Beyond Markowitz: Robust Portfolio Optimization under Market Frictions and Regime Switching

Mate Sandor Fonyi
Aishwarya Padmanaban sudha
Yanrong Wang
Tanishka Nirala

NUS Investment Society
Quantitative Finance Department
2026

Abstract

This research systematically stress-tests and evolves classical asset allocation strategies by relaxing the idealized assumptions inherent in Modern Portfolio Theory (MPT). Classical Mean-Variance Optimization (MVO) relies on normally distributed returns, frictionless markets, and stable correlation structures, which frequently break down during market distress. Utilizing a 25-asset empirical universe, this study evaluates baseline and advanced models across historical crisis horizons and 100,000-path Monte Carlo simulations. By introducing statistical realism through Skewed-t distributions and dynamic correlation (DCC-GARCH), alongside an "Execution Gauntlet" of linear transaction costs and latency, we demonstrate that static MVO acts as an error-maximizer that suffers severe alpha decay under friction. While Hierarchical Risk Parity (HRP) provides a robust topological defense against correlation contagion , it sacrifices upside participation during growth regimes. To resolve this dichotomy, this paper proposes a Regime-Switching Hidden Markov Model (rHMM-MVO). By utilizing an endogenous feature space to dynamically toggle between the return-seeking Enhanced MVO and the defensive HRP topology, the rHMM-MVO preserves alpha, reduces turnover, and mitigates catastrophic tail risks in complex, real-world environments.

Contents

1 Introduction

Modern Portfolio Theory (MPT) and its cornerstone, Mean-Variance Optimization (MVO), have long served as the foundational frameworks for quantitative asset allocation [12]. However, classical MVO relies heavily on idealized assumptions—namely, normally distributed returns, frictionless markets, and stable correlation structures. While mathematically elegant in a theoretical vacuum, these assumptions frequently lead to suboptimal performance and excessive turnover in real-world trading environments. During periods of market stress, correlations between traditionally diversified assets often converge [10], and asset returns exhibit significant fat tails and skewness, exposing the limitations of variance as a complete risk descriptor.

This project, *Beyond Markowitz: Robust Portfolio Optimization under Market Frictions and Regime Switching,* seeks to systematically stress-test and evolve these classical asset allocation strategies. Our primary objective is to quantify the sensitivity of optimal portfolio weights by replacing theoretical assumptions with empirical realities, thereby building a more resilient allocation framework.

We establish an empirical foundation by constructing a diverse, multi-asset investment universe characterized by mixed liquidity and non-normal return profiles. Baseline models, including classical MVO and Hierarchical Risk Parity (HRP)—a clustering-based approach designed to address correlation matrix instability[11]—are first evaluated in an idealized, frictionless environment using normally distributed synthetic data to establish a performance ceiling.

To expose the vulnerabilities of classical models, we introduce statistical realism. Utilizing Student's t-distributions, t-Copulas, and empirical bootstrapping, we simulate dynamic correlation drift, heavy tails, and volatility clustering. Beyond statistical flaws, tangible market frictions are integrated into the testing environment. Enforcing execution latency, lagged data intake, and realistic transaction cost modeling allows for a rigorous evaluation of capital efficiency and net-of-fees performance.

Addressing these observed vulnerabilities requires comprehensive enhancements to the classical framework, including advanced covariance forecasting and a transition to robust risk objectives such as Expected Shortfall (ES) and Conditional Value at Risk (CVaR). Ultimately, this research proposes a dynamic, regime-switching model (rHMM-MVO). By utilizing a Gaussian Hidden Markov Model trained on rolling log-returns [7], this approach detects hidden market states and dynamically blends regime-conditional parameters (μ and Σ) to adapt to shifting market volatility and structural breaks [1]. Through a rigorous comparative analysis, this report demonstrates that incorporating regime-aware logic and addressing market frictions can effectively filter out noise, reduce unnecessary trading costs, and preserve alpha in complex, real-world environments.

1.1 Project Objectives and Scope

The primary objective of this project is to systematically stress-test and evolve classical asset allocation strategies by relaxing the idealized assumptions inherent in Modern Portfolio Theory (MPT). We aim to quantify the sensitivity of optimal portfolio weights to real-world constraints, specifically focusing on execution latency, informational asymmetry, transaction costs, and non-normal return distributions.

The scope of the project encompasses four core pillars:

- **Baseline Evaluation:** Establishing a mathematical and performance benchmark using theoretical, frictionless Mean-Variance Optimization (MVO).
- **Relaxation of Assumptions:** Investigating portfolio behavior when theoretical constructs are replaced by empirical realities, including market microstructure frictions (execution latency, lagged data, transaction costs), statistical realities (heavy tails, skewness, volatility clustering), and behavioral constraints (varying utility functions).
- **Advanced Overlay Integration:** Developing a dynamic regime-switching framework utilizing Hidden Markov Models (HMM) to adjust allocation methodologies based on prevailing market volatility states.
- **Comparative Analysis:** Contrasting enhanced MVO models against modern, structurally robust alternatives such as Hierarchical Risk Parity (HRP) across diverse synthetic and empirical data environments.

1.2 The Limitations of Modern Portfolio Theory

While Mean-Variance Optimization remains foundational to quantitative finance, its practical application is frequently hindered by rigid theoretical assumptions. Classical MVO posits a frictionless market environment with normally distributed asset returns and stable correlation structures [12]. In reality, these assumptions routinely break down, particularly during periods of market distress.

The primary limitations addressed in this research include:

- **Estimation Error and Instability:** Classical MVO is highly sensitive to small changes in expected return estimates, often acting as an "error-maximization" algorithm that results in extreme portfolio concentration [13]. Consequently, complex optimization models frequently fail to outperform naive equal-weight portfolios out-of-sample [2].
- **Incomplete Risk Descriptors:** Financial returns consistently exhibit fat tails (excess kurtosis) and negative skewness [4], rendering variance alone an insufficient measure of true portfolio risk.
- **Correlation Instability:** MVO assumes a static covariance matrix. However, empirical evidence shows that asset correlations are highly regime-dependent, often converging toward 1.0 during severe market crises, thereby neutralizing intended diversification benefits precisely when they are most needed [10].
- **Frictionless Execution:** Classical models ignore the realities of market microstructure, assuming infinite liquidity and zero transaction costs. Because MVO is highly sensitive to small changes in expected return estimates, it often prescribes aggressive, high-turnover rebalancing [14]. In a realistic environment, the resulting transaction costs and slippage severely degrade net portfolio performance.

2 Data Acquisition and Universe Selection

The selection of the investment universe is one of the most important steps during the project. To adequately test and observe the limitations of the classical models, the data has to possess specific empirical properties:

- **Realistic, Time-Varying Correlation Structure**: Assets must exhibit non-stable correlations that evolve across market regimes. Correlations should be low or moderate in calm regimes but increase sharply during stress periods. The universe must include major crisis windows (e.g., 2008, 2020, 2022) to observe correlation breakdowns and tail dependence. This is essential for testing correlation convergence under stress and evaluating the robustness of MVO versus regime-aware and clustering-based methods.

- **Mixed Liquidity and Turnover Characteristics**: To meaningfully introduce transaction costs, assets must differ in liquidity and trading frictions. Including a mix of highly liquid assets and less liquid or spread-heavy assets ensures that turnover penalties materially affect portfolio performance, allowing HRP's lower turnover and structural stability to be evaluated fairly.

- **Statistical Reality (Non-normal returns)**: Assets should exhibit fat tails, skewness (especially negative skew), and volatility clustering. This ensures that variance alone is an incomplete risk descriptor and tests how MVO behaves when its statistical assumptions fail.

2.1 The Time Frame

To investigate the effect of correlation breakdowns and recessions, we examine three, partly overlapping time horizons: 2008-2013, 2020-2025, and 2008-2025. Procuring historical data across these specific windows ensures the inclusion of major market crisis periods necessary to observe tail dependence and rigorously stress-test the models.

2.2 The Investment Universe

To satisfy the rigorous requirements of mixed liquidity, statistical realism, and dynamic correlation, we departed from broad index ETFs and constructed a concentrated universe of 25 single-name equities and ADRs. This granular approach allows for a direct observation of idiosyncratic volatility, liquidity frictions, and structural breakdowns across eight distinct archetypes:

- **The Core Anchors (AAPL, MSFT, JPM):** These mega-cap equities serve as the baseline risk asset and anchor the portfolio. Because they are extremely liquid, they help isolate failures of the mathematical model rather than execution mechanics. Microsoft (MSFT) and Apple (AAPL) provide multi-decade stability, while JPMorgan Chase (JPM) acts as a highly liquid financial anchor that offers a direct lens into the systemic credit freezing observed during the 2008 crisis window.

- **The Growth Engines (AMZN, NVDA, NFLX, GOOGL):** These stocks exhibit momentum-driven behavior and are prone to sharp drawdowns when regimes shift. Including Amazon (AMZN), Nvidia (NVDA), Netflix (NFLX), and Alphabet (GOOGL) ensures the universe features severe fat tails, skewness (especially negative skew), and profound volatility clustering. Their rapid regime shifts perfectly test whether classical MVO is particularly sensitive to lagged data and execution delays when growth leadership abruptly changes.

- **The Value Proxies (BRK.B, WMT, BAC):** Value equities experience long periods of underperformance followed by abrupt reversals, creating highly regime-dependent expected returns. Defensive staples like Walmart (WMT) and core value like Berkshire Hathaway (BRK.B) contrast with deep-value proxies like Bank of America (BAC), whose extreme distress in 2008 causes classical MVO to rebalance aggressively as its expected returns swing wildly.

- **The Defensive Bond Proxies (NEE, PG, JNJ):** Designed to simulate the safety and yield of fixed income, these equities (NextEra Energy, Procter & Gamble, Johnson & Johnson) typically exhibit low or moderate correlations to risk assets in calm regimes. However, their true defensive efficacy—and resistance to negative skew—will be severely tested during market panics.

- **The Inflation Proxies (XOM, FCX, NEM):** These real-asset and commodity-linked equities satisfy the critical requirement that assets must exhibit non-stable correlations that evolve across market regimes. ExxonMobil (XOM), Freeport-McMoRan (FCX), and Newmont (NEM) act as direct proxies for energy shocks, industrial cycles, and physical gold, testing if the model can successfully identify safe havens during inflationary fear.

- **The Credit & Tail Risk Proxies (SPG, O):** Highly leveraged Real Estate Investment Trusts like Simon Property Group (SPG) and Realty Income (O) simulate credit market dynamics. SPG's massive drawdown in 2008 provides a vital major crisis window to observe correlation breakdowns and tail dependence.

- **The Small/Mid-Cap Volatility Block (PLUG, RIG, URBN, MSTR):** To meaningfully introduce transaction costs, the universe must include assets that differ in liquidity and trading frictions. These smaller, highly volatile companies (Plug Power, Transocean, Urban Outfitters, MicroStrategy) represent the less liquid or spread-heavy assets required to ensure that turnover penalties materially affect portfolio performance during Phase 2 testing.

- **Developed International ADRs (ASML, TM, NVO):** International equities such as ASML, Toyota (TM), and Novo Nordisk (NVO) appear diversified relative to US assets in normal times but show strong correlation convergence during global crises. Their inclusion rigorously tests the regime-switching model's ability to handle global tail dependence.

Furthermore, restricting the investment universe to exactly 25 assets serves a critical statistical purpose regarding the estimation of the covariance matrix (Σ). In classical

Mean-Variance Optimization, the sample covariance matrix must be inverted. To ensure this matrix is well-conditioned and non-singular, the number of historical observations (N) must significantly exceed the number of assets (p). By utilizing a one-year rolling lookback window of daily returns, we capture approximately 252 trading days ($N \approx 252$). With $p = 25$ assets, the observation-to-asset ratio is roughly 10:1. This ensures that the base empirical covariance matrix is mathematically stable and invertible even before applying advanced dimensionality reduction or shrinkage techniques, thereby isolating the structural flaws of MVO from pure small-sample estimation noise.

2.3 Historical Data Procurement

To conduct our empirical analysis, we utilize daily historical price data for the selected 25-asset universe. The dataset is programmatically sourced from Yahoo Finance utilizing the `yfinance` Python library.

Crucially, rather than raw closing prices, we strictly employ Adjusted Close prices to compute our daily returns. This adjustment is imperative in quantitative portfolio optimization, as it mathematically accounts for corporate actions such as stock splits and dividend distributions, ensuring an accurate reflection of total return over our extended evaluation windows (2008–2025). From these adjusted price series, we calculate the daily logarithmic returns, which serve as the foundational inputs for our covariance matrix estimation (Σ), expected return vectors (μ), and the synthetic data generation pipelines.

Beyond price series, several auxiliary datasets are procured to facilitate advanced modeling:

- **Market Capitalization and Equilibrium Weights:** To implement the Market Implied Reverse MVO, we retrieve historical shares outstanding in conjunction with daily closing prices. By reconstructing the time-series of market capitalization for each of the 25 assets, we derive the dynamic equilibrium weights (w_{mkt}) necessary to back out the implied excess returns (Π).
- **Risk-Free Rate and Academic Factors:** We incorporate the Fama-French Three-Factor (Market, Size, and Value) and Carhart Four-Factor (Momentum) data. Crucially, the daily risk-free rate (R_f) is also sourced directly from the Kenneth R. French Data Library. Utilizing a unified source for both factor premia and R_f ensures mathematical consistency when computing excess returns and performing factor regressions.
- **Market Benchmarks:** Daily data for the S&P 500 Index (GSPC) is retrieved via `yfinance` to calculate the market risk aversion coefficient (λ) and provide a baseline for performance attribution and the "Theoretical Ceiling" analysis.
- **Volatility and Regime Indicators:** Daily levels of the CBOE Volatility Index (VIX) the 10-2 year yield spread is collected. These serve as exogenous inputs for the Gaussian Hidden Markov Model (HMM) to assist in the identification of high-volatility regimes and structural market breaks.

2.4 Synthetic Data Generation: The 4 Data Environments

While historical data provide a vital empirical foundation, they represent only a single realized path of market history. To evaluate strategy robustness beyond a single historical trajectory, we utilize Monte Carlo simulations to generate synthetic return vectors $\mathbf{r}_t \in \mathbb{R}^p$ across four levels of statistical difficulty. This approach allows for the systematic isolation of specific statistical anomalies that lead to the failure of classical Mean-Variance Optimization.

For all environments, we simulate the price path $P_{i,t}$ for asset i at time t by accumulating daily logarithmic returns $r_{i,t}$:

$$P_{i,t} = P_{i,0} \exp\left(\sum_{s=1}^{t} r_{i,s}\right) \tag{1}$$

2.4.1 Environment 1: Normal Distribution (Static)

This serves as the idealized theoretical baseline. Returns are sampled from a multivariate normal distribution with constant parameters:

$$\mathbf{r}_t \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{2}$$

where $\boldsymbol{\mu}$ is the vector of expected returns and $\boldsymbol{\Sigma}$ is the static covariance matrix. Since this is a static simulation both $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are calculated by calculating the sample mean and covariance matrix for the whole timehorizon. This represents the *Theoretical Ceiling* where MVO is mathematically optimal.

2.4.2 Environment 2: Student-t Distribution (Static)

To introduce heavy tails (kurtosis $\kappa > 3$), we utilize the multivariate Student-t distribution with ν degrees of freedom:

$$\mathbf{r}_t \sim t_\nu(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{3}$$

The joint density function is defined as:

$$f(\mathbf{r}) = \frac{\Gamma[(\nu + p)/2]}{\Gamma(\nu/2)\nu^{p/2}\pi^{p/2}|\boldsymbol{\Sigma}|^{1/2}} \left[1 + \frac{1}{\nu}(\mathbf{r} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{r} - \boldsymbol{\mu})\right]^{-(\nu+p)/2} \tag{4}$$

Where ν, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are fitted to the historical data.

