Computational Approaches to Stochastic Portfolio Optimization Under Regime-Switching Dynamics

by

Elena Marchetti

Department of Statistical Science
in the Graduate School of
Duke University

Date: ______________

Approved by:

Dr. Jonathan Whitfield, Supervisor

Dr. Priya Nair

Dr. Marcus Chen

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor of Philosophy
in the Department of Statistical Science
in the Graduate School of Duke University

2026

Abstract

This dissertation develops computational frameworks for optimizing investment portfolios when financial markets exhibit abrupt shifts in volatility and correlation structure — a phenomenon known as regime switching. Traditional mean-variance optimization, rooted in Gaussian and stationary assumptions, systematically underestimates tail risk during turbulent regimes and over-allocates to assets whose correlations converge in downturns. We propose a class of hidden Markov regime-switching models coupled with stochastic gradient-based solvers that jointly estimate latent regimes and construct forward-looking efficient frontiers.

The core contribution is a two-stage procedure: first, a Bayesian changepoint detection algorithm segments historical returns into distinct volatility regimes using reversible-jump Markov chain Monte Carlo; second, a scenario-based convex optimizer — implemented via the alternating direction method of multipliers — solves the resulting non-Gaussian portfolio problem. We prove that the optimizer converges to the true conditional value-at-risk frontier under mild regularity conditions on the regime process.

Empirical evaluation uses synthetic data calibrated to the post-2008 regulatory environment and backtests on a 30-year panel of equity, fixed-income, and commodity indices. The regime-aware portfolios reduce maximum drawdown by 18–24% relative to static Markowitz benchmarks during stress periods while maintaining comparable Sharpe ratios in calm markets. We extend the framework to incorporate transaction costs, short-sale constraints, and ESG tilts, demonstrating its flexibility for institutional asset allocators. The methods are packaged as an open-source `R` library, `regimePort`, with bindings to `Stan` for posterior sampling and

`CVXR` for disciplined convex programming.

Keywords: portfolio optimization, regime switching, hidden Markov models, convex optimization, conditional value-at-risk, Bayesian changepoint detection, ADMM, risk management.

Dedication

To my parents, who taught me that curiosity

is the only compass worth following.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

Modern portfolio theory, as formalized by Markowitz [10], remains the dominant paradigm for institutional asset allocation. Its elegance — distill risk into a covariance matrix, maximize expected return per unit of variance — has proven remarkably durable across seven decades of practice. Yet the 2008 financial crisis, the COVID-19 volatility spike of 2020, and the rate-hiking cycle of 2022–2023 each exposed the same structural vulnerability: correlations that appear modest in calm regimes collapse toward unity precisely when diversification is most needed.

This dissertation addresses a concrete question: can we build portfolio optimizers that are *regime-aware* — that anticipate shifts in the return-generating process rather than treating history as a single stationary draw? We answer affirmatively by combining tools from Bayesian time-series modeling and convex optimization into a unified computational pipeline.

1.2 The Regime-Switching Perspective

Financial returns exhibit well-documented stylized facts that violate the i.i.d. Gaussian benchmark: volatility clustering, fat tails, leverage effects, and time-varying correlations [4]. Hidden Markov models (HMMs) provide a parsimonious framework for capturing these features. A K-state HMM posits that returns $\mathbf{r}_t \in \mathbb{R}^p$ are drawn from:

$$\mathbf{r}_t \mid s_t = k \;\sim\; \mathcal{N}_p(\boldsymbol{\mu}_k,\ \boldsymbol{\Sigma}_k)\,, \qquad k \in \{1, \dots, K\}, \tag{1.1}$$

where the latent state s_t evolves according to a first-order Markov chain with transition matrix $\mathbf{Q} = (q_{ij})$, and $q_{ij} = \Pr(s_t = j \mid s_{t-1} = i)$.

Definition 1 (Regime). A *regime* is a persistent state of the return-generating process characterized by a distinct pair $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$. Regimes typically correspond to macroeconomic environments: expansion (low volatility, positive drift), contraction (elevated volatility, negative drift), and crisis (extreme volatility, correlation breakdown).

The critical insight for portfolio construction is that the forward-looking distribution of returns is a *mixture*:

$$\Pr(\mathbf{r}_{T+1} \mid \mathcal{F}_T) = \sum_{k=1}^{K} \xi_T(k) \;\cdot\; \mathcal{N}_p(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \tag{1.2}$$

where $\xi_T(k) = \Pr(s_{T+1} = k \mid \mathcal{F}_T)$ is the filtered state probability given information up to time T. When regimes are persistent, ξ_T concentrates mass on a single component, but the *possibility* of transition fattens the tails.

1.3 Contributions

This dissertation makes four principal contributions:

1. **Changepoint-regularized HMM estimation.** We introduce a reversible-jump MCMC sampler that jointly infers the number of regimes K and their parameters, penalizing superfluous states through a birth–death process on the transition matrix. This avoids the ad-hoc information-criterion selection common in the literature.

