Regime Switching Model Calibration

Grant Holtes

March 25, 2026

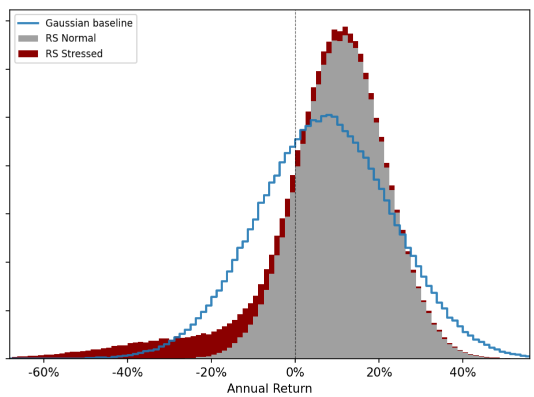

1 Introduction

Regime-switching models, also known as Markov-switching models, are a method for generating simulated return data from a non-normal distribution with conditional correlations. The model specifies two or more distributions, or "regimes", from which returns are generated, with the active regime changing at each time step within a simulation according to specified transition probabilities. Typically the regimes represent a "normal market" and a "stressed market".

Each regime is described as a multivariate normal distribution with a specified mean, volatility, and correlation structure between assets. While each regime is itself normally distributed, the unconditional distribution produced by the mixture exhibits fat tails, negative skewness, and time-varying volatility. A key advantage of this framework is that it allows semi-technical investment practitioners to directly critique the assumptions underpinning each regime and the switching process, which is more difficult for alternatives such as copula models or stochastic volatility models.

A core challenge with regime-switching models is their integration with an investment committee process that relies on capital market assumptions (CMAs). CMAs typically specify an expected return and volatility over some horizon (e.g. 10 years). This is straightforward to accommodate in a single-regime Gaussian return-generating process, but the presence of multiple regimes means that the unconditional moments implied by the mixture will not automatically match the CMA targets.

This note specifies a method for calibrating a two-regime Markov-switching model to a set of CMAs for expected geometric return and volatility. The approach generalises naturally to more than two regimes.

The method works by treating the normal regime's expected return and volatility as unknown variables. These are optimised numerically so that the overall simulation's expected geometric returns and volatilities, computed across all regime-blended paths, match the capital market assumptions exactly. All other inputs are specified directly by the investment team.

2 Model and Calibration Method

2.1 Inputs

The inputs required from the CMA process are:

1. Expected geometric returns over the simulated horizon for each asset class, $\bar{g}_k$.
2. Expected volatility for each asset class, $\bar{\sigma}_k$.

The regime-switching assumptions required are:

1. The arithmetic return $\mu_k^{(s)}$, volatility $\sigma_k^{(s)}$, and correlation matrix $\Omega^{(s)}$ in the stressed regime.
2. The correlation matrix in the normal market regime, $\Omega^{(n)}$.
3. The probabilities of transitions between regimes: p_{ns} (normal to stressed) and p_{sn} (stressed to normal).[1]

Additionally, a return floor $r_{\min}$ may be specified to prevent any single-year return from falling below a given threshold (e.g. -95%). The normal-regime arithmetic mean $\boldsymbol{\mu}^{(n)}$ and volatility $\boldsymbol{\sigma}^{(n)}$ are left as free variables and are not inputs. These are calibrated so that the blended simulation reproduces the target unconditional moments $\bar{g}_k$ and $\bar{\sigma}_k$, which is the core of the calibration process.

2.2 Simulation Methodology

We generate S Monte Carlo paths of N-year annual returns for K asset classes under the two-regime Markov-switching framework described above. The procedure consists of five steps.

Step 1: Regime path generation. For each path $s = 1, \ldots, S$, we simulate an N-step Markov chain $\{R_{s,t}\}_{t=1}^{N}$ with state space $\{0, 1\}$ (normal, stressed). The initial state is drawn from the stationary distribution $\pi^{(s)} = p_{ns}/(p_{ns} + p_{sn})$, and subsequent transitions follow

$$P(R_{s,t} = 1 \mid R_{s,t-1} = 0) = p_{ns}, \qquad P(R_{s,t} = 0 \mid R_{s,t-1} = 1) = p_{sn}.$$

This produces an $S \times N$ binary mask, $R_{s,t}$, that determines which regime is at play in each year of each simulation.

[1] Or equivalently the expected regime duration, which is $1/P(\text{switch out of regime})$.