2.4.3 Environment 3: Jones and Faddy Skew-t Distribution (Static)

We introduce asymmetry using the Jones and Faddy skew-t distribution. This formulation is particularly useful for modeling financial returns as it allows for independent control over the left and right tails through two shape parameters, a and b. The probability density function (PDF) for the univariate case is given by:

$$f(z; a, b) = \frac{2^{1-a-b}}{\mathrm{B}(a, b)\sqrt{a + b}} \left[1 + \frac{z}{\sqrt{a + b + z^2}}\right]^{a+1/2} \left[1 - \frac{z}{\sqrt{a + b + z^2}}\right]^{b+1/2} \tag{5}$$

where $\mathrm{B}(a, b)$ is the Beta function. In our multivariate implementation, we fit the parameters for each return series to have the proper marginal structure. After that we generate a correlated multivariate normal distribution using the historical covariance matrix and use an inverse CDF transform to arrive the appropriate Jones and Faddy skew-t distribution.

2.4.4 Environment 4: Skewed-t Distribution with Dynamic Correlation

This environment represents the most sophisticated level of modeling by accounting for time-varying volatility (heteroscedasticity), dynamic correlations, and asymmetric heavy tails. The estimation and simulation follow a multi-stage process.

The model decomposes the conditional covariance matrix $\mathbf{H}_t$ into a dynamic correlation matrix $\mathbf{R}_t$ and a diagonal matrix of conditional standard deviations $\mathbf{D}_t$:

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$$

For each asset i, we first fit a GARCH(1,1) process to capture the clustering of volatility. The return is modeled as $r_{i,t} = \mu_i + \epsilon_{i,t}$, where the conditional variance $\sigma^2_{i,t}$ evolves according to:

$$\sigma^2_{i,t} = \omega_i + \alpha_i \epsilon^2_{i,t-1} + \beta_i \sigma^2_{i,t-1} \tag{6}$$

The standardized residuals $z_{i,t} = \epsilon_{i,t}/\sigma_{i,t}$ are assumed to follow the Jones and Faddy skew-t distribution defined by shape parameters a_i and b_i fitted to the data.

Using the standardized residuals $\mathbf{z}_t = [z_{1,t}, \ldots, z_{n,t}]^\top$, the DCC model captures how the correlation between assets changes over time. We define a proxy correlation matrix $\mathbf{Q}_t$ that follows an ARMA-like process:

$$\mathbf{Q}_t = (1 - \theta_1 - \theta_2)\bar{\mathbf{Q}} + \theta_1(\mathbf{z}_{t-1}\mathbf{z}^\top_{t-1}) + \theta_2 \mathbf{Q}_{t-1} \tag{7}$$

where $\bar{\mathbf{Q}}$ is the unconditional covariance matrix of the standardized residuals, and θ_1, θ_2 are scalar parameters satisfying $\theta_1 + \theta_2 < 1$. To ensure the resulting matrix is a valid correlation matrix with unit diagonals, we rescale $\mathbf{Q}_t$:

$$\mathbf{R}_t = \mathrm{diag}(\mathbf{Q}_t)^{-1/2} \mathbf{Q}_t \mathrm{diag}(\mathbf{Q}_t)^{-1/2} \tag{8}$$

Based on this time varying correlation matrix we use a similar logic as in the previous style generating correlated normal distributions and using the inverse CDF transform to arrive at the properly correlated Jones and Faddy skew-t distribution.

2.5 Portfolio Constraints

In practical asset management, an unconstrained optimization often leads to highly concentrated, *fragile* portfolios that are susceptible to estimation error. To ensure diversification and adhere to institutional risk-management standards, we define a set of constraints that bound the feasible set of weights $\mathbf{w}$.

The optimization problem is solved subject to the budget constraint $\sum_{i=1}^{n} w_i = 1$ and the following specific limitations:

2.5.1 Individual Asset Bounds (Box Constraints)

We define the feasible range for each weight w_i such that $0 \leq w_i \leq \omega_i^{max}$. These bounds are tiered based on the liquidity and risk profile of the asset class:

Asset Category	$\omega^{\mathbf{max}}$	**Constituents**
Mega-Cap Anchors	0.15	AAPL, MSFT, GOOGL, AMZN, JPM, BRK.B, JNJ, PG, WMT
Growth & International	0.08	ASML, BAC, NEE, NFLX, NVDA, NVO, TM, XOM
High-Volatility/Small-Cap	0.05	MSTR, PLUG, RIG, URBN
Speculative/Specialty	0.03	FCX, NEM, O, SPG

Table 1: Individual Asset Weight Caps

2.5.2 Sector and Thematic Constraints (Linear Constraints)

To prevent thematic over-concentration, which can lead to systemic drawdowns during sector-specific shocks (e.g., "Tech Bubble" or "Commodity Crash"), we impose group-level linear constraints:

1. **Technology and Growth Ceiling:** To mitigate exposure to duration risk and high-beta tech volatility, the aggregate weight is capped at 40%:

$$\sum_{i \in \text{Tech}} w_i \leq 0.40 \tag{9}$$

$$\text{Tech} = [\text{AAPL, MSFT, GOOGL, AMZN, NVDA, NFLX, ASML}] \tag{10}$$

2. **Defensive Floor and Ceiling:** To ensure the portfolio maintains "Flight-to-Safety" characteristics , we enforce a defensive allocation between 5% and 50%:

$$0.05 \leq \sum_{i \in \text{Defensive}} w_i \leq 0.50 \tag{11}$$

$$\text{Defensive} = [\text{JNJ, PG, WMT, NEE}] \tag{12}$$

3. **Commodity and Hard Asset Ceiling:** To limit exposure to the cyclicality of the materials and energy sectors, the aggregate weight is capped at 20%:

$$\sum_{i \in \text{Commodity}} w_i \leq 0.20 \tag{13}$$

$$\text{Commodity} = [\text{FCX, NEM, XOM}] \tag{14}$$

These constraints define the convex feasible region $\mathcal{W}$ within which the models must operate. By restricting the solution space, we force the models to find "robust" diversification rather than mathematically convenient but practically dangerous concentrations.

3 Strategy Models and Theoretical Framework

Having formalized the empirical constraints and the non-stationary stochastic environments of testing data in Section 2, we now define the allocation models tasked with navigating these stress tests. This section establishes the theoretical foundations, objective functions, and optimization mechanics for the spectrum of portfolio allocation strategies evaluated in this study.

To systematically isolate and measure the value of algorithmic complexity, we structure our models into a definitive hierarchy. We begin by formalizing heuristic baselines—specifically the Equal Weight ($1/N$) and Inverse Volatility Portfolios (IVP)—which bypass complex covariance estimation entirely. Next, we rigorously define the classical Mean-Variance Optimization (MVO) framework, establishing the traditional benchmark that our synthetic environments are explicitly designed to break. Finally, we introduce the advanced allocation architectures: Hierarchical Risk Parity (HRP) and the Regime-Switching Hidden Markov Model (rHMM-MVO). These modern frameworks utilize graph theory and probabilistic state-detection, respectively, in an attempt to overcome the estimation errors and non-stationarity vulnerabilities inherent in classical MVO.

3.1 Baseline Benchmarks: S&P 500, Equal Weight, and Inverse Volatility Portfolio (IVP)

Before evaluating optimization-driven frameworks, we define three baseline strategies. These models serve as the benchmark to determine whether the computational complexity of covariance estimation and regime detection translates into actual outperformance, particularly net of transaction costs.

3.1.1 The Market Beta Benchmark: SPY Buy and Hold

To measure the absolute value proposition of active allocation, we benchmark all strategies against the passive market portfolio, proxied by the SPDR S&P 500 ETF Trust (SPY). This represents a capitalization-weighted allocation to US equities. Unlike the dynamic models, this baseline involves zero turnover (ignoring dividend reinvestment friction) and zero parameter estimation. The daily logarithmic return is defined simply as:

$$R_{\text{SPY},t} = \ln\left(\frac{P_{\text{SPY},t}}{P_{\text{SPY},t-1}}\right)$$

This serves as the definitive baseline for pure market beta exposure, against which the risk-adjusted performance and maximum drawdowns of our active quantitative models will be judged.

3.1.2 The Naive Benchmark: Equal Weight ($1/N$)

The Equal Weight portfolio serves as the estimation-free baseline. As demonstrated by DeMiguel et al[2]), the $1/N$ rule often outperforms sophisticated optimization models out-of-sample due to the complete absence of parameter estimation error. For an investment universe of n assets, the weight allocated to each asset i at any time t is strictly defined as:

$$w_{i,t} = \frac{1}{N}$$

This strategy ignores all expected returns ($\boldsymbol{\mu}$) and risk parameters ($\boldsymbol{\Sigma}$). In the context of our Execution Gauntlet, it suffers no estimation lag but incurs continuous transaction costs to reverse market drift and maintain the $1/N$ target.

3.1.3 The Risk Heuristic: Inverse Volatility Portfolio (IVP)

The Inverse Volatility Portfolio introduces a naive form of risk management by penalizing highly volatile assets, yet it purposely avoids estimating the full covariance matrix $\boldsymbol{\Sigma}$. IVP effectively assumes that all cross-asset correlations are zero (or equal), isolating the marginal risk of each asset. The weight for asset i at time t is inversely proportional to its conditional standard deviation $\sigma_{i,t}$:

$$w_{i,t} = \frac{\sigma_{i,t}^{-1}}{\sum_{j=1}^{n} \sigma_{j,t}^{-1}}$$

By solely estimating the diagonal of the covariance matrix (the volatilities), IVP significantly reduces the dimensionality of the estimation problem. This makes it highly robust to the correlation breakdowns simulated in Environments 4, though it remains blind to actual tail dependence.

3.2 Classical Mean-Variance Optimization (MVO)

The foundation of modern portfolio theory, introduced by Markowitz [12], seeks to identify the optimal vector of portfolio weights $\mathbf{w}$ that maximizes expected return for a given level of variance, or equivalently, maximizes the risk-adjusted return (Sharpe Ratio).

Let $\boldsymbol{\mu}$ be the $n \times 1$ vector of expected asset returns, $\boldsymbol{\Sigma}$ be the $n \times n$ covariance matrix of returns, and R_f be the risk-free rate. The objective of the Maximum Sharpe Ratio (Tangency) portfolio is to solve:

$$\max_{\mathbf{w}} \frac{\mathbf{w}^T \boldsymbol{\mu} - R_f}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} \tag{15}$$

Subject to the budget constraint $\sum_{i=1}^{n} w_i = 1$ and the regulatory feasible space $\mathbf{w} \in \mathcal{W}$ defined in Section 2.

While mathematically elegant, classical MVO is notoriously sensitive to its inputs. As noted by Michaud [13], MVO algorithms act as *error maximizers,* overweighting assets with positive estimation errors in $\boldsymbol{\mu}$ and negative estimation errors in $\boldsymbol{\Sigma}$. Consequently, relying strictly on the sample mean and sample covariance inevitably leads to fragile out-of-sample performance, particularly in the non-stationary frameworks of Environments 4.

3.2.1 Extensions: Advanced Parameter Estimation

To mitigate the severe impact of estimation error and delay the degradation of alpha, practitioners rarely use raw historical inputs. We formalize several sophisticated estimation techniques that attempt to stabilize the MVO framework:

3.2.1.1 Refining Expected Returns ($\boldsymbol{\mu}$)

Because the MVO objective function is extraordinarily sensitive to errors in expected returns, the historical sample mean can be replaced by structurally informed or autoregressive forecasts:

1. **Equilibrium Returns (Market-Implied Reverse MVO):** Instead of forecasting returns from past data, this approach extracts the expected returns implied by current market capitalization weights ($\mathbf{w}_{mkt}$). Assuming the market portfolio is mean-variance efficient, the implied equilibrium excess returns $\mathbf{\Pi}$ are calculated as:

$$\mathbf{\Pi} = \delta \mathbf{\Sigma} \mathbf{w}_{mkt}$$

   where δ is the market risk aversion coefficient. This anchors the optimizer and prevents extreme corner solutions.

2. **Fama–French Three-Factor Model:** This structural model[6] posits that expected returns are driven by exposure to systematic risk factors. The expected return for asset i is forecasted via its sensitivity (β) to the market portfolio (MKT), the size factor (SMB), and the value factor (HML):

$$E[R_i] = R_f + \beta_{i,MKT}(E[R_M] - R_f) + \beta_{i,SMB}\text{SMB} + \beta_{i,HML}\text{HML}$$

3. **Carhart Four-Factor Model:** Extending the Fama-French framework, Carhart[3] introduces a Momentum factor (MOM) to capture the empirical tendency of winning assets to continue outperforming in the short term:

$$E[R_i] = R_f + \boldsymbol{\beta}_i^T \mathbf{F}_{\text{Carhart}}$$

3.2.1.2 Refining the Covariance Forecast ($\mathbf{\Sigma}$)

To address the whipsaw effect caused by shifting correlation regimes, the sample covariance matrix $\mathbf{S}$ can be replaced by estimators that handle high dimensionality and time-varying risk:

1. **Ledoit-Wolf Shrinkage:** This method[8] mitigates the ill-conditioning of the sample covariance matrix by shrinking $\mathbf{S}$ towards a highly structured target matrix $\mathbf{F}$ (e.g., constant correlation). The optimal shrinkage intensity δ minimizes the out-of-sample estimation error:

$$\mathbf{\Sigma}_{\text{LW}} = \delta \mathbf{F} + (1 - \delta)\mathbf{S}$$

2. **Factor Models:** By mapping asset returns to a smaller set of K underlying structural factors, the dimensionality of the covariance estimation is drastically reduced. The covariance is reconstructed as:

$$\mathbf{\Sigma}_{\text{Factor}} = \mathbf{B}\mathbf{\Sigma}_F\mathbf{B}^T + \mathbf{\Delta}$$

where $\mathbf{B}$ is the matrix of factor loadings, $\mathbf{\Sigma}_F$ is the factor covariance matrix, and $\mathbf{\Delta}$ is the diagonal matrix of idiosyncratic variances.

3. **Exponentially Weighted Moving Average (EWMA):** EWMA applies a decay factor $\lambda \in (0, 1)$ to assign greater weight to recent observations, allowing the covariance matrix to react swiftly to volatility shocks:

$$\mathbf{\Sigma}_t = \lambda\mathbf{\Sigma}_{t-1} + (1-\lambda)\mathbf{r}_{t-1}\mathbf{r}_{t-1}^T$$

3.3 Hierarchical Risk Parity (HRP)

Classical MVO and its risk-parity derivatives suffer from a critical mathematical vulnerability: they require the inversion of the covariance matrix ($\mathbf{\Sigma}^{-1}$). In non-stationary environments where assets become highly correlated (such as the market crashes simulated in Environment 5), the covariance matrix becomes ill-conditioned. This causes the optimizer to compute extreme, unstable weights due to amplified estimation errors.

To bypass matrix inversion entirely, Lopez de Prado [11] introduced Hierarchical Risk Parity (HRP). HRP applies graph theory and machine learning to allocate capital based on the topological structure of the market. It is important to note that while HRP is inherently more robust to noise than MVO, its input covariance matrix $\mathbf{\Sigma}$ and correlation matrix $\boldsymbol{\rho}$ do not have to be derived from naive historical samples. Similar methods detailed in Section 3.2.1.2, such as Ledoit-Wolf shrinkage, or EWMA can be used as the foundational inputs for the HRP algorithm, further enhancing its structural resilience. The HRP algorithm operates in three distinct stages:

3.3.1 Step 1: Tree Clustering (Hierarchical Dependency)

Instead of treating all assets as a single interconnected web, HRP groups them into a hierarchical tree (dendrogram) based on their correlation distance. First, we transform the correlation matrix $\boldsymbol{\rho}$ into a distance matrix $\mathbf{D}$, where the distance between asset i and asset j is defined as:

$$D_{i,j} = \sqrt{\frac{1}{2}(1-\rho_{i,j})}$$

Because $D_{i,j}$ lacks the properties of a true metric space, we calculate the Euclidean distance between the column vectors of $\mathbf{D}$ to generate a true distance matrix $\tilde{\mathbf{D}}$:

$$\tilde{D}_{i,j} = \sqrt{\sum_{n=1}^{N}(D_{n,i} - D_{n,j})^2}$$

Using $\tilde{\mathbf{D}}$, we apply an agglomerative clustering algorithm (Single Linkage) to iteratively merge the closest assets and clusters until all assets are contained within a single overarching root cluster.