2. **Regime-conditional CVaR optimization.** We formulate the portfolio problem as

minimizing conditional value-at-risk under a mixture-of-Gaussians return model, transforming a non-convex objective into a second-order cone program solvable via ADMM. Convergence guarantees are established under α-mixing conditions on the regime chain.

3. **Empirical calibration and backtesting.** Using a synthetic data engine calibrated to post-2008 market structure and out-of-sample tests on a multi-asset panel spanning 1994–2024, we document consistent improvements in downside-risk metrics without sacrificing upside capture.

4. **Open-source software.** The `regimePort` package for `R` provides a modular implementation interfacing with `Stan` for Bayesian computation and `CVXR` for convex optimization, lowering the barrier for practitioners.

1.4 Thesis Outline

Chapter 2 surveys the literature on regime-switching models in finance, convex risk measures, and computational Bayesian methods. Chapter 3 develops the changepoint-regularized HMM estimator with full algorithmic details and convergence diagnostics. Chapter 4 formulates the regime-conditional CVaR portfolio problem and derives the ADMM algorithm. Chapter 5 presents the simulation study and empirical backtests. Chapter 6 discusses extensions to transaction costs, ESG constraints, and multi-period rebalancing. Chapter 7 concludes with limitations and directions for future work.

Chapter 2

Literature Review

2.1 Markowitz and Its Discontents

Markowitz [10] framed portfolio selection as a quadratic program: choose weights $\mathbf{w} \in \mathbb{R}^p$ to

$$\min_{\mathbf{w}} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \quad \text{s.t.} \quad \mathbf{w}^\top \boldsymbol{\mu} \geq \mu_0, \quad \mathbf{w}^\top \mathbf{1} = 1. \tag{2.1}$$

The solution delivers the familiar efficient frontier — a parabola in mean-variance space. Its weaknesses are equally familiar: $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are notoriously difficult to estimate (the "Markowitz error-maximization" problem [11]), and the Gaussian assumption ignores tail dependence.

Subsequent decades produced a rich literature on robust optimization [5], shrinkage estimation [9], and Black–Litterman blending [3]. These refinements improve *static* estimates but do not address the core problem: the data-generating process itself is non-stationary.

2.2 Hidden Markov Models in Finance

Hamilton [8] introduced regime-switching to economics, modeling GDP growth as alternating between expansion and recession states. The framework migrated to finance through applications in volatility modeling [14], interest rate dynamics [1], and tactical asset allocation [7].

Estimation is typically performed via the expectation-maximization (EM) algorithm or Bayesian MCMC. The EM approach is fast but provides only point estimates; MCMC yields full posterior distributions but scales poorly with the number of assets p. A key open problem — which Chapter 3 addresses — is determining K, the number of regimes, without resorting to information criteria that penalize complexity in an ad-hoc manner.

Table 2.1: Comparison of regime-switching estimation methods.

Method	**Selects K?**	**Uncertainty**	**Scalability**
EM + BIC	Yes (criterion)	Point estimate only	$\mathcal{O}(TKp^2)$
Hamilton filter	No	Filtered probabilities	$\mathcal{O}(TK^2)$
Gibbs sampler	No	Full posterior	$\mathcal{O}(TK^2p^2)$ per iteration
RJ-MCMC (this work)	Yes (inference)	Full posterior incl. K	$\mathcal{O}(TK^2_{\max}p^2)$ per iter.

2.3 Convex Risk Measures

Value-at-Risk (VaR) — the α-quantile of the loss distribution — became the regulatory standard through the Basel Accords. Its theoretical deficiencies are well documented: VaR is not a coherent risk measure [2] because it fails the sub-additivity axiom. Conditional Value-at-Risk (CVaR), defined as the expected loss exceeding VaR,

$$\mathrm{CVaR}_\alpha(L) = \mathbb{E}[L \mid L \geq \mathrm{VaR}_\alpha(L)]\,, \tag{2.2}$$

is coherent and convex [12]. Rockafellar and Uryasev [13] showed that CVaR minimization can be cast as a linear program when the return distribution is represented by scenarios — an insight we exploit in Chapter 4, extending it to mixture-of-Gaussians representations.

Definition 2 (Coherent Risk Measure). A risk measure $\rho : \mathcal{X} \to \mathbb{R}$ is *coherent* if it satisfies (i) monotonicity: $X \leq Y \Rightarrow \rho(X) \geq \rho(Y)$, (ii) sub-additivity: $\rho(X+Y) \leq \rho(X) + \rho(Y)$, (iii) positive homogeneity: $\rho(\lambda X) = \lambda\rho(X)$ for $\lambda \geq 0$, and (iv) translational invariance: $\rho(X+c) = \rho(X) - c$.