Step 2: Stressed-regime return generation. For each year t, let $\mathcal{S}_t = \{s : R_{s,t} = 1\}$ denote the set of paths in the stressed regime.

We draw $|\mathcal{S}_t|$ samples from a multivariate normal distribution with mean $\boldsymbol{\mu}^{(s)}$ and covariance $\Sigma^{(s)} = D(\boldsymbol{\sigma}^{(s)})\,\Omega^{(s)}\,D(\boldsymbol{\sigma}^{(s)})$, where $D(\cdot)$ denotes the diagonal matrix.

Correlated samples are obtained via the Cholesky factorisation

$$\Sigma^{(s)} = L^{(s)} L^{(s)\top}$$

with

$$\mathbf{r}_{s,t} = L^{(s)}\mathbf{z}_{s,t} + \boldsymbol{\mu}^{(s)}$$

for standard normal $\mathbf{z}_{s,t} \sim \mathcal{N}(\mathbf{0}, I_K)$ of shape $|St| \times K$, for each asset in the K assets.

To eliminate sampling error in the covariance, we apply a Cholesky correction to each year's cross-section. Let $\hat{\Sigma}_t$ be the sample covariance of $\{\mathbf{r}_{s,t}\}_{s\in\mathcal{S}_t}$ and $\hat{\Sigma}_t = \hat{L}_t\hat{L}_t^\top$ its Cholesky factorisation. The corrected samples are

$$\tilde{\mathbf{r}}_{s,t} = \left(\hat{L}_t^{-1}\,\mathbf{r}_{s,t}\right)^\top L^{(s)\top},$$

which ensures $\mathrm{Cov}(\tilde{\mathbf{r}}_{\cdot,t}) = \Sigma^{(s)}$ exactly for the stressed returns in each year t.

Step 3: Normal-regime return generation. We generate $S \times N$ standard normal draws and apply the same Cholesky procedure using an initial covariance matrix constructed from $\Omega^{(n)}$ with an arbitrary initial volatility. After per-year covariance correction, each asset's returns are globally standardised to zero mean and unit variance. This produces a set of standardised correlated normals $\mathbf{Z}^{(n)} \in \mathbb{R}^{S\times N\times K}$ with the correlation structure of $\Omega^{(n)}$ baked in. For any candidate normal-regime mean $\boldsymbol{\mu}^{(n)}$ and volatility $\boldsymbol{\sigma}^{(n)}$, the normal-regime returns are then

$$r^{(n)}_{s,t,k} = \mu^{(n)}_k + \sigma^{(n)}_k \cdot Z^{(n)}_{s,t,k},$$

which preserves the target correlation since scaling and shifting are affine operations on the marginals.

Step 4: Calibration of normal-regime parameters. The normal-regime arithmetic mean $\boldsymbol{\mu}^{(n)}$ and volatility $\boldsymbol{\sigma}^{(n)}$ are determined by finding the required values such that the overall sample geometric return and volatility match the CMA geometric return g and volatility σ for all k assets.

This is achieved by solving

$$\min_{\boldsymbol{\mu}^{(n)},\boldsymbol{\sigma}^{(n)}} \sum_{k=1}^{K} \left[(\hat{g}_k - \bar{g}_k)^2 + (\hat{\sigma}_k - \bar{\sigma}_k)^2 \right],$$

where the blended return for sample s in year t is

$$r_{s,t,k} = \begin{cases} \tilde{r}^{(s)}_{s,t,k} & \text{if } R_{s,t} = 1, \\ r^{(n)}_{s,t,k} & \text{if } R_{s,t} = 0, \end{cases}$$

with a return floor applied as $r_{s,t,k} \leftarrow \max(r_{s,t,k},\, r_{\min})$.
The simulated geometric return is

$$\hat{g}_k = \bar{\mathbb{E}}_s \left[\left(\prod_{t=1}^{N} (1 + r_{s,t,k}) \right)^{1/N} \right] - 1$$

and the simulated volatility is

$$\hat{\sigma}_k = \mathrm{std}(\{r_{s,t,k}\}_{s,t})$$

The optimisation is performed via Nelder–Mead with adaptive restarts.

Step 5: Output. The final output is the $S \times N \times K$ array of blended annual returns, together with the fitted normal-regime parameters. The return floor is applied within the optimisation loop, ensuring the calibrated moments account for any truncation effects.

3 Example Simulation

3.1 Input Assumptions

The following setup defines a example, with all inputs being for demonstration purposes only, without any historical basis or return forecasting. The simulation is calibrated to a universe of five asset classes over a 20-year horizon with 20,000 Monte Carlo paths. A return floor of −95% is applied to all single-year returns.