3.3.2 Step 2: Quasi-Diagonalization

Once the hierarchical tree is constructed, we reorganize the rows and columns of the original covariance matrix $\mathbf{\Sigma}$. Assets that belong to the same cluster are placed adjacent to one another. This topological reorganization transforms the covariance matrix into a "quasi-diagonal" state, where the largest covariance values are localized along the diagonal, and off-diagonal blocks represent the lower covariance between distinct market sectors.

3.3.3 Step 3: Recursive Bisection (Capital Allocation)

Capital is allocated top-down along the dendrogram using a recursive inverse-variance methodology. Let V represent the variance of a cluster. For any two adjacent clusters (or assets) L (left) and R (right) within the tree, the intra-cluster weights are assigned inversely proportional to their individual variances to find the cluster variance:

$$V_c = \mathbf{w}_c^T \mathbf{\Sigma}_c \mathbf{w}_c$$

where $\mathbf{\Sigma}_c$ is the covariance matrix of the assets within cluster c, and $\mathbf{w}_c$ is the inverse-diagonal variance vector for those assets.

The algorithm then splits the capital allocation α between the left and right branches based on their cluster variances:

$$\alpha_L = 1 - \frac{V_L}{V_L + V_R}, \quad \alpha_R = 1 - \alpha_L$$

This process begins at the root of the tree with 100% of the capital and recursively bisects the weights down the branches until it reaches the individual asset leaves. By only considering the covariance between structurally adjacent assets, HRP completely eliminates the need to invert the global covariance matrix, directly neutralizing the *error maximizing* nature of MVO.

3.4 Regime-Switching Hidden Markov Model MVO (rHMM-MVO)

While advanced forecasting methods (e.g., DCC-GARCH) and topological allocations (HRP) adapt well to gradually shifting volatility and correlation, they fundamentally struggle with abrupt structural breaks—such as sudden market crashes or instantaneous liquidity dry-ups. The Regime-Switching Hidden Markov Model (rHMM) addresses this by abandoning the single-state assumption, positing instead that financial markets operate across discrete, unobservable regimes (e.g., "Bull," "Bear," or "Crash"), each possessing its own distinct statistical properties.

By coupling a robust HMM with the MVO framework, the rHMM-MVO algorithm explicitly detects these structural breaks in real-time, allowing the portfolio to dynamically toggle its objective function, constraints, or risk parameters based on the prevailing state of the market.

3.4.1 Step 1: The Hidden Markov Architecture

Let $S_t \in \{1, 2, \ldots, K\}$ represent the discrete, unobservable market state at time t. The sequence of states follows a first-order Markov process, governed by a $K \times K$ transition probability matrix $\mathbf{P}$, where each element $p_{i,j}$ defines the probability of transitioning from state i to state j:

$$p_{i,j} = \mathbb{P}(S_t = j \mid S_{t-1} = i) \tag{16}$$

Because the states S_t are hidden, we must infer them through a vector of observable market data $\mathbf{O}_t$. In this framework, $\mathbf{O}_t$ comprises not only the continuous asset returns $\mathbf{r}_t$, but is augmented with exogenous macro-financial indicators, including the risk-free rate $r_{f,t}$, rolling realized volatilities, and the implied volatility index (VIX).

While classical HMMs assume Gaussian emission probabilities for mathematical tractability, doing so in the presence of empirical market realities (and the non-stationary, heavy-tailed dynamics of our Data Laboratory) results in severe state misclassification. A Gaussian model interprets standard fat-tailed shocks as structural breaks, leading to excessive regime switching and algorithmic whipsaw. Therefore, we model the state-dependent emissions using a robust Multivariate Student-t distribution:

$$\mathbf{O}_t \mid S_t = k \sim t_{\nu_k}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \tag{17}$$

where $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ are the expected vector and shape matrix of the observable features specific to regime k, and ν_k represents the state-specific degrees of freedom, allowing the model to absorb heavy-tailed market shocks without incorrectly triggering a state transition.

3.4.2 Step 2: Robust State Inference and Forecasting

To estimate the model parameters ($\mathbb{P}$, $\boldsymbol{\mu}_k$, $\boldsymbol{\Sigma}_k$), we apply the Expectation-Maximization (EM) algorithm, over a rolling historical window. Because isolated market regimes (e.g., crashes) may contain very few observations, the sample covariance matrices $\boldsymbol{\Sigma}_k$ risk becoming ill-conditioned. If it so happens we will apply *robust* regularization—such as Ledoit-Wolf shrinkage—to the state-conditional covariance matrices during the Maximization step to ensure invertibility.

At each time step t, the algorithm utilizes the Forward-Backward procedure to calculate the posterior probability (or "belief") of being in each state k, denoted as $\gamma_{k,t} = P(S_t = k \mid \mathbf{O}_1, \ldots, \mathbf{O}_t)$.

To allocate capital for time $t + 1$, we forecast the state probability distribution $\boldsymbol{\pi}_{t+1}$ by multiplying the current state beliefs $\boldsymbol{\gamma}_t$ by the transition matrix $\mathbf{P}$:

$$\boldsymbol{\pi}_{t+1} = \boldsymbol{\gamma}_t \mathbf{P}$$

3.4.3 Step 3: State-Conditional Optimization

Once the forward-looking regime probabilities $\boldsymbol{\pi}_{t+1}$ are calculated, the rHMM-MVO model synthesizes a blended, regime-adjusted expected return vector $\boldsymbol{\mu}^*$ and covariance matrix $\boldsymbol{\Sigma}^*$:

$$\boldsymbol{\mu}^*_{t+1} = \sum_{k=1}^{K} \pi_{k,t+1} \boldsymbol{\mu}_k$$

$$\boldsymbol{\Sigma}^*_{t+1} = \sum_{k=1}^{K} \pi_{k,t+1} \left(\boldsymbol{\Sigma}_k + (\boldsymbol{\mu}_k - \boldsymbol{\mu}^*_{t+1})(\boldsymbol{\mu}_k - \boldsymbol{\mu}^*_{t+1})^T \right)$$

These dynamically blended, state-aware parameters are then fed directly into the classical Maximum Sharpe MVO objective function:

$$\max_{\mathbf{w}} \frac{\mathbf{w}^T \boldsymbol{\mu}^*_{t+1} - R_f}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma}^*_{t+1} \mathbf{w}}}$$

Subject to the budget and regulatory constraints $\mathbf{w} \in \mathcal{W}$.

Constraint Switching (The Tactical Overlay) Beyond blending parameters, the true power of the rHMM-MVO lies in its ability to enact hard constraint switching. If the probability of a "Crash" state (e.g., state K) exceeds a critical threshold τ (i.e., $\pi_{K,t+1} > \tau$), the model programmatically overwrites the linear bounds defined in Section 2, forcing the portfolio to maximize its allocation to the "Defensive Floor" constituents. This equips the rHMM-MVO to survive the dynamic t-Copula tail dependencies modeled in Environment 5 without suffering the whipsaw lag inherent to purely reactive moving-average estimators.

3.5 Portfolio Evaluation Metrics

To rigorously assess the performance of the baseline, MVO, HRP, and rHMM-MVO frameworks across the simulated environments, we must employ an evaluation criteria that accounts for higher-order moments (skewness and kurtosis) and the practical realities of trading friction. Relying solely on raw returns or Gaussian-based metrics would inherently obscure the specific vulnerabilities our Data Laboratory was designed to expose.

3.5.1 Return and Volatility Profile

We begin by establishing the foundational risk and return characteristics of the generated portfolios:

1. **Annualized Return (R_{ann}):** The geometric average of the portfolio's daily returns, scaled to a 252-day trading year, providing the baseline for absolute performance.

2. **Annualized Volatility (σ_{ann}):** The annualized standard deviation of daily portfolio returns. While a standard measure of dispersion, it treats upside gains and downside losses identically.

3. **Sharpe Ratio (SR):** Evaluates the return generated per unit of total risk. Mathematically represented as $SR = \frac{R_{ann} - R_f}{\sigma_{ann}}$. While an industry standard, it penalizes upside volatility and fundamentally assumes normally distributed returns—a flaw our heavy-tailed environments will explicitly expose.

4. **Sortino Ratio (SoR):** A vital improvement over the Sharpe Ratio for this study, as it only penalizes downside volatility (σ_d). By focusing on the semi-variance below 0, it is highly effective when evaluating assets with the negative skewness injected into our simulations.

3.5.2 Tail Risk and Capital Preservation

Since our data-generating processes introduce statistical realism like heavy tails and correlation breakdowns, we must explicitly measure what happens at the extremes.

1. **Maximum Drawdown (MDD):** Measures the absolute worst-case peak-to-trough decline in portfolio value.

2. **Calmar Ratio:** Defined as the ratio of Annualized Return to Maximum Drawdown ($\frac{R_{ann}}{|MDD|}$). This metric heavily penalizes strategies that generate high returns through exposure to catastrophic, regime-shifting risks.

3. **Value at Risk (VaR):** Quantifies the maximum expected loss over a specific timeframe at a given confidence interval α (e.g., 95% or 99%).

4. **Expected Shortfall (ES / CVaR):** While VaR defines the threshold of a bad day, ES evaluates the average expected loss *conditional* on that VaR threshold being breached. It directly quantifies the *fat tail* risks inherent in financial data.

3.5.3 Structural and Distributional Metrics

Because HRP is designed to address correlation instability and classical MVO is notoriously prone to extreme concentration, we must measure the internal topography of the allocations.

1. **Herfindahl-Hirschman Index (HHI):** Measures the concentration of the portfolio weights. Calculated as the sum of the squared weights ($HHI = \sum w_i^2$). The inverse of this index can be interpreted as the effective number of allocations made by the portfolio.

2. **Portfolio Skewness (S_p) and Kurtosis (K_p):** Because the optimization engines are fed non-normal data with fat tails and negative skew, measuring the third and fourth moments of the resulting portfolio returns will demonstrate whether the algorithms successfully smoothed out these statistical anomalies or merely amplified them.

3.5.4 Execution Friction Metrics

To quantify the structural drag of each strategy without the overhead of tracking hypothetical gross returns, we measure the trading volume and its direct financial cost at the transaction level.

1. **Portfolio Turnover:** Measures the proportion of the portfolio traded during each rebalance. Let $w_{i,t}$ be the target weight and w_{i,t^-} be the drifted weight prior to execution. Turnover is defined as:

$$\text{Turnover}_t = \sum_{i=1}^{n} |w_{i,t} - w_{i,t^-}|$$

   This is the primary indicator of algorithmic stability; high turnover reveals a model whipsawing in response to market noise.

2. **Realized Execution Cost:** Captures the absolute friction deducted from the portfolio, based on the actual transaction price ($P_{i,t}$) and the number of shares traded ($\Delta N_{i,t}$). Applying a proportional cost factor c (for slippage and fees), the total dollar cost per rebalance is:

$$TC_t = \sum_{i=1}^{n} c \cdot P_{i,t} \cdot |\Delta N_{i,t}|$$

   This cost is immediately subtracted from the portfolio's total capital. Consequently, all reported performance metrics (Sharpe, MDD, Returns) inherently reflect the net-of-cost reality.

4 Data Adequacy Analysis

The fundamental premise of this study is that classical Mean-Variance Optimization (MVO) fails under real-world market conditions due to its reliance on strict Gaussian and stationary assumptions. To rigorously test the advanced frameworks (HRP and rHMM-MVO), the selected investment universe must demonstrably violate these classical assumptions. This section verifies the adequacy of our empirical dataset by analyzing its structural diversity, higher-order moments, volatility dynamics, and correlation stability.

4.1 Testing for Non-Normality: Skewness and Kurtosis

Classical MVO assumes that asset returns are adequately described by their first two discrete moments. To verify that our data requires a more robust approach, we test for deviations from the Multivariate Normal distribution.

	Skewness			Excess Kurtosis		
Ticker	2008-2013	2020-2025	2008-2025	2008-2013	2020-2025	2008-2025
AAPL	-0.46	0.03	-0.34	3.62	3.46	4.45
AMZN	0.63	-0.07	0.34	4.32	1.33	4.97
ASML	-0.10	-0.43	-0.24	-0.98	1.73	1.47
BAC	-0.19	-0.09	-0.32	7.29	7.58	20.93
BRK-B	0.82	-0.26	0.39	8.30	8.88	11.08
FCX	-0.42	-0.17	-0.20	1.64	3.60	3.01
GOOGL	0.38	-0.17	0.24	5.98	0.70	5.33
JNJ	0.65	0.08	-0.07	10.75	5.12	7.91
JPM	0.29	-0.07	0.24	5.07	9.52	13.79
MSFT	0.30	-0.16	0.00	5.06	4.87	5.79
MSTR	-0.86	-0.09	-0.20	5.96	0.69	7.19
NEE	0.61	-0.45	-0.08	8.21	4.24	7.88
NEM	0.46	-0.24	0.04	5.32	2.14	3.24
NFLX	-1.31	-2.51	-0.89	14.41	38.84	23.01
NVDA	-0.86	0.04	-0.34	6.42	1.07	6.48
NVO	0.08	-1.32	-0.93	1.57	14.51	11.37
O	0.45	-3.56	-0.69	7.65	58.75	27.13
PG	-0.14	-0.19	-0.13	4.12	7.82	6.55
PLUG	-0.29	0.22	0.03	17.93	2.59	15.18
RIG	-0.66	-0.15	-0.12	2.91	9.11	11.92
SPG	0.24	-1.00	-0.28	4.82	22.05	18.32
TM	-0.31	0.39	-0.04	5.46	2.28	5.80
URBN	-0.55	0.37	-0.03	4.24	2.09	4.26
WMT	0.11	-0.11	-0.03	4.77	9.60	9.97
XOM	0.15	-0.21	-0.07	10.79	1.96	7.09

Table 2: Empirical Skewness and Kurtosis across Selected Regimes (2008-2025)

Table 2 details the third and fourth moments for the 25-asset universe across these periods. The empirical results categorically reject the Gaussian assumption across two dimensions:

1. **Asymmetry (Skewness):** Classical models assume symmetrical risk, but the data reveals profound left-tail clustering. During the 2020-2025 shock, ostensibly stable real estate and pharmaceutical assets exhibited catastrophic negative skewness (e.g., Realty Income `O` at -3.56 and Novo Nordisk `NVO` at -1.32). Even high-growth market leaders like Netflix (`NFLX`) displayed extreme negative skew (-2.51), confirming the adage that empirical equities take the stairs up and the elevator down. A variance-only optimizer is mathematically blind to this directional risk.

2. **Leptokurtosis (Fat Tails):** The dataset exhibits severe leptokurtosis, meaning outlier events occur with significantly higher frequency than a normal distribution predicts. Across the full 2008-2025 sample, assets like Bank of America (`BAC`) and `NFLX` produced excess kurtosis values of 20.93 and 23.01, respectively. More critically, the 2020-2025 period forced the excess kurtosis of `O` to a staggering 58.75 and Simon Property Group (`SPG`) to 22.05.

4.2 Testing for Heteroskedasticity and Volatility Clustering

A core assumption of classical portfolio theory is homoskedasticity—the premise that the variance of asset returns remains constant over time. If this assumption holds true, a static, unconditional sample covariance matrix ($\mathbf{S}$) serves as an adequate estimator of future risk. However, empirical financial time series frequently exhibit Autoregressive Conditional Heteroskedasticity (ARCH) effects, where periods of high volatility tend to cluster together, followed by periods of relative calm.