Chapter 3

Bayesian Changepoint-Regularized HMM

3.1 Model Specification

We observe a p-dimensional return series $\{\mathbf{r}_t\}_{t=1}^{T}$. The generative model is:

Hierarchical Model

$$K \sim \text{Poisson}^{+}(\lambda_0) \tag{3.1}$$

$$q_{ii} \sim \text{Beta}(a_0, b_0), \qquad i = 1, \ldots, K \tag{3.2}$$

$$\boldsymbol{\mu}_k \sim \mathcal{N}_p(\mathbf{0}, \tau^{-1}\mathbf{I}), \qquad k = 1, \ldots, K \tag{3.3}$$

$$\boldsymbol{\Sigma}_k \sim \mathcal{W}_p^{-1}(\nu_0, \mathbf{S}_0), \qquad k = 1, \ldots, K \tag{3.4}$$

$$s_1 \sim \text{Categorical}(\boldsymbol{\pi}_0)$$

$$s_t \mid s_{t-1} \sim \text{Categorical}(\mathbf{q}_{s_{t-1},\cdot}), \qquad t = 2, \ldots, T$$

$$\mathbf{r}_t \mid s_t \sim \mathcal{N}_p(\boldsymbol{\mu}_{s_t}, \boldsymbol{\Sigma}_{s_t}), \qquad t = 1, \ldots, T$$

The innovation is the prior on K: a zero-truncated Poisson with rate λ_0, combined with a reversible-jump mechanism [6] that proposes birth, death, and split-merge moves on the state space.

3.2 Algorithm

Algorithm **??** sketches the main loop. At each iteration, with probabilities b_K, d_K, and $1 - b_K - d_K$, we propose a birth, death, or within-model move respectively.

Table 3.1: Reversible-jump MCMC for regime-switching HMM.

Step	**Operation**
1	Initialize $K^{(0)}$, $\boldsymbol{\theta}^{(0)} = \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \mathbf{Q}\}$, and $s_{1:T}^{(0)}$
2	**for** $m = 1, \ldots, M$ **do**
3	Sample $u \sim \text{Uniform}(0, 1)$
4	**if** $u < b_K$: propose birth move, accept with prob. α_{birth}
5	**else if** $u < b_K + d_K$: propose death move, accept with prob. α_{death}
6	**else**: update $\boldsymbol{\theta}$ via Gibbs, sample $s_{1:T}$ via forward-filtering backward-sampling
7	**end for**

The birth move draws a new state $(K + 1)$ by splitting an existing regime randomly; the death move merges two adjacent regimes. Acceptance probabilities are derived from the usual Metropolis–Hastings–Green ratio, with the Jacobian of the dimension-changing transformation absorbed into the proposal ratio [6].

3.3 Simulation Study

We evaluate the sampler on synthetic data with known ground truth. Table 3.2 reports recovery rates for K across 500 replications under three data-generating processes.

Table 3.2: Regime count recovery across simulation settings ($N = 500$ replications).

True K	$T = 500$	$T = 1000$	$T = 2000$	**Avg. Runtime (s)**
$K = 2$	0.94	0.97	0.99	42.3
$K = 3$	0.88	0.93	0.97	67.8
$K = 4$	0.79	0.87	0.94	98.2

Recovery is excellent for $K \leq 3$ with $T \geq 1000$, broadly adequate for $K = 4$, and the sampler correctly avoids over-fitting (the mean posterior K exceeds the truth by at most 0.18 across settings).

Bibliography

[1] A. Ang and G. Bekaert, "International asset allocation with regime shifts," *Review of Financial Studies*, vol. 15, no. 4, pp. 1137–1187, 2002.

[2] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.

[3] F. Black and R. Litterman, "Global portfolio optimization," *Financial Analysts Journal*, vol. 48, no. 5, pp. 28–43, 1992.

[4] R. Cont, "Empirical properties of asset returns: stylized facts and statistical issues," *Quantitative Finance*, vol. 1, no. 2, pp. 223–236, 2001.

[5] F. J. Fabozzi, P. N. Kolm, D. A. Pachamanova, and S. M. Focardi, *Robust Portfolio Optimization and Management*. Hoboken, NJ: Wiley, 2007.

[6] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.

[7] M. Guidolin and A. Timmermann, "Asset allocation under multivariate regime switching," *Journal of Economic Dynamics and Control*, vol. 31, no. 11, pp. 3503–3544, 2007.

[8] J. D. Hamilton, "A new approach to the economic analysis of nonstationary time series and the business cycle," *Econometrica*, vol. 57, no. 2, pp. 357–384, 1989.

[9] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *Journal of Multivariate Analysis*, vol. 88, no. 2, pp. 365–411, 2004.

[10] H. Markowitz, "Portfolio selection," *Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.

[11] R. O. Michaud, "The Markowitz optimization enigma: is 'optimized' optimal?" *Financial Analysts Journal*, vol. 45, no. 1, pp. 31–42, 1989.

[12] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[13] R. T. Rockafellar and S. Uryasev, "Conditional value-at-risk for general loss distributions," *Journal of Banking and Finance*, vol. 26, no. 7, pp. 1443–1471, 2002.

[14] A. Timmermann, "Moments of Markov switching models," *Journal of Econometrics*, vol. 96, no. 1, pp. 75–111, 2000.