Capital market assumptions. Table 1 presents the expected geometric returns and volatility from the CMA process, alongside the stressed-regime arithmetic means and volatility.

Table 1: Asset class return and volatility assumptions.

	Unconditional (CMA)		Stressed Regime	
Asset Class	Geo. Return	Volatility	Arith. Mean	Volatility
Australian Equity	6.0%	16.0%	−18.0%	22.0%
International Equity	7.0%	15.0%	−15.0%	20.0%
Bonds	4.0%	5.0%	5.0%	8.0%
Property	4.5%	14.0%	−17.0%	22.0%
Cash	3.0%	2.0%	3.0%	3.0%

Regime transition probabilities. The Markov transition probabilities are $p_{ns} = 3/26 \approx 11.5\%$ and $p_{sn} = 75.0\%$, implying expected regime durations of approximately 8.7 years (normal) and 1.3 years (stressed).

Correlation matrices. Tables 2 and 3 present the assumed correlation structures for the normal and stressed regimes respectively. The stressed regime features materially higher correlations among growth assets, consistent with the empirical observation that diversification benefits diminish during market dislocations. Bonds and Cash correlations are held constant across regimes.

Table 2: Normal-regime correlation matrix, $\Omega^{(n)}$.

	Aus Eq.	Intl Eq.	Bonds	Property	Cash
Australian Equity	1.00	0.60	0.10	0.55	0.00
International Equity	0.60	1.00	0.05	0.50	0.00
Bonds	0.10	0.05	1.00	0.15	0.40
Property	0.55	0.50	0.15	1.00	0.00
Cash	0.00	0.00	0.40	0.00	1.00

Table 3: Stressed-regime correlation matrix, $\Omega^{(s)}$.

	Aus Eq.	Intl Eq.	Bonds	Property	Cash
Australian Equity	1.00	0.95	0.40	0.90	0.00
International Equity	0.95	1.00	0.35	0.90	0.00
Bonds	0.40	0.35	1.00	0.30	0.40
Property	0.90	0.90	0.30	1.00	0.00
Cash	0.00	0.00	0.40	0.00	1.00

3.2 Simulation results

Figure 1 shows the distributions of yearly returns under the regime switching model, which display the expected fatter tails and high skew for the equity and property asset classes, which have severe outcomes in the stressed regime.