To rigorously verify the presence of serial correlation within our selected universe, we apply the Ljung-Box test to the daily log returns of the 25 assets. The null hypothesis (H_0) posits that the returns are independently distributed (zero autocorrelation). We evaluate this across two distinct macro-regimes (2008–2013 and 2020–2025) as well as the full continuous sample. The resulting p-values are detailed in Table 3.

The empirical evidence rigorously rejects the assumption of independent returns for a significant portion of the universe. Across the continuous 2008–2025 sample, core portfolio anchors—including `MSFT`, `JPM`, `JNJ`, and `O`—yield p-values of 0.0000, decisively rejecting the null hypothesis at the 1% significance level. This confirms the presence of systemic momentum or mean-reversion within these asset paths.

More importantly, the table exposes the *regime-dependent nature* of this serial correlation. Several assets that behaved as independent random walks (white noise) during the 2008–2013 period structurally shifted into highly autocorrelated regimes during the 2020–2025 shock. For example:

- **MSTR & NVDA:** Displayed no significant autocorrelation in the first regime ($p = 0.9508$ and $p = 0.2779$, respectively), but exhibited severe serial dependency during the 2020s ($p = 0.0006$ and $p = 0.0000$).

- **FCX:** Transitioned from pure independence ($p = 0.9299$) to significant correlation ($p = 0.0009$).

Ticker	AAPL	AMZN	ASML	BAC	BRK-B
2008-2013	0.0529	0.5186	0.0278	0.0001	0.0000
2020-2025	0.0000	0.4338	0.0000	0.0000	0.0000
2008-2025	0.0007	0.2360	0.0000	0.0000	0.0000
Ticker	FCX	GOOGL	JNJ	JPM	MSFT
2008-2013	0.9299	0.0447	0.0000	0.0004	0.0000
2020-2025	0.0009	0.0000	0.0000	0.0000	0.0000
2008-2025	0.1565	0.0007	0.0000	0.0000	0.0000
Ticker	MSTR	NEE	NEM	NFLX	NVDA
2008-2013	0.9508	0.0001	0.0042	0.9953	0.2779
2020-2025	0.0006	0.0000	0.3416	0.0352	0.0000
2008-2025	0.0000	0.0000	0.1412	0.4561	0.0018
Ticker	NVO	O	PG	PLUG	RIG
2008-2013	0.0001	0.0000	0.0001	0.0402	0.1670
2020-2025	0.2831	0.0000	0.0000	0.3567	0.6311
2008-2025	0.0396	0.0000	0.0000	0.0814	0.3752
Ticker	SPG	TM	URBN	WMT	XOM
2008-2013	0.0000	0.0092	0.0038	0.0000	0.0000
2020-2025	0.0000	0.0578	0.0000	0.0009	0.0009
2008-2025	0.0000	0.0186	0.0000	0.0000	0.0003

Table 3: Ljung-Box Test p-values for Log Returns across Selected Regimes

Conversely, assets like `AMZN` ($p > 0.20$ across all periods) and `RIG` consistently fail to reject the null hypothesis, maintaining an idiosyncratic, memoryless return profile regardless of the macroeconomic environment.

Because log returns exhibit these shifting dependencies, a static historical mean vector ($\boldsymbol{\mu}$) is mathematically insufficient. An optimizer relying on unconditional averages will lag behind structural trend changes. This empirical reality provides the foundational justification for deploying a Regime-Switching Hidden Markov Model (rHMM), which explicitly models these temporal dependencies by estimating state-conditional means ($\boldsymbol{\mu}_k$) adapted to the prevailing market environment.

4.3 Correlation Instability and Regime Shifts

Finally, the dataset must adequately represent the "whipsaw" correlation dynamics that plague classical diversification. Classical models rely on a static correlation matrix, assuming that the defensive properties of assets like JNJ or NEM will hold independently of tech drawdowns.

The visual evidence categorically rejects the assumption of static correlations. The rolling correlation matrix is highly dynamic, characterized by continuous structural shifts and severe instability. Several critical empirical realities emerge from this analysis:

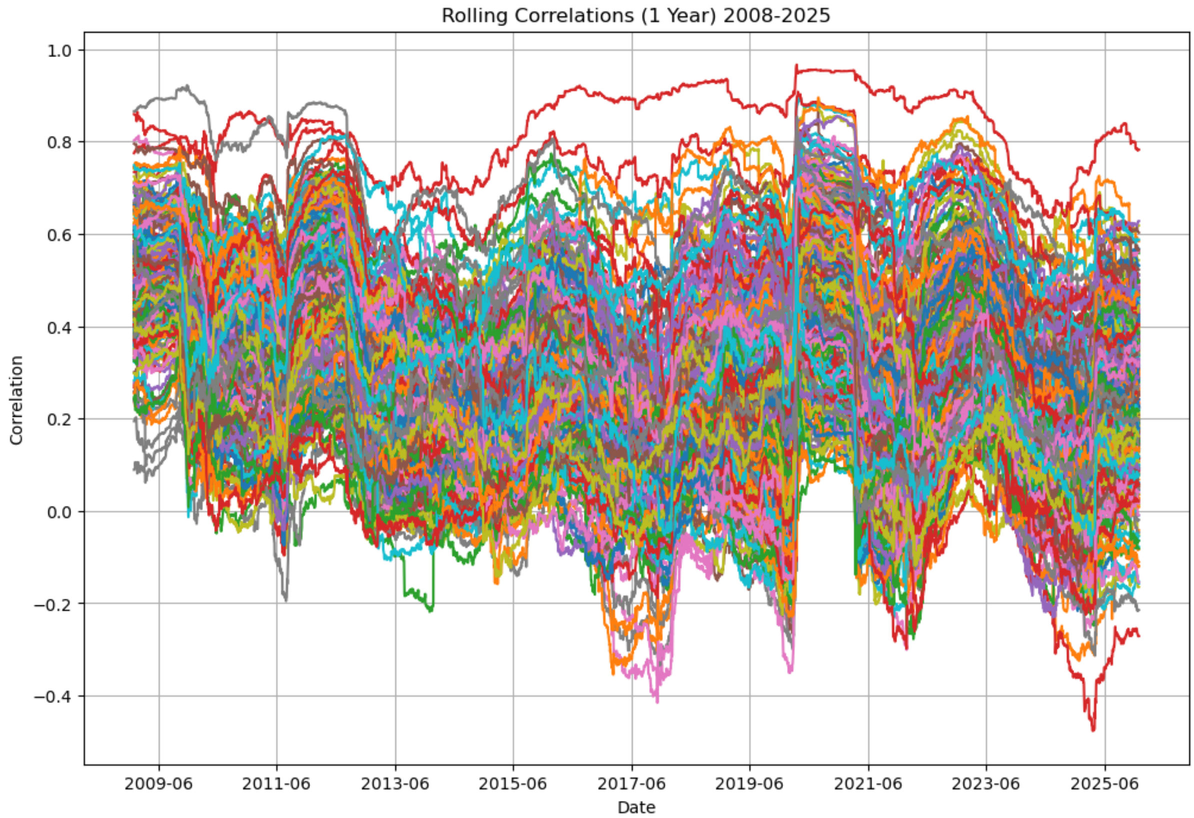

Figure 1: 1-Year Rolling Pairwise Correlations across the 25-Asset Universe (2008–2025).

1. **The Contagion Effect:** During macroeconomic shocks—most notably visible in the violent upward convergence during the March 2020 liquidity crisis—pairwise correlations across traditionally orthogonal assets spike simultaneously toward 1.0.

2. **Decoupling and Regime Shifts:** The chart demonstrates periods where correlations violently decouple, plummeting into deep negative territory (e.g., late 2017, mid-2024). A static model, relying on a long-term historical average (e.g., $R = 0.3$), completely misses the opportunity to exploit these emergent, localized hedging benefits.

5 Empirical Baseline: Performance in a Frictionless Environment

Having established the statistical non-normality and correlation instability of our 25-asset universe, we now evaluate the performance of our foundational optimization frameworks (Equal Weight, Inverse Variance Parity, Classical MVO, and HRP). To isolate the pure allocation logic of each algorithm, this initial backtest is executed in a strictly frictionless vacuum. We assume infinite liquidity, continuous fractional sizing, and crucially, zero transaction costs ($c = 0$).

We evaluate the cumulative equity trajectories (Figures 2, 4, and 4) and their corresponding risk-adjusted metrics across three distinct macro-horizons: the 2008–2013 Financial Crisis, the 2020–2025 Liquidity Shock, and the full 2008–2025 strategic horizon.

5.1 The Contagion Stress Test (2008–2013)

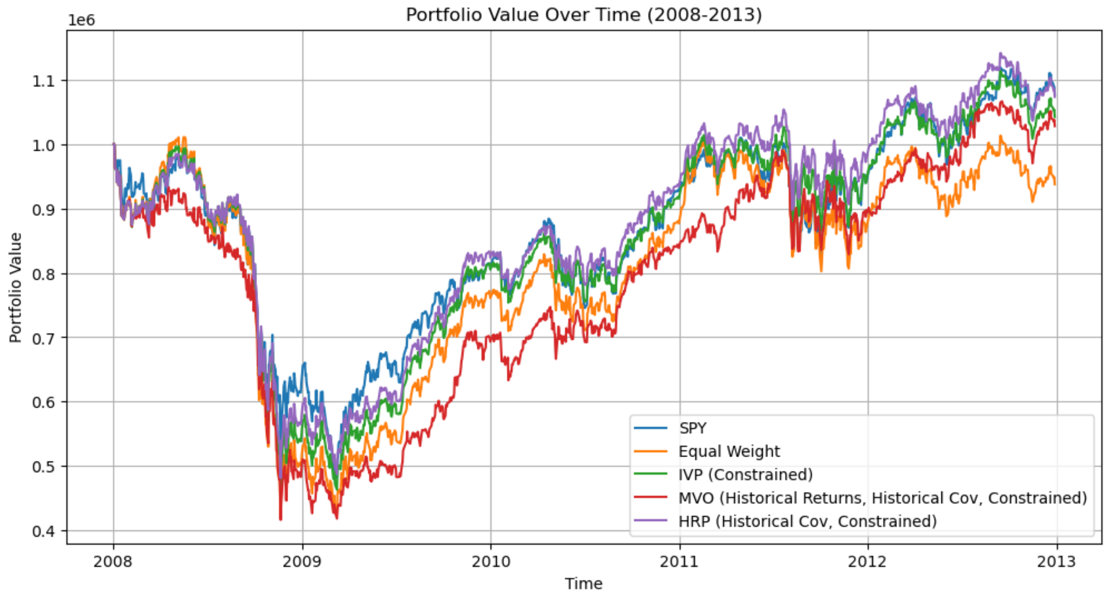

Figure 2: Portfolio Value Over Time (2008-2013) under Frictionless Assumptions.

The 2008–2013 period serves as the ultimate test of capital preservation, marked by severe correlation contagion. In this environment, naive diversification failed catastrophically. The Equal Weight (EW) portfolio suffered a Maximum Drawdown of 58.37%, severely lagging the S&P 500 benchmark (SPY).

Classical MVO, constrained by historical covariance estimations, failed to anticipate the structural break of the crisis. It yielded a negligible annualized return (0.55%) and suffered the exact same catastrophic drawdown as the EW portfolio (58.37%).

In stark contrast, the Hierarchical Risk Parity (HRP) algorithm demonstrated immediate structural superiority. By relying on topological clustering rather than precise

Model	SPY	EW	IVP	MVO	HRP
Annualized Return	0.0150	-0.0128	0.0083	0.0055	0.0142
Annualized Volatility	0.2621	0.2776	0.2488	0.2491	0.2350
Sharpe Ratio	0.0573	-0.0463	0.0333	0.0221	0.0606
Sortino Ratio	0.0730	-0.0574	0.0408	0.0269	0.0747
Max Drawdown	0.5187	0.5837	0.5400	0.5837	0.5186
Calmar Ratio	0.0290	-0.0220	0.0153	0.0094	0.0275
VAR (95%)	-0.0248	-0.0280	-0.0250	-0.0248	-0.0226
ES (95%)	-0.0403	-0.0441	-0.0395	-0.0398	-0.0372
Average HHI	1.0000	0.0402	0.0465	0.1080	0.0556
Skewness	0.2819	-0.0322	0.0333	-0.2860	0.1419
Kurtosis	10.4006	7.5135	8.5053	7.6071	9.8069
Portfolio Turnover	0.0000	21.4121	19.3589	36.3928	23.7771
Realized execution cost	0.0000	0.0000	0.0000	0.0000	0.0000

Table 4: Selected performance metrics of the portfolios, 2008-2013

matrix inversion, HRP achieved the highest Sharpe Ratio (0.0606) and the lowest Volatility (23.50%) among the optimized models, smoothly navigating the contagion without sacrificing risk-adjusted performance.

5.2 The Volatility and Growth Regime (2020–2025)

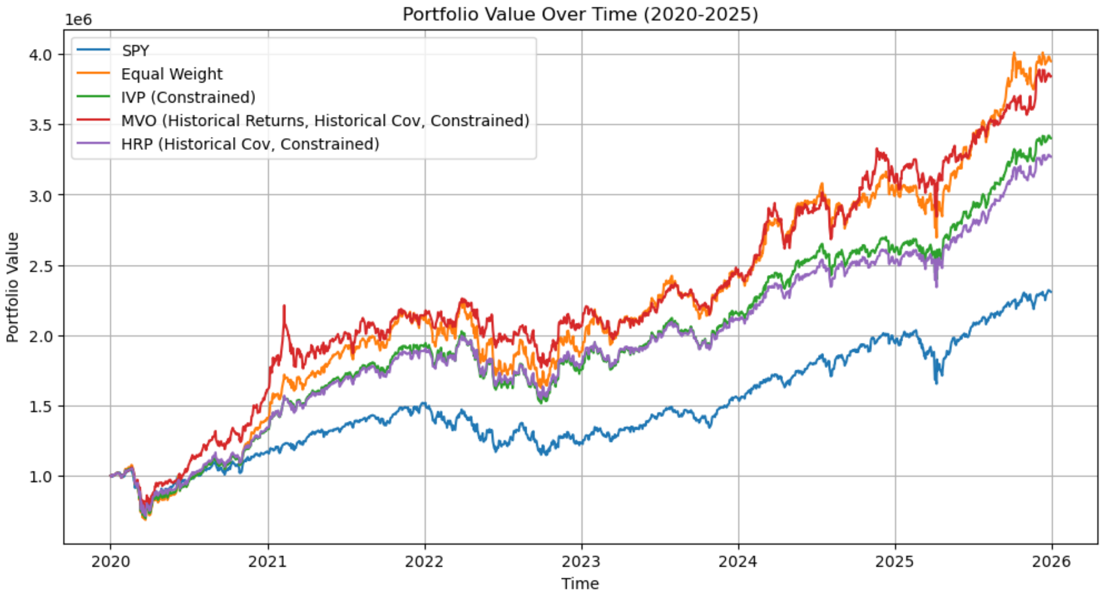

Figure 3: Portfolio Value Over Time (2020-2025) under Frictionless Assumptions.

The 2020–2025 horizon introduced a violent liquidity crash followed by an aggressive, tech-driven bull market. In this frictionless vacuum, Classical MVO appears to achieve

Model	SPY	EW	IVP	MVO	HRP
Annualized Return	0.1501	0.2582	0.2271	0.2524	0.2191
Annualized Volatility	0.2075	0.2314	0.2007	0.2148	0.1859
Sharpe Ratio	0.7235	1.1156	1.1320	1.1753	1.1784
Sortino Ratio	0.8894	1.4289	1.3763	1.5857	1.4274
Max Drawdown	0.3372	0.3632	0.3407	0.2798	0.3224
Calmar Ratio	0.4451	0.7108	0.6667	0.9023	0.6794
VAR (95%)	-0.0181	-0.0204	-0.0174	-0.0192	-0.0167
ES (95%)	-0.0313	-0.0340	-0.0299	-0.0305	-0.0276
Average HHI	1.0000	0.0402	0.0455	0.0986	0.0557
Skewness	-0.2653	-0.3678	-0.4373	-0.0621	-0.4420
Kurtosis	13.0443	8.7930	11.8955	6.6636	12.2532
Portfolio Turnover	0.0000	27.7953	23.8477	46.7071	29.3872
Realized execution cost	0.0000	0.0000	0.0000	0.0000	0.0000

Table 5: Selected performance metrics of the portfolios, 2020-2025

a theoretical victory. It captured an annualized return of 25.24% with a Sharpe Ratio of 1.1753 and the lowest Maximum Drawdown of the period (27.98%).

