- Haran, M. and Hughes, J. (2016), *batchmeans: Consistent Batch Means Estimation of Monte Carlo Standard Errors*, Denver, CO. R package version 1.0-3.
- Hardy, M. (2001), ‘A regime-switching model of long-term stock returns’, *North American Actuarial Journal* **5**(2), 41–53.
- Hardy, M. (2002), ‘Bayesian risk management for equity-linked insurance’, *Scandinavian Actuarial Journal* **2002**(3), 185–211.
- Hardy, M. (2003), *Investment Guarantees: Modeling and Risk Management for Equity Linked Life Insurance*, John Wiley and Sons.
- Hartman, B. M. and Groendyke, C. (2013), ‘Model selection and averaging in financial risk management’, *North American Actuarial Journal* **17**(3), 216–228.
- Hartman, B. M. and Heaton, M. J. (2011), ‘Accounting for regime and parameter uncertainty in regime-switching models’, *Insurance: Mathematics and Economics* **49**(3), 429 – 437.
URL: <http://www.sciencedirect.com/science/article/pii/S0167668711000801>
- Hsu, C.-W., Sinay, M. S. and Hsu, J. S. (2012), ‘Bayesian estimation of a covariance matrix with flexible prior specification’, *Annals of the Institute of Statistical Mathematics* **64**(2), 319–342.
- Jasra, A., Holmes, C. and Stephens, D. (2005), ‘Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling’, *Statistical Science* **20**(1), 50–67.
- Jones, G. L., Haran, M., Caffo, B. S. and Neath, R. (2006), ‘Fixed-width output analysis for markov chain monte carlo’, *Journal of the American Statistical Association* **101**(476), 1537–1547.
- Kolkiewicz, A. W. and Lin, F. S. (2017), ‘Pricing surrender risk in ratchet equity-index annuities under regime-switching lévy processes’, *North American Actuarial Journal* **21**(3), 433–457.
- Leonard, T. and Hsu, J. S. (1992), ‘Bayesian inference for a covariance matrix’, *The Annals of Statistics* **20**(4), 1669–1696.
- Liechty, J. C., Liechty, M. W. and Müller, P. (2004), ‘Bayesian correlation estimation’, *Biometrika* **91**(1), 1–14.
- Liechty, M., Liechty, J. and Müller, P. (2009), ‘The Shadow Prior’, *Journal of Computational and Graphical Statistics* **18**(2), 368–383.
- Ng, A. C.-Y. and Li, J. S.-H. (2011), ‘Valuing variable annuity guarantees with the multivariate esscher transform’, *Insurance: Mathematics and Economics* **49**(3), 393–400.
- Ng, A. C.-Y. and Li, J. S.-H. (2013), ‘Pricing and hedging variable annuity guarantees with multiasset stochastic investment models’, *North American Actuarial Journal* **17**(1), 41–62.
- R Core Team (2016), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

URL: <https://www.R-project.org/>

Ritter, C. and Tanner, M. A. (1992), ‘Facilitating the Gibbs sampler: the Gibbs stopper and the griddy-Gibbs sampler’, *Journal of the American Statistical Association* **87**(419), 861–868.

Wang, H. and Pillai, N. S. (2013), ‘On a class of shrinkage priors for covariance matrix estimation’, *Journal of Computational and Graphical Statistics* **22**(3), 689–707.

Yahoo! Inc. (2017), ‘Yahoo! finance’.

URL: <http://finance.yahoo.com/>

Yang, R. and Berger, J. O. (1994), ‘Estimation of a covariance matrix using the reference prior’, *The Annals of Statistics* **22**(3), 1195–1211.

Appendix A: Hyperparameter Settings

For both the health and bank simulation and the VA application, we used the following hyperparameter settings:

$$\alpha_\gamma = 20,$$

$$\beta_\gamma = 1,$$

$$\sigma_\lambda^2 = 1,$$

$$\mu_\lambda = 0,$$

$$\alpha_\sigma = 50,$$

$$\beta_\sigma = 0.05,$$

$$\tau^2 = 10,$$

$$\nu^2 = 0.0001.$$

In all cases the hyperparameters were chosen to be relatively non-informative where possible. Some needed to be adjusted to improve convergence.

Appendix B: Inference Details

This appendix gives the parameter estimates and their associated MCMC standard errors in the common correlation and grouped correlations models for the health and bank data. The standard errors (given in parentheses) are estimated using the batch means method (Jones et al. 2006, Haran and Hughes 2016).

Asset Symbol	Common Correlation Model ($K = 1$)		Grouped Correlations Model ($K = 2$)	
	Primary Regime	Secondary Regime	Primary Regime	Secondary Regime
UNH	-0.00771 (0.00007)	0.01338 (0.00053)	-0.00444 (0.00002)	-0.01919 (0.00013)
AET	-0.00748 (0.00013)	0.02065 (0.00090)	-0.00175 (0.00003)	-0.05142 (0.00020)
CI	-0.00681 (0.00012)	0.01371 (0.00081)	-0.00183 (0.00003)	-0.05549 (0.00019)
HUM	-0.00772 (0.00015)	0.01913 (0.00100)	-0.00245 (0.00003)	-0.04417 (0.00021)
BAC	0.00703 (0.00025)	-0.04039 (0.00137)	0.00129 (0.00003)	0.00192 (0.00015)
KEY	0.00470 (0.00020)	-0.03080 (0.00112)	-0.00062 (0.00003)	0.02340 (0.00014)
WFC	0.00179 (0.00013)	-0.02105 (0.00067)	-0.00103 (0.00003)	0.00120 (0.00012)
C	0.00601 (0.00021)	-0.03522 (0.00118)	0.00113 (0.00003)	0.00166 (0.00018)
JPM	0.00256 (0.00017)	-0.02718 (0.00089)	-0.00079 (0.00003)	-0.00485 (0.00016)

Table 4.: Posterior means for μ_p parameters in primary and secondary regimes ($K = 1, K = 2$)

Appendix C: Fund Information

This appendix lists the investment options used for our variable annuity analysis; these investment options were taken from the Allianz Vision Variable Annuity product.

Symbol	Company	Category
DGAGX	Dreyfus	Large Blend
FSICX	Fidelity	Multisector Bond
MDDVX	BlackRock	Large Value
MDLOX	BlackRock	World Allocation
MMUFX	MFS	Utilities
MRBFX	MFS	Intermediate-Term Bond
MWTRX	Metropolitan West	Intermediate-Term Bond
OIGAX	Oppenheimer	Foreign Large Growth
PAAIX	PIMCO	Tactical Allocation
PCRIX	PIMCO	Commodities Broad Basket
PEBIX	PIMCO	Emerging Markets Bond
PGBIX	PIMCO	World Bond
PGOVX	PIMCO	Long Government
PHIYX	PIMCO	High Yield Bond
PIGLX	PIMCO	World Bond
PISIX	PIMCO	Foreign Large Blend
PRRIX	PIMCO	Inflation-Protected Bond
PTLDX	PIMCO	Short-Term Bond
PTTRX	PIMCO	Intermediate-Term Bond

Table 5.: List of investment options used in the variable annuity product example. Company and investment categories for each fund are also shown.