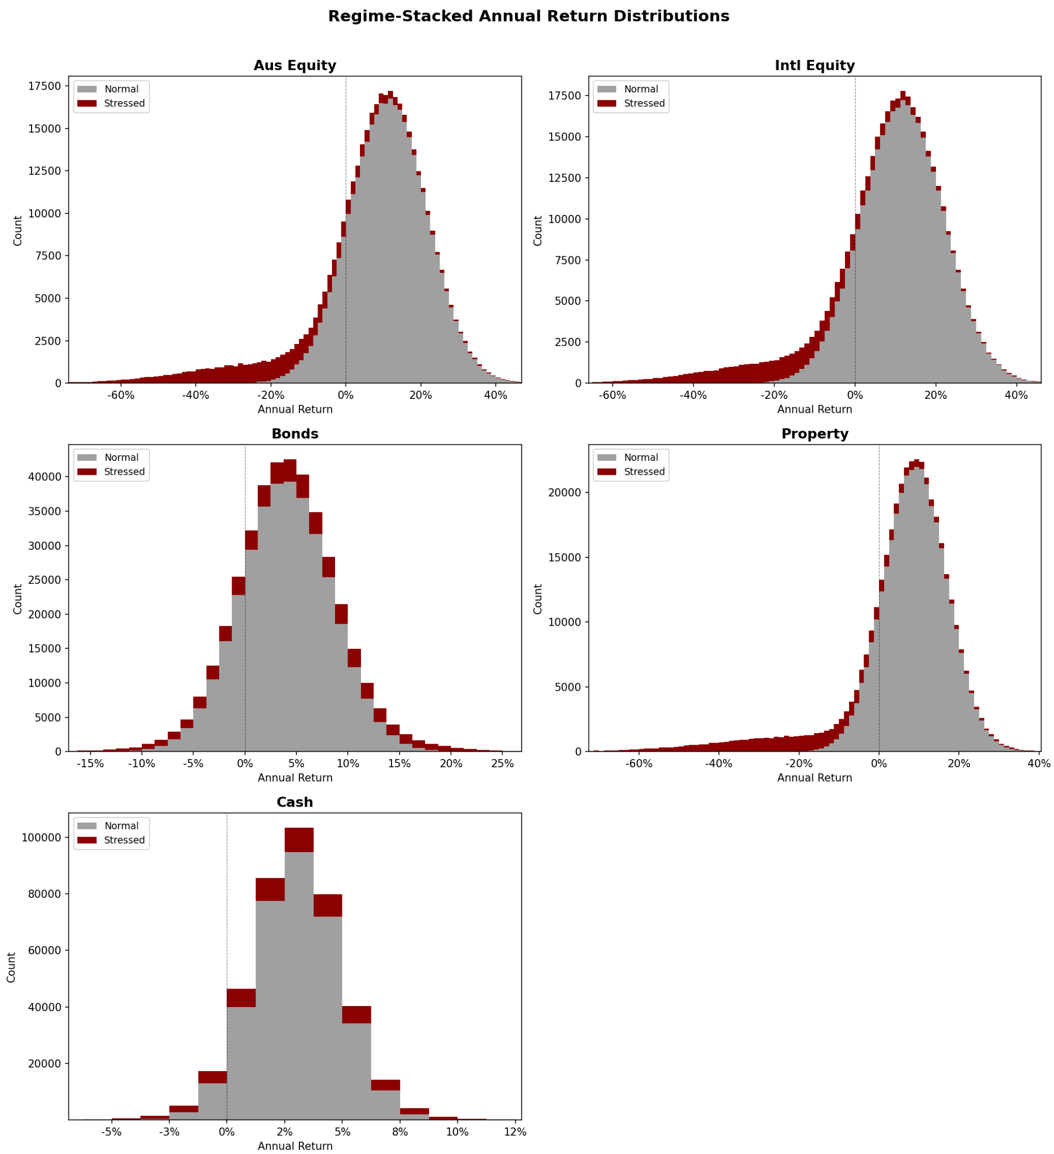

Figure 1: Distribution of yearly returns across all simulations, split by regime

3.3 Gaussian Baseline

A single-regime Gaussian baseline is constructed for comparison. The baseline uses the same target geometric returns and volatilities, with the unconditional correlation matrix implied by the regime-switching simulation. The same standardised-normal generation, Cholesky covariance correction, and moment-matching optimisation procedure are applied, differing only in the absence of regime switching. This ensures the baseline matches the regime-switching model on all first and second unconditional moments, as confirmed by the matching first two columns of Table 4, isolating the contribution of regime structure to higher-order properties.

Table 4: Moment comparison: regime-switching model vs. single-regime Gaussian baseline. Both models are calibrated to identical geometric returns and volatilities. Skewness and excess kurtosis isolate the higher-order effects of regime switching.

Asset Class	Model	Geo. Return p.a.	Volatility	Skewness	Excess Kurtosis
Australian Equity	Regime-Switching	6.00%	16.00%	−1.441	3.591
	Gaussian	6.00%	16.00%	0.003	−0.006
International Equity	Regime-Switching	7.00%	15.00%	−1.286	3.037
	Gaussian	7.00%	15.00%	−0.005	0.001
Bonds	Regime-Switching	4.00%	5.00%	0.122	1.155
	Gaussian	4.00%	5.00%	0.001	−0.005
Property	Regime-Switching	4.50%	14.00%	−1.920	5.728
	Gaussian	4.50%	14.00%	0.000	−0.001
Cash	Regime-Switching	3.00%	2.00%	−0.019	0.705
	Gaussian	3.00%	2.00%	0.002	0.006

These differences result in larger maximum drawdown across the equity and property asset classes, as shown in Figure 2, and more extreme left tail events, as in Table 5.

Table 5: Tail risk comparison: probability of annual returns falling below selected thresholds. The Gaussian baseline exhibits higher probabilities of moderate losses but materially lower probabilities of severe losses, reflecting the absence of fat tails.