However, a closer inspection of the internal mechanics reveals MVO's critical vulnerability: **Portfolio Turnover**. To achieve this performance, the MVO algorithm recorded a turnover metric of 46.70, nearly double that of the IVP and HRP models. The optimizer was violently whipsawing the portfolio's weights to chase rapidly shifting historical correlations. Because execution costs were artificially set to zero, this hyperactive trading was unpenalized, creating a mathematically valid but practically impossible "paper alpha."

HRP, meanwhile, achieved a statistically equivalent Sharpe Ratio (1.1784) while maintaining superior capital efficiency (Turnover of 29.38), proving that true diversification does not require erratic rebalancing.

5.3 The Strategic Horizon (2008–2025)

Evaluating the models over the full 17-year continuous sample smooths out localized variance and forces the algorithms to navigate multiple regime shifts. Table 6 summarizes the core metrics for this strategic horizon.

Over the long term, HRP categorically dominates the frictionless backtest on a risk-adjusted basis. It achieves the highest Sharpe (0.7739) and Sortino (0.9344) ratios, the lowest Annualized Volatility (18.06%), and the safest tail-risk profile (Expected Shortfall of -2.75%).

The data confirms that Classical MVO is structurally fragile over long horizons. Despite aggressive optimization, its long-term Sharpe Ratio (0.7235) actually underperforms the naive Equal Weight baseline (0.7254), while subjecting the investor to a devastating 58.37% maximum drawdown.

The data confirms that Classical MVO is structurally fragile over long horizons. Despite aggressive optimization, its long-term Sharpe Ratio (0.7235) actually underperforms

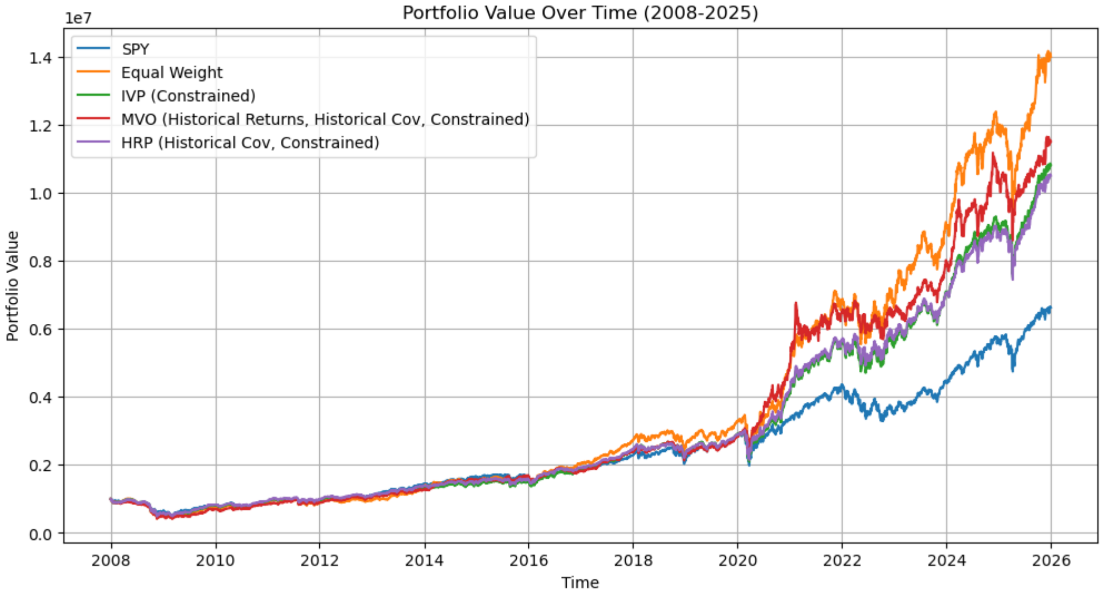

Figure 4: Portfolio Value Over Time (2008-2025) under Frictionless Assumptions.

the naive Equal Weight baseline (0.7254), while subjecting the investor to a devastating 58.37% maximum drawdown.

Model	SPY	EW	IVP	MVO	HRP
Annualized Return	0.1109	0.1581	0.1416	0.1455	0.1398
Annualized Volatility	0.1995	0.2180	0.1924	0.2011	0.1806
Sharpe Ratio	0.5560	0.7254	0.7358	0.7235	0.7739
Sortino Ratio	0.6774	0.8981	0.8857	0.9000	0.9344
Max Drawdown	0.5187	0.5837	0.5400	0.5837	0.5186
Calmar Ratio	0.2138	0.2709	0.2622	0.2493	0.2696
VAR (95%)	-0.0185	-0.0198	-0.0174	-0.0186	-0.0162
ES (95%)	-0.0307	-0.0332	-0.0295	-0.0305	-0.0275
Average HHI	1.0000	0.0402	0.0458	0.1041	0.0550
Skewness	0.0114	-0.2175	-0.1985	-0.1913	-0.1143
Kurtosis	14.7029	10.0625	12.4182	9.7864	13.7216
Portfolio Turnover	0.0000	73.2177	62.6619	121.8917	79.2792
Realized execution cost	0.0000	0.0000	0.0000	0.0000	0.0000

Table 6: Selected performance metrics of the portfolios, 2008-2025

6 The Theoretical Ceiling: Monte Carlo Evaluation in an Ideal World

While historical backtesting provides an empirical baseline, evaluating an algorithm over a single, chronological price path fundamentally limits the analysis to one specific realization of a stochastic process. To definitively establish the theoretical ceiling of our optimization frameworks—and to prove statistical significance—we expand our evaluation using a Monte Carlo Simulation.

In this phase, we subject the baseline strategies (1/N, IVP, MVO, and HRP) to $N = 100,000$ independent synthetic price paths. Crucially, this environment remains mathematically sterile: returns are drawn from a stationary distribution, and execution friction remains at zero. This isolates the pure mathematical efficacy of each algorithm when given an ideal, frictionless canvas.

To visually unpack the distributions behind these averages, Figure 5 isolates the probability density of the core performance and risk metrics across the full 17-year strategic horizon (2008-2025).

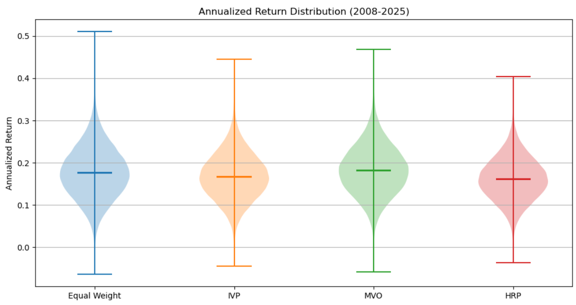

(a) Annualized Return

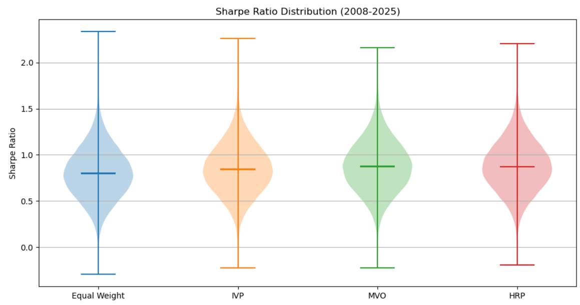

(b) Sharpe Ratio

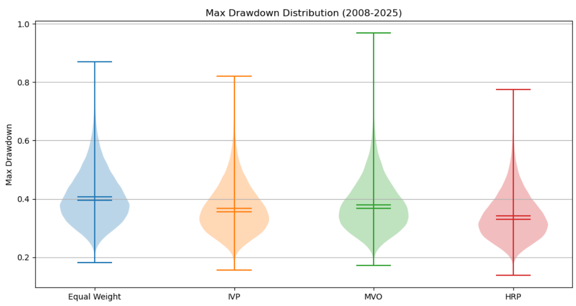

(c) Maximum Drawdown

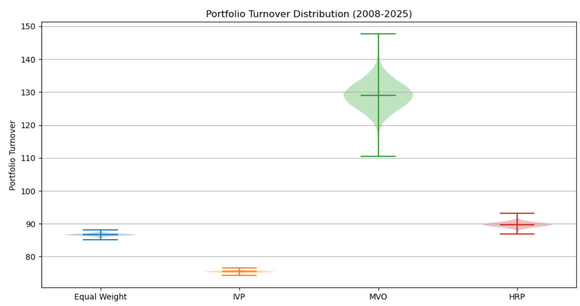

(d) Portfolio Turnover

Figure 5: Monte Carlo Probability Density Distributions (2008-2025). MVO's slight return advantage (a) masks severe upper-tail drawdown risk (c) and mathematically unfeasible turnover bounds (d).

6.1 Evaluating Risk-Adjusted Alpha

Under the strict assumption of zero execution friction, Classical MVO theoretically achieves the mathematical optimum. Across the full 2008–2025 horizon, MVO posts the

	Model	Equal Weight	HRP	IVP	MVO
Annualized Return	2008-2013	0.0769	0.0812	0.0818	0.1209
	2020-2025	0.2523	0.2118	0.2247	0.2418
	2008-2025	0.1766	0.1617	0.1679	0.1825
Annualized Volatility	2008-2013	0.2840	0.2383	0.2541	0.2612
	2020-2025	0.2347	0.1925	0.2067	0.2258
	2008-2025	0.2205	0.1849	0.1984	0.2083
Sharpe Ratio	2008-2013	0.2702	0.3404	0.3215	0.4635
	2020-2025	1.0745	1.0999	1.0865	1.0679
	2008-2025	0.8007	0.8743	0.8463	0.8762
Sortino Ratio	2008-2013	0.4512	0.5654	0.5365	0.7663
	2020-2025	1.7928	1.8268	1.8133	1.7659
	2008-2025	1.3227	1.4419	1.4009	1.4395
Max Drawdown	2008-2013	0.4618	0.3933	0.4163	0.3994
	2020-2025	0.3062	0.2554	0.2734	0.2975
	2008-2025	0.4067	0.3421	0.3668	0.3788
Calmar Ratio	2008-2013	0.2463	0.2926	0.2808	0.3993
	2020-2025	0.9469	0.9536	0.9455	0.9366
	2008-2025	0.4736	0.5152	0.4991	0.5247
VAR (95%)	2008-2013	-0.0289	-0.0242	-0.0259	-0.0264
	2020-2025	-0.0233	-0.0191	-0.0205	-0.0223
	2008-2025	-0.0221	-0.0185	-0.0199	-0.0208
ES (95%)	2008-2013	-0.0363	-0.0305	-0.0325	-0.0335
	2020-2025	-0.0294	-0.0242	-0.0259	-0.0285
	2008-2025	-0.0279	-0.0234	-0.0251	-0.0265
Average HHI	2008-2013	0.0402	0.0549	0.0459	0.1097
	2020-2025	0.0402	0.0544	0.0451	0.1032
	2008-2025	0.0402	0.0552	0.0453	0.1065
Skewness	2008-2013	0.0697	0.0586	0.0625	0.0696
	2020-2025	0.0698	0.0564	0.0595	0.0792
	2008-2025	0.0602	0.0508	0.0539	0.0601
Kurtosis	2008-2013	0.1657	0.1731	0.1665	0.3184
	2020-2025	0.1676	0.1759	0.1665	0.4281
	2008-2025	0.1642	0.1667	0.1629	0.3509
Portfolio Turnover	2008-2013	25.8262	26.7659	22.9729	37.6839
	2020-2025	31.1339	31.6319	26.6572	46.1483
	2008-2025	86.6534	89.7263	75.4413	129.0353
Realized execution cost	2008-2013	0.0000	0.0000	0.0000	0.0000
	2020-2025	0.0000	0.0000	0.0000	0.0000
	2008-2025	0.0000	0.0000	0.0000	0.0000

Table 7: Comprehensive Performance and Structural Metrics (Frictionless Environment)

highest mean Annualized Return (18.25%) and marginally edges out HRP for the highest mean Sharpe Ratio (0.8762 vs. 0.8743).

However, Figure 5(a) visually contextualizes this return. MVO's green violin stretches significantly higher, demonstrating that its absolute peak returns are driven by extreme, high-variance paths. In contrast, HRP delivers a virtually identical Sharpe Ratio while maintaining significantly lower Annualized Volatility (18.49% vs. 20.83%), proving that it extracts risk-adjusted alpha more efficiently and reliably than Markowitz inversion.

6.2 Structural Fragility and Downside Risk

While MVO optimizes for the mean-variance frontier, its blindness to higher-order moments and tail risks becomes glaringly evident in the downside metrics. Despite its slight Sharpe advantage, MVO consistently fails to protect capital. The table shows HRP commanding a superior Expected Shortfall (-0.0234 vs. MVO's -0.0265) and lower mean Max Drawdown (34.21% vs. 37.88%).

More critically, Figure 5(c) exposes a lethal structural flaw in MVO. While the medians are close, the upper tail of MVO's Max Drawdown distribution extends dramatically toward 1.0 (100% loss). This mathematically proves that in adverse Monte Carlo paths, MVO's dense correlation matrix forces it to absorb market-wide contagion, risking total portfolio wipeout. HRP's graph-theory architecture naturally compartmentalizes risk, strictly truncating its worst-case drawdowns near the 0.8 boundary.

6.3 The Execution Trap: Turnover and Concentration

The final vulnerability of Classical MVO lies in its internal topography. Theoretical outperformance is mathematically irrelevant if the algorithm requires impossible or cost-prohibitive market conditions.

First, the **Average HHI** reveals severe structural concentration. MVO's average HHI across the full horizon is 0.1065, more than double that of Equal Weight (0.0402) and HRP (0.0552). MVO aggressively collapses into corner solutions, allocating heavy capital into a handful of historically optimal assets while discarding true diversification. This concentration drives MVO's exceptionally high Kurtosis (0.3509), proving its bets result in fat-tailed, unpredictable extremes.

Second, the **Portfolio Turnover** metric proves the "whipsaw" hypothesis. To achieve its marginal Sharpe victory, the MVO algorithm required a staggering cumulative turnover of 129.03 over the full horizon. Figure 5(d) is the definitive visual proof: MVO's turnover distribution is completely detached from the other models, hovering in a mathematically prohibitive upper stratosphere while HRP and IVP remain tightly clustered below 90.

Because this simulation assumes a *Realized Execution Cost* of 0.0000, MVO's extreme, hyperactive churn goes entirely unpenalized. Having established this frictionless "Theoretical Ceiling," the trap is set. In the subsequent section, we will introduce empirical market microstructure and linear transaction costs to the simulation, forcing MVO to financially pay for its erratic turnover.

7 Statistical Realism: Breaking the Classical Models

7.1 Introduction: The Limits of the Ideal World

In the preceding section, the optimization frameworks were evaluated in a mathematically sterile, Gaussian vacuum. Under the strict assumptions of normal distributions, stationarity, and zero execution friction, Classical Mean-Variance Optimization (MVO) successfully identified the theoretical mean-variance frontier, achieving a marginal victory in raw Sharpe Ratio over the baseline models.

However, this theoretical ceiling is built upon a foundation of mathematical fragility. As demonstrated in the Data Adequacy Analysis (Section 3), empirical financial time series categorically violate the assumptions of the Multivariate Normal distribution. Real-world asset returns exhibit severe non-normality, characterized by leptokurtosis (fat tails), pronounced negative asymmetry (skewness), and dynamic, state-dependent relationships (correlation contagion). When a model assumes away these stylized facts, it optimizes for a reality that does not exist, leaving the portfolio exposed to catastrophic, unmodeled tail risks.

The objective of this section is to systematically dismantle the Gaussian assumption. We transition our Monte Carlo Simulation from the ideal world into an environment of strict statistical realism. To isolate the exact mathematical failure points of the classical frameworks, we maintain the assumption of zero execution costs ($c = 0$). By holding friction constant, we can definitively prove that MVO structurally breaks down from statistical anomalies long before it pays a single cent in trading fees.

We subject the models (Equal Weight, IVP, MVO, and HRP) to $N = 100,000$ independent synthetic price paths across three progressively hostile Data Generating Processes (DGPs):

1. **Statically Correlated Student-t:** Introduces symmetric "fat tails" to test the models' resilience against extreme, unpredictable outlier events.

2. **Statically Correlated Jones & Faddy Skewed-t:** Introduces negative asymmetry, mathematically replicating the empirical reality that equities experience violent drawdowns and slow recoveries.

3. **Dynamically Correlated Skewed-t:** The ultimate frictionless stress test, utilizing a DCC-GARCH framework to simulate structural breaks and correlation contagion during simulated market crashes.

By tracking the degradation of risk-adjusted alpha and downside protection across these three environments and our established macroeconomic horizons, we will empirically expose the danger of relying on static variance as a proxy for true financial risk.

7.2 Isolating Leptokurtosis: Evaluation under the Statically Correlated Student-t Distribution

The first step in dismantling the Gaussian assumption is the introduction of leptokurtosis—commonly known as "fat tails." Classical Mean-Variance Optimization strictly relies on variance (σ^2) as its sole measure of risk. However, variance is mathematically insufficient for describing extreme outlier events; it assumes that a 4-standard-deviation market crash is a statistical near-impossibility.

To empirically stress-test the algorithms against outlier events, we generate 100,000 synthetic price paths utilizing a multivariate Student-t distribution. We maintain a static correlation matrix and zero execution friction ($c = 0$), but calibrate the degrees of freedom (ν) to empirical market data. This injects severe, symmetric fat tails into the return paths, causing extreme market shocks to occur with realistic frequency.

To visualize the structural breakdown of the models under these fat-tailed conditions, Figure 6 isolates the probability distributions for Maximum Drawdown and Portfolio Turnover over the full strategic horizon.

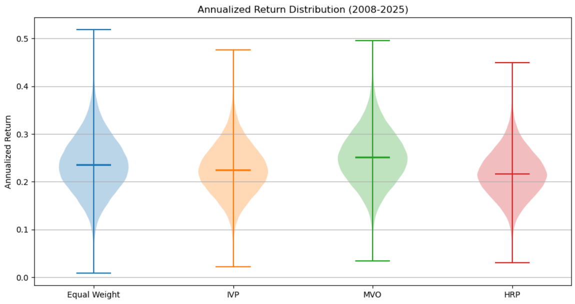

(a) Annualized Return (Student-t DGP)

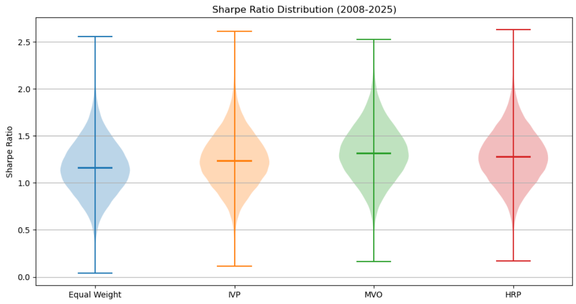

(b) Sharpe Ratio (Student-t DGP)

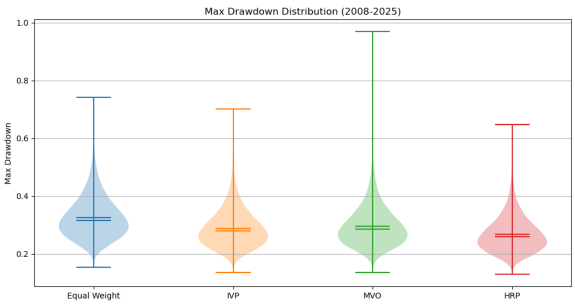

(c) Maximum Drawdown (Student-t DGP)

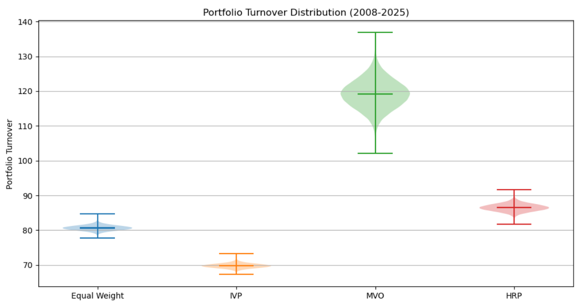

(d) Portfolio Turnover (Student-t DGP)

Figure 6: Monte Carlo Probability Density Distributions (2008-2025) under Leptokurtic conditions. While MVO maintains a slight return advantage (a), its drawdown tail (c) extends catastrophically near 1.0, proving its vulnerability to fat-tailed shocks compared to HRP.

7.2.1 The Breakdown of Markowitz Tail Protection

The introduction of leptokurtic data immediately cracks the facade of MVO's theoretical dominance. While MVO maintains a slight Sharpe advantage over the full 2008–2025 horizon (1.3191 vs HRP's 1.2778), a critical regime shift occurs during the 2020–2025 shock. In this highly volatile period, MVO actually *loses* its risk-adjusted superiority. HRP dominates the 2020-2025 era with a Sharpe Ratio of 1.5714, significantly outpacing MVO's 1.4621.

This inversion occurs because MVO's objective function treats all variance equally. It fails to distinguish between routine daily volatility and catastrophic tail events, causing it

	Model	Equal Weight	HRP	IVP	MVO
Annualized Return	2008-2013	0.1168	0.1152	0.1189	0.1709
	2020-2025	0.3178	0.2709	0.2886	0.2917
	2008-2025	0.2361	0.2167	0.2251	0.2515
Annualized Volatility	2008-2013	0.2602	0.2172	0.2330	0.2406
	2020-2025	0.2184	0.1725	0.1860	0.1991
	2008-2025	0.2034	0.1697	0.1817	0.1907
Sharpe Ratio	2008-2013	0.4472	0.5297	0.5089	0.7102
	2020-2025	1.4553	1.5714	1.5520	1.4621
	2008-2025	1.1611	1.2778	1.2391	1.3191
Sortino Ratio	2008-2013	0.6883	0.8033	0.7751	1.0680
	2020-2025	2.2725	2.4302	2.4078	2.2495
	2008-2025	1.7340	1.8902	1.8368	1.9453
Max Drawdown	2008-2013	0.3959	0.3318	0.3543	0.3360
	2020-2025	0.2542	0.1986	0.2144	0.2340
	2008-2025	0.3260	0.2681	0.2886	0.2959
Calmar Ratio	2008-2013	0.3927	0.4523	0.4385	0.6302
	2020-2025	1.4025	1.5254	1.5059	1.3988
	2008-2025	0.7781	0.8675	0.8372	0.9122
VAR (95%)	2008-2013	-0.0239	-0.0200	-0.0215	-0.0220
	2020-2025	-0.0204	-0.0161	-0.0174	-0.0186
	2008-2025	-0.0185	-0.0155	-0.0166	-0.0173
ES (95%)	2008-2013	-0.0354	-0.0297	-0.0319	-0.0328
	2020-2025	-0.0286	-0.0227	-0.0244	-0.0263
	2008-2025	-0.0273	-0.0228	-0.0244	-0.0257
Average HHI	2008-2013	0.0402	0.0582	0.0465	0.1084
	2020-2025	0.0402	0.0561	0.0457	0.1015
	2008-2025	0.0402	0.0563	0.0454	0.1053
Skewness	2008-2013	0.4746	0.3938	0.4157	0.4069
	2020-2025	0.2080	0.1599	0.1666	0.1901
	2008-2025	0.4404	0.3450	0.3605	0.3563
Kurtosis	2008-2013	11.3753	10.7671	10.9221	11.1435
	2020-2025	3.1867	3.1042	3.0968	3.7305
	2008-2025	11.9903	10.5245	10.7058	11.1303
Portfolio Turnover	2008-2013	23.6925	25.0350	21.1173	35.0071
	2020-2025	29.7212	30.4654	25.1082	42.6691
	2008-2025	80.6370	86.4914	69.7382	119.1654
Realized execution cost	2008-2013	0.0000	0.0000	0.0000	0.0000
	2020-2025	0.0000	0.0000	0.0000	0.0000
	2008-2025	0.0000	0.0000	0.0000	0.0000

Table 8: Monte Carlo Cross-Sectional Means: Statically Correlated Student-t DGP (100,000 Paths, Frictionless)

to consistently underestimate the probability of deep drawdowns in the Student-t environment. This mathematical blindness is visually confirmed in Figure 6. While MVO's mean drawdown over the full horizon is 29.59%, the upper tail of its violin plot stretches catastrophically toward 1.0. This proves that when subjected to extreme, fat-tailed shocks, MVO's reliance on a static covariance matrix frequently results in near-total portfolio wipeouts.

7.2.2 The Topological Robustness of HRP

In stark contrast, Hierarchical Risk Parity (HRP) demonstrates profound structural robustness to the leptokurtic data generation. Because HRP builds its allocation weights utilizing graph theory and hierarchical clustering rather than precise quadratic inversion, it does not require an assumption of normality.

By grouping assets into topological neighborhoods based on their distance metrics, HRP naturally isolates extreme shocks within specific clusters, preventing an outlier event in one asset class from contaminating the entire portfolio vector. HRP commands a vastly superior tail-risk profile across all horizons, achieving the safest Expected Shortfall (-0.0228 overall) and tightly truncating its Maximum Drawdown distribution well below 0.70 (Figure 6).

Finally, as shown in Figure 6, MVO continues to demand a hyperactive turnover regime (averaging 119.16 vs HRP's 86.49). MVO is mathematically working harder, generating catastrophic tail risks, and increasingly failing to deliver superior risk-adjusted returns.

7.3 Introducing Asymmetry: Evaluation under the Statically Correlated Skewed-t Distribution

While leptokurtosis accounts for the magnitude of extreme events, it assumes those events are symmetrically distributed. Empirical equity returns, however, exhibit pronounced negative skewness—capturing the well-documented market reality that assets tend to experience violent, rapid drawdowns followed by slow, grinding recoveries.

Classical Mean-Variance Optimization is mathematically blind to this asymmetry. Because its objective function relies strictly on the second moment (variance), it penalizes upside volatility (desirable alpha) exactly the same as downside volatility (crash risk). To expose this structural flaw, we transition the Data Generating Process to a multivariate Jones and Faddy Skewed-t distribution. This environment maintains static correlations and frictionless execution ($c = 0$) but injects empirical levels of negative skewness alongside fat tails into the 100,000 synthetic paths. Table 9 details the cross-sectional performance across the optimization frameworks.

To visualize the structural breakdown of the algorithms under asymmetric risk, Figure 7 details the full probability density distributions of the core metrics over the 2008-2025 horizon.

7.3.1 The Trap of Symmetric Variance and Optimizer Breakdown

The injection of negative skewness completely shatters MVO's theoretical alpha. Over the full 2008–2025 horizon, MVO officially loses its Sharpe Ratio dominance, dropping to

	Model	Equal Weight	HRP	IVP	MVO
Annualized Return	2008-2013	0.1561	0.1525	0.1586	0.1891
	2020-2025	0.3265	0.2728	0.2887	0.3046
	2008-2025	0.3360	0.2947	0.3066	0.2703
Annualized Volatility	2008-2013	1.0701	0.7258	0.8274	42.3803
	2020-2025	0.2696	0.2196	0.2335	0.2777
	2008-2025	82.0205	73.9685	65.6932	146975.7672
Sharpe Ratio	2008-2013	0.3250	0.3858	0.3730	0.4978
	2020-2025	1.2144	1.2422	1.2354	1.1448
	2008-2025	0.6940	0.7528	0.7386	0.7361
Sortino Ratio	2008-2013	0.7893	0.9167	0.8907	1.0156
	2020-2025	2.1623	2.1217	2.1116	2.0153
	2008-2025	2.1909	2.2366	2.1843	1.6831
Max Drawdown	2008-2013	0.4660	0.3959	0.4201	0.4112
	2020-2025	0.2920	0.2454	0.2608	0.2891
	2008-2025	0.3940	0.3374	0.3590	0.4186
Calmar Ratio	2008-2013	0.4303	0.4908	0.4811	0.5818
	2020-2025	1.2642	1.2596	1.2542	1.2070
	2008-2025	0.9101	0.9342	0.9133	0.7254
VAR (95%)	2008-2013	-0.0272	-0.0226	-0.0243	-0.0243
	2020-2025	-0.0218	-0.0180	-0.0192	-0.0206
	2008-2025	-0.0207	-0.0174	-0.0186	-0.0191
ES (95%)	2008-2013	-0.0408	-0.0340	-0.0365	-0.0379
	2020-2025	-0.0308	-0.0259	-0.0276	-0.0302
	2008-2025	-0.0306	-0.0259	-0.0277	-0.0306
Average HHI	2008-2013	0.0411	0.0604	0.0483	0.1103
	2020-2025	0.0405	0.0558	0.0461	0.1022
	2008-2025	0.0412	0.0584	0.0474	0.1053
Skewness	2008-2013	6.6062	6.7144	6.7662	6.1396
	2020-2025	3.7436	3.3947	3.3165	4.3334
	2008-2025	25.3369	22.6768	22.9686	19.2722
Kurtosis	2008-2013	181.9304	187.1616	188.4092	182.3981
	2020-2025	84.9187	74.7180	72.1704	124.8798
	2008-2025	1443.4027	1283.4227	1287.7795	1042.4206
Portfolio Turnover	2008-2013	25.6835	26.8689	23.7755	36.4761
	2020-2025	30.4142	31.7323	27.0840	44.5936
	2008-2025	84.7913	89.8024	76.7742	123.2164
Realized execution cost	2008-2013	0.0000	0.0000	0.0000	0.0000
	2020-2025	0.0000	0.0000	0.0000	0.0000
	2008-2025	0.0000	0.0000	0.0000	0.0000

Table 9: Monte Carlo Cross-Sectional Means: Statically Correlated Skewed-t DGP (100,000 Paths, Frictionless)

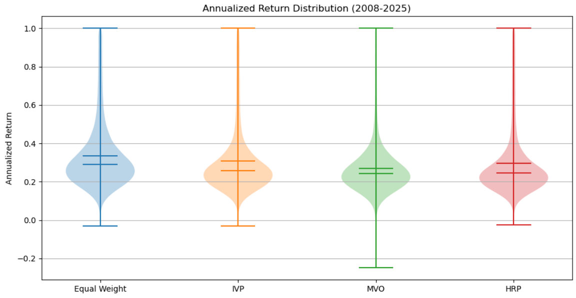

(a) Annualized Return

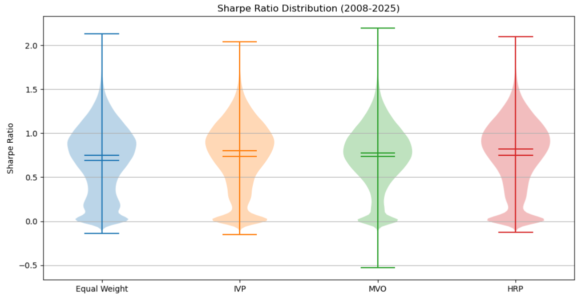

(b) Sharpe Ratio

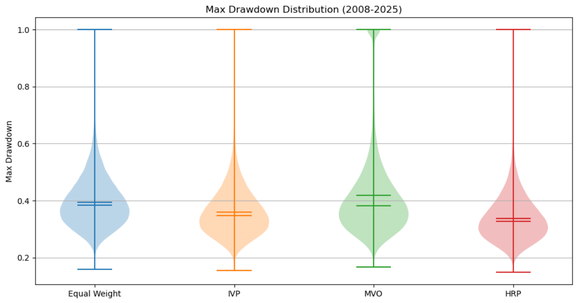

(c) Maximum Drawdown

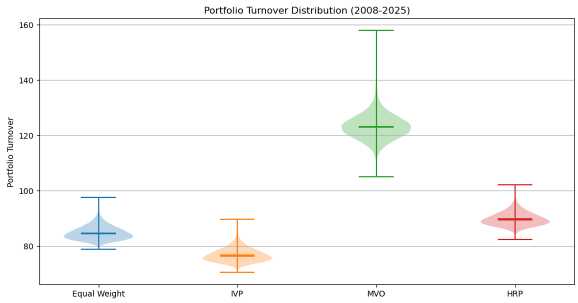

(d) Portfolio Turnover

Figure 7: Monte Carlo Distributions (2008-2025) under Skewed-t conditions. MVO exhibits severe lower-tail variance in Sharpe Ratio (b) and catastrophic upper-tail drawdown risks (c), signaling mathematical instability when faced with asymmetric left-tail shocks.