		$P(r <$ threshold$)$			
Asset Class	Model	-10%	-20%	-30%	-40%
Australian Equity	Regime-Switching	10.30%	6.25%	3.91%	2.10%
	Gaussian	14.18%	4.45%	1.00%	0.16%
International Equity	Regime-Switching	9.36%	5.42%	3.02%	1.40%
	Gaussian	11.47%	3.12%	0.57%	0.07%
Bonds	Regime-Switching	0.46%	0.01%	—	—
	Gaussian	0.24%	0.00%	—	—
Property	Regime-Switching	8.94%	5.93%	3.70%	1.96%
	Gaussian	13.55%	3.47%	0.59%	0.05%
Cash	Regime-Switching	—	—	—	—
	Gaussian	—	—	—	—

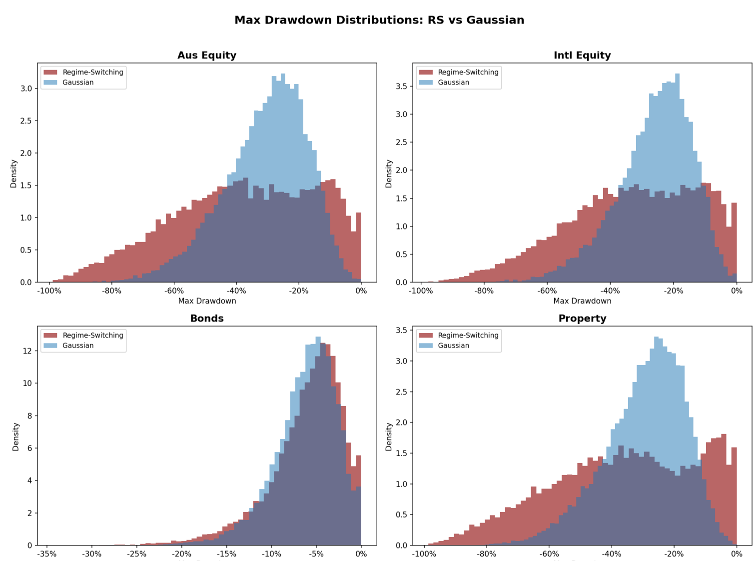

Figure 2: Distribution maximum drawdown by simulation, Regime-Switching and Gaussian models

4 Limitations

As volatility is zero-bounded, it is possible to define combinations of CMAs and stressed volatility that cannot be reconciled with any feasible normal market parameters. This can be detected by the solver failing to match the CMA inputs and the fitted normal market volatility being near zero. If this is the case, the stressed market assumptions and / or the regime transition probabilities should be reviewed.

Note:

Views expressed are the author's, and may differ from those of JANA investments. This material does not constitute investment advice and should not be relied upon as such. Investors should seek independent advice before making investment decisions. Past performance cannot guarantee future results. The charts and tables are shown for illustrative purposes only.

5 Appendix

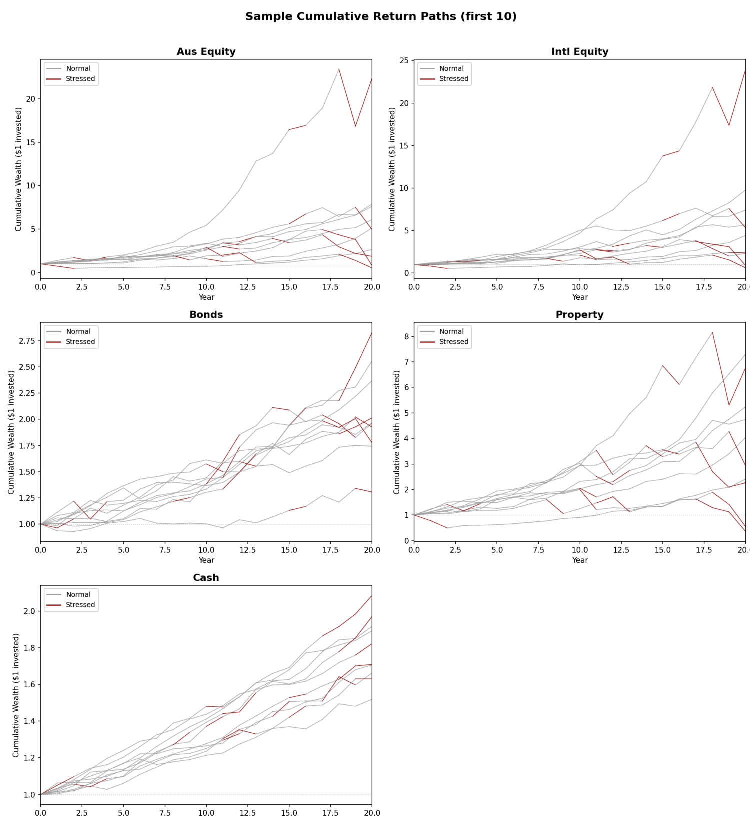

Figure 3: Cumulative return for the first 10 simulations by asset class, split by regime