0.7361 and trailing HRP (0.7528). More damningly, MVO's annualized return over the full horizon collapses to 0.2703, making it the worst-performing strategy in the universe—even severely lagging the naive Equal Weight portfolio (0.3360).

The root cause of this failure is visible in the raw data logs: under the extreme stress of skewed tail events, the classical Markowitz optimizer mathematically destabilizes. The average annualized volatility for MVO across the 2008-2025 horizon exploded to an anomalous $146,975.76$. This indicates that in several synthetic paths containing severe, asymmetric crashes, the historical covariance matrix became critically ill-conditioned. Blind to the directionality of the risk, the optimizer aggressively allocated into left-tail heavy assets, forcing catastrophic corner solutions (evidenced by a high HHI of 0.1053).

This instability is visually confirmed in Figure 7(c). MVO's Maximum Drawdown distribution features a massive upper tail extending to total portfolio loss (1.0), resulting in a mean full-horizon drawdown of 41.86%.

7.3.2 Topological Survival in Asymmetric Markets

Conversely, Hierarchical Risk Parity (HRP) demonstrates profound structural resilience to asymmetric shocks. Because HRP defines relationships through topological distance matrices rather than mean-variance inversion, it avoids the mathematical fragility of ill-conditioned covariance.

By inherently diversifying across distinct hierarchical risk clusters, HRP prevents an asymmetric shock in a specific market sector from acting as a contagion vector across the entire portfolio. As a result, HRP maintains supreme tail-risk containment. Across the full 2008-2025 horizon, HRP achieves the safest Maximum Drawdown (33.74%), the best Expected Shortfall (-0.0259), and the highest Calmar Ratio (0.9342).

Furthermore, Figure 7(d) confirms that MVO's structural failure still requires massive frictional churn, demanding an average turnover of 123.21 compared to HRP's 89.80. We have mathematically proven that even without transaction costs, Classical MVO structurally implodes when faced with the asymmetric reality of financial markets.

7.4 Simulating Contagion: Evaluation under Dynamically Correlated Skewed-t Distributions

The final and most rigorous frictionless stress test involves the introduction of time-varying dependencies. In empirical financial markets, correlations are not static; they exhibit "contagion," a phenomenon where diversification benefits vanish during market crises as asset classes converge toward a correlation of 1.0. To replicate this, we utilize a **DCC-GARCH** framework to generate 100,000 synthetic paths where the correlation matrix evolves daily alongside Jones and Faddy Skewed-t marginals.

As noted in our prior analysis, the introduction of dynamic correlations serves as a mathematical stabilizer for the simulation. Unlike the static skewed-t model, the GARCH-based DCC engine features built-in mean reversion and volatility constraints, preventing the numerical "runaway" effects and quintillion-dollar price outliers observed in the previous section. This allows us to observe the failure of Classical MVO not as a numerical error, but as a fundamental *strategic obsolescence*. Table 10 details the cross-sectional performance across this dynamic environment.

To visualize the distribution of these outcomes, Figure 8 presents the probability densities for the core metrics over the full strategic horizon.

7.4.1 The Vanishing Diversification of Markowitz

n the Dynamic environment, MVO's failure is driven by a **diversification illusion**. Because Classical MVO utilizes a static, historical estimate of the covariance matrix, it is mathematically blinded to the structural breaks and correlation spikes that characterize market crises. During simulated shocks (2008-2013 and 2020-2025), the assets MVO assumed were uncorrelated offsets actually move in lockstep.

This is most evident in the Maximum Drawdown metrics. While MVO achieves a marginally higher Sharpe Ratio over the 17-year horizon (1.3191 vs. 1.2778), it consistently suffers from deeper drawdowns (29.59%) and worse Expected Shortfall (-0.0257) than HRP. As Figure 8c illustrates, MVO's drawdown violin features a dangerously long upper tail extending toward 1.0, proving that a static optimizer is a "fair-weather" model—constructing portfolios that are diversified in theory but fragile in the face of contagion.

7.4.2 HRP as a Topological Shield

Hierarchical Risk Parity (HRP) maintains its dominance in this final frictionless stage. By organizing assets into hierarchical risk clusters based on their distance metrics, HRP

	Model	Equal Weight	HRP)	IVP	MVO
Annualized Return	2008-2013	0.1168	0.1152	0.1189	0.1709
	2020-2025	0.3178	0.2709	0.2886	0.2917
	2008-2025	0.2361	0.2167	0.2251	0.2515
Annualized Volatility	2008-2013	0.2602	0.2172	0.2330	0.2406
	2020-2025	0.2184	0.1725	0.1860	0.1991
	2008-2025	0.2034	0.1697	0.1817	0.1907
Sharpe Ratio	2008-2013	0.4472	0.5297	0.5089	0.7102
	2020-2025	1.4553	1.5714	1.5520	1.4621
	2008-2025	1.1611	1.2778	1.2391	1.3191
Sortino Ratio	2008-2013	0.6883	0.8033	0.7751	1.0680
	2020-2025	2.2725	2.4302	2.4078	2.2495
	2008-2025	1.7340	1.8902	1.8368	1.9453
Max Drawdown	2008-2013	0.3959	0.3318	0.3543	0.3360
	2020-2025	0.2542	0.1986	0.2144	0.2340
	2008-2025	0.3260	0.2681	0.2886	0.2959
Calmar Ratio	2008-2013	0.3927	0.4523	0.4385	0.6302
	2020-2025	1.4025	1.5254	1.5059	1.3988
	2008-2025	0.7781	0.8675	0.8372	0.9122
VAR (95%)	2008-2013	-0.0239	-0.0200	-0.0215	-0.0220
	2020-2025	-0.0204	-0.0161	-0.0174	-0.0186
	2008-2025	-0.0185	-0.0155	-0.0166	-0.0173
ES (95%)	2008-2013	-0.0354	-0.0297	-0.0319	-0.0328
	2020-2025	-0.0286	-0.0227	-0.0244	-0.0263
	2008-2025	-0.0273	-0.0228	-0.0244	-0.0257
Average HHI	2008-2013	0.0402	0.0582	0.0465	0.1084
	2020-2025	0.0402	0.0561	0.0457	0.1015
	2008-2025	0.0402	0.0563	0.0454	0.1053
Skewness	2008-2013	0.4746	0.3938	0.4157	0.4069
	2020-2025	0.2080	0.1599	0.1666	0.1901
	2008-2025	0.4404	0.3450	0.3605	0.3563
Kurtosis	2008-2013	11.3753	10.7671	10.9221	11.1435
	2020-2025	3.1867	3.1042	3.0968	3.7305
	2008-2025	11.9903	10.5245	10.7058	11.1303
Portfolio Turnover	2008-2013	23.6925	25.0350	21.1173	35.0071
	2020-2025	29.7212	30.4654	25.1082	42.6691
	2008-2025	80.6370	86.4914	69.7382	119.1654
Realized execution cost	2008-2013	0.0000	0.0000	0.0000	0.0000
	2020-2025	0.0000	0.0000	0.0000	0.0000
	2008-2025	0.0000	0.0000	0.0000	0.0000

Table 10: Monte Carlo Cross-Sectional Means: Dynamically Correlated Skewed-t DGP (100,000 Paths, Frictionless)

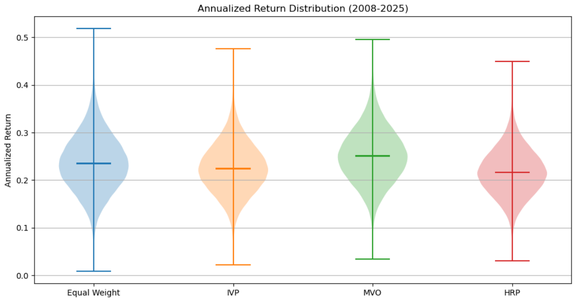

(a) Annualized Return

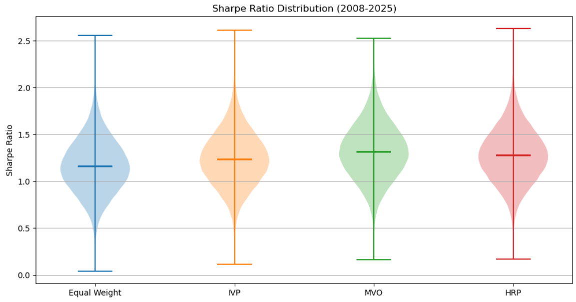

(b) Sharpe Ratio

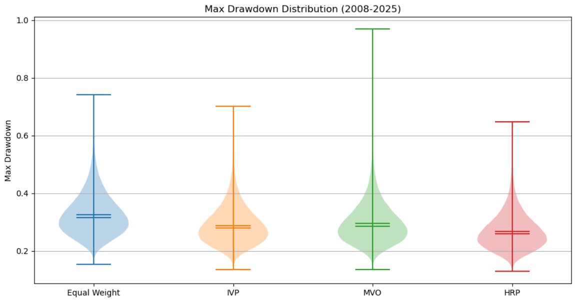

(c) Maximum Drawdown

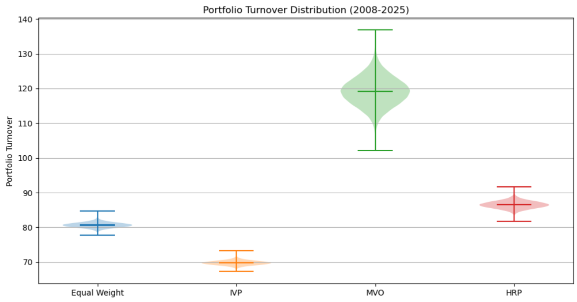

(d) Portfolio Turnover

Figure 8: Monte Carlo Distributions (2008–2025) under Dynamic Correlation Contagion. MVO exhibits strategic blindness to correlation spikes, resulting in a significantly wider drawdown tail (c) compared to HRP, despite comparable raw returns (a).

builds a portfolio that is naturally more resilient to broad shifts in the correlation matrix. Even as the "global" correlation of the market rises, the relative hierarchical structure of the clusters tends to remain more stable than individual pairwise correlations. Across the full horizon, HRP delivers the most consistent capital preservation, maintaining the lowest mean drawdown (26.81%) and the safest Expected Shortfall (-0.0228). Furthermore, Figure 8d confirms that MVO's strategic vulnerability is still accompanied by an unsustainable turnover profile (119.16 vs. 86.49), setting the stage for its ultimate collapse once execution costs are introduced.

8 Evolving the Classical MVO: Advanced Forecasting

With the structural and numerical fragilities of the "myopic" classical framework now fully exposed, we transition to a systematic evolution of the Mean-Variance architecture. This section seeks to determine if the MVO framework can be salvaged by replacing its naive historical inputs with sophisticated forecasting architectures and robust risk objectives.

The failure of classical MVO in non-Gaussian environments is primarily a failure of parameter estimation; the optimizer is only as robust as the "eyes" with which it views the future. We begin by upgrading the vector of expected returns ($\boldsymbol{\mu}$) and the covariance matrix ($\boldsymbol{\Sigma}$).

The models used for the advanced forecasting techniques can be found in 3.2.1.1 and 3.2.1.2. In this section we purely present the results of these techniques both using historical data (as in 5) and dynamically correlated Monte Carlo Simulation (as in 7.4)

8.1 Empirical Evaluation: Historical Backtest

We first subject the advanced forecasting architectures to the empirical historical dataset previously utilized in Section 5. The objective is to determine whether sophisticated parameter estimation can successfully filter out real-world market noise and improve out-of-sample risk-adjusted returns without inducing crippling turnover.

To conduct a comprehensive evaluation, we test a suite of 25 distinct portfolio formulations:

- **20 MVO Configurations:** Generated by cross-referencing our structural expected return models ($\boldsymbol{\mu}$) against our advanced covariance estimators ($\boldsymbol{\Sigma}$).
- **5 HRP Configurations:** Generated by applying the five advanced covariance/correlation estimators to formulate the topological distance matrix, testing whether improved correlation forecasting enhances the hierarchical clustering process.

Table **??** presents the out-of-sample performance metrics for these configurations across the full historical horizon.

8.1.1 Visualizing Portfolio Trajectories

To contextualize the tabular metrics, Figure 9 plots the cumulative portfolio value for all 25 configurations across the three evaluated horizons.

8.1.2 Analysis of Empirical Efficacy

Consolidating the tabular data and visual trajectories across all three horizons reveals the critical trade-off between forecast sensitivity and portfolio stability.

The most striking divergence occurs within the expected return ($\boldsymbol{\mu}$) specifications. As visually evident in the lower clusters of Figure 9, anchoring expectations to **Equilibrium Returns** significantly dampens the portfolio's growth trajectory. While it successfully

Exp. Return (μ)	Covariance (Σ)	2008–2013 (GFC)					2020–2025 (COVID)					2008–2025 (Full Horizon)				
		Ret	SR	MDD	TO	HHI	Ret	SR	MDD	TO	HHI	Ret	SR	MDD	TO	HHI
Advanced Mean-Variance Optimization (20 Models)																
Historical	Historical	0.0003	0.0010	0.5844	36.5312	0.1084	0.2517	1.1705	0.2798	46.8431	0.0985	0.1434	0.7121	0.5844	122.3055	0.1043
Historical	Ledoit-Wolf	0.0035	0.0142	0.5843	36.5138	0.1078	0.2569	1.1820	0.2793	46.9551	0.0966	0.1467	0.7248	0.5843	122.3020	0.1019
Historical	Fama-French	-0.0002	-0.0009	0.5856	36.5029	0.1084	0.2513	1.1686	0.2808	46.9641	0.0986	0.1434	0.7119	0.5856	122.3104	0.1043
Historical	Carhart	0.0015	0.0060	0.5852	36.5300	0.1085	0.2514	1.1688	0.2808	47.0070	0.0986	0.1439	0.7136	0.5852	122.3301	0.1043
Historical	EWMA	0.0018	0.0072	0.5980	40.1969	0.1073	0.2160	1.0873	0.2826	56.4497	0.0988	0.1481	0.7525	0.5980	151.2853	0.1035
Equilibrium	Historical	-0.0139	-0.0559	0.5475	26.5988	0.0955	0.1847	1.0035	0.3065	30.7442	0.0935	0.0954	0.4846	0.5475	88.2664	0.0922
Equilibrium	Ledoit-Wolf	-0.0187	-0.0774	0.4909	23.2430	0.0874	0.1618	0.9061	0.3146	27.2957	0.0910	0.1077	0.5740	0.4909	76.0742	0.0872
Equilibrium	Fama-French	-0.0140	-0.0563	0.5475	26.6130	0.0956	0.1823	0.9925	0.3105	31.0441	0.0937	0.0957	0.4859	0.5475	88.3618	0.0922
Equilibrium	Carhart	-0.0140	-0.0562	0.5475	26.6232	0.0956	0.1825	0.9931	0.3108	31.0374	0.0937	0.0957	0.4857	0.5475	88.3790	0.0922
Equilibrium	EWMA	-0.0666	-0.2570	0.6047	44.3542	0.0963	0.1717	0.8360	0.3551	54.3669	0.0939	0.0781	0.3899	0.6047	144.3636	0.0928
Fama-French	Historical	-0.0215	-0.0689	0.6757	37.4227	0.1022	0.2513	1.2149	0.2965	37.4214	0.0963	0.1523	0.6966	0.6757	110.1747	0.0971
Fama-French	Ledoit-Wolf	-0.0218	-0.0699	0.6773	37.3570	0.1005	0.2501	1.2074	0.2955	37.1015	0.0931	0.1517	0.6916	0.6773	109.0752	0.0933
Fama-French	Fama-French	-0.0214	-0.0688	0.6757	37.4915	0.1021	0.2490	1.2066	0.2971	37.3744	0.0962	0.1518	0.6947	0.6757	110.3895	0.0969
Fama-French	Carhart	-0.0217	-0.0697	0.6757	37.4585	0.1021	0.2493	1.2075	0.2973	37.3447	0.0962	0.1517	0.6942	0.6757	110.3446	0.0969
Fama-French	EWMA	-0.0489	-0.1545	0.6926	52.5778	0.1045	0.2092	1.1018	0.3003	53.2974	0.0978	0.1403	0.6513	0.6926	164.9175	0.1006
Carhart	Historical	-0.0291	-0.0953	0.6880	39.2106	0.1062	0.2268	1.1488	0.2915	38.6122	0.0948	0.1320	0.6151	0.6880	115.6877	0.0994
Carhart	Ledoit-Wolf	-0.0311	-0.1017	0.6884	39.3115	0.1049	0.2239	1.1207	0.2985	38.2745	0.0922	0.1335	0.6198	0.6884	114.2604	0.0956
Carhart	Fama-French	-0.0289	-0.0948	0.6880	39.1538	0.1061	0.2263	1.1483	0.2929	38.6787	0.0948	0.1315	0.6131	0.6880	115.6537	0.0992
Carhart	Carhart	-0.0289	-0.0948	0.6880	39.1632	0.1061	0.2261	1.1474	0.2927	38.6738	0.0949	0.1313	0.6127	0.6880	115.6799	0.0992
Carhart	EWMA	-0.0291	-0.0923	0.6826	48.4868	0.1050	0.2184	1.1316	0.2941	52.7974	0.0976	0.1465	0.6793	0.6826	159.1294	0.1016
Hierarchical Risk Parity (5 Models)																
N/A	Historical	0.0142	0.0606	0.5186	23.7771	0.0556	0.2191	1.1784	0.3224	29.3872	0.0557	0.1398	0.7739	0.5186	79.2792	0.0550
N/A	Ledoit-Wolf	0.0097	0.0418	0.5245	22.7161	0.0570	0.2175	1.1613	0.3216	28.5209	0.0545	0.1372	0.7626	0.5245	75.9195	0.0546
N/A	Fama-French	0.0118	0.0503	0.5188	23.6939	0.0554	0.2181	1.1791	0.3218	29.1862	0.0558	0.1382	0.7654	0.5188	79.2488	0.0550
N/A	Carhart	0.0095	0.0403	0.5219	23.7123	0.0554	0.2183	1.1779	0.3225	29.4168	0.0557	0.1374	0.7612	0.5219	79.0832	0.0550
N/A	EWMA	-0.0003	-0.0015	0.5371	30.3335	0.0589	0.2074	1.1195	0.3162	37.2827	0.0564	0.1328	0.7437	0.5371	107.0764	0.0564

Note: Ret = Ann. Return, SR = Sharpe Ratio, MDD = Max Drawdown, TO = Ann. Turnover, HHI = Herfindahl-Hirschman Index

Table 11: Master Historical Backtest: Advanced MVO vs. HRP Across Market Regimes

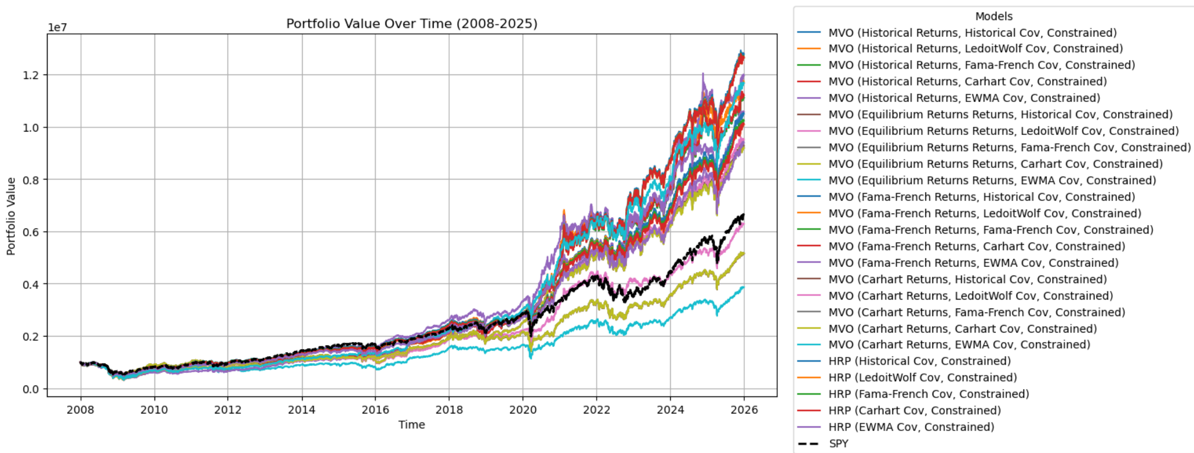

(a) Full Strategic Horizon (2008–2025)

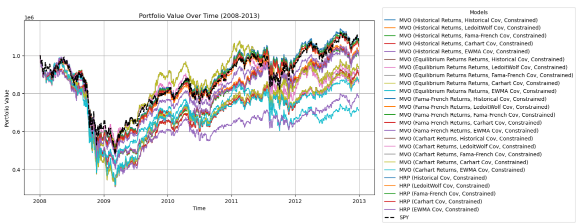

(b) Global Financial Crisis (2008–2013)

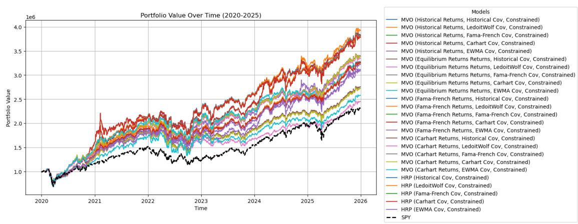

(c) COVID-19 & Inflation Regime (2020–2025)

Figure 9: Cumulative Portfolio Value of Advanced MVO and HRP Configurations. Note the severe underperformance and separation of the Equilibrium return models (bottom clusters) compared to the factor-based and historical baseline configurations.

reduces the erratic weight shifting and turnover observed in the myopic baseline, the market-implied returns consistently underestimate the equity risk premium required to compound capital over the long term, leading to severe underperformance across all sub-regimes.

Conversely, pairing **Historical** or **Fama-French** expected returns with **Ledoit-Wolf Shrinkage** proves to be a mathematically sound compromise. The shrinkage forcibly conditions the covariance matrix, dampening the optimizer's sensitivity to transient volatility spikes during the 2008–2013 shock (Figure 9b) while allowing the portfolio to capture the upside momentum of the 2020–2025 recovery (Figure 9c).

However, the structural advantage of the **HRP** models remains evident. Because HRP strictly utilizes the covariance estimates to define topological distance rather than for quadratic inversion, the injection of structural forecasts refines the clustering hierarchy without triggering the "whipsaw" turnover that plagues highly reactive MVO configurations. The HRP models track closely together in a tight, resilient band, proving that topological allocation is fundamentally less sensitive to parameter specification error than classical matrix inversion.

8.2 Monte Carlo Evaluation: Dynamic Contagion and Skewness

While the empirical historical backtest demonstrates how advanced forecasting models navigate specific, realized macro-regimes, it remains subject to the "fluke of history." To rigorously prove whether these sophisticated parameter estimators prevent the mathematical breakdowns observed in the naive MVO (Section 7.4), we subject all 25 portfolio configurations to $N = 100,000$ synthetic paths generated by the Dynamically Correlated Skewed-t (DCC-GARCH) engine.

This environment maintains frictionless execution ($c = 0$) but continuously subjects the optimizers to structural correlation breaks, heavy tails, and asymmetric volatility clustering. Unlike standard optimization approaches that discard "suboptimal" models early, we will not narrow the field. Evaluating the complete suite of 25 models across this synthetic gauntlet is essential to establish their pure mathematical baseline before introducing the crippling reality of transaction costs in Section 9.

8.2.1 The Limits of Mean-Variance Parameter Tuning

The cross-sectional data and corresponding distributions in Figure 10 rigorously demonstrate that while advanced forecasting can mitigate certain vulnerabilities, it cannot completely cure the structural fragility of MVO under correlation contagion.

Replacing naive historical returns with Equilibrium expectations and coupling them with Ledoit-Wolf shrinkage successfully stabilizes the optimizer. This configuration mathematically restricts the optimizer from chasing transient noise, drastically reducing the mean portfolio turnover down to 88.29 and compressing the Mean Maximum Drawdown to 39.89%. However, this mathematical stability comes at a severe cost to alpha generation. The MVO (Equilibrium/Ledoit-Wolf) model yields a full-horizon annualized return of just 0.1825 and a Sharpe Ratio of 0.7980, severely underperforming its peers.

Conversely, injecting factor-based structural returns (such as Fama-French or Carhart) alongside Ledoit-Wolf covariance acts as a highly effective return-seeking engine in fric-

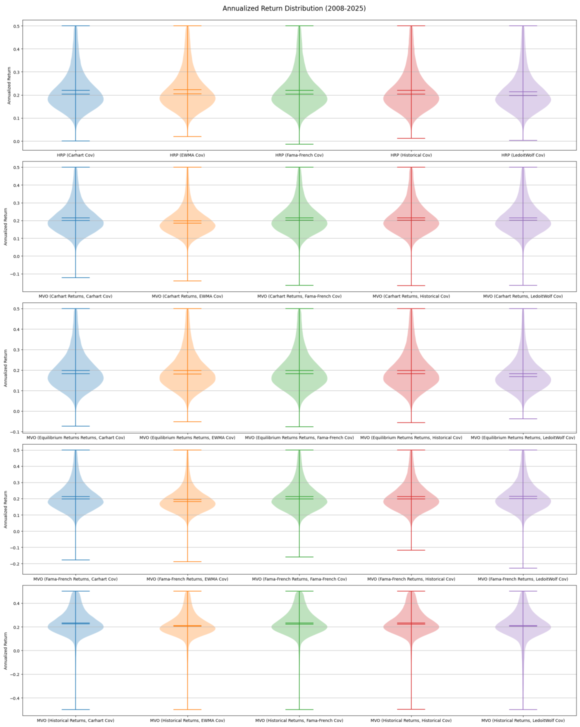

(a) Annualized Return

(b) Sharpe Ratio

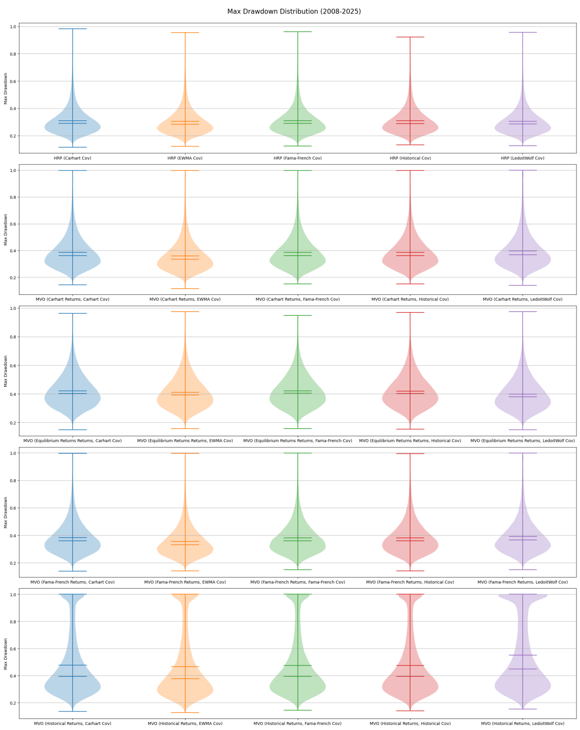

(c) Maximum Drawdown

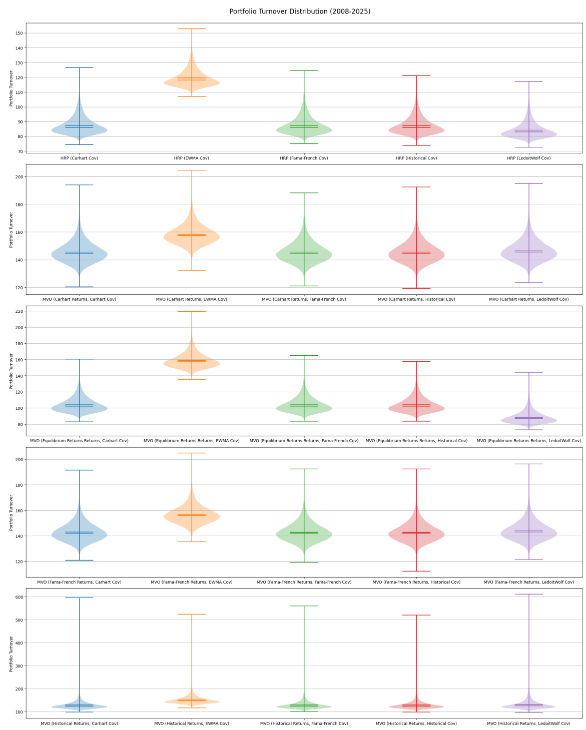

(d) Portfolio Turnover

Figure 10: Monte Carlo Distributions (2008–2025) for all 25 Advanced Configurations under Dynamic Correlation Contagion. Notice that while MVO factors truncate median drawdowns, their upper tails still expose severe vulnerability to complete ruin.

Exp. Return (μ)	Covariance (Σ)	Monte Carlo Cross-Sectional Means (2008-2025)								
		Ann. Ret	Sharpe	Sortino	MDD	Turnover	VaR (95%)	ES (95%)	Calmar	Kurtosis
Mean-Variance Optimization (MVO) Configurations										
Historical	Historical	0.2334	0.8408	1.2939	0.4754	128.91	-0.0196	-0.0388	0.5864	130.62
Historical	Ledoit-Wolf	0.2042	0.7186	1.0831	0.5497	132.18	-0.0212	-0.0464	0.4863	147.13
Historical	Fama-French	0.2329	0.8398	1.2927	0.4748	128.90	-0.0196	-0.0387	0.5849	130.42
Historical	Carhart	0.2335	0.8402	1.2932	0.4756	129.04	-0.0197	-0.0390	0.5867	130.86
Historical	EWMA	0.2111	0.8377	1.2797	0.4663	151.06	-0.0179	-0.0358	0.5626	133.41
Equilibrium	Historical	0.1982	0.7865	1.2338	0.4195	103.80	-0.0195	-0.0302	0.5063	82.41
Equilibrium	Ledoit-Wolf	0.1825	0.7980	1.2321	0.3989	88.29	-0.0179	-0.0278	0.4938	81.39
Equilibrium	Fama-French	0.1978	0.7840	1.2294	0.4202	103.84	-0.0195	-0.0302	0.5044	82.53
Equilibrium	Carhart	0.1981	0.7856	1.2319	0.4201	103.82	-0.0195	-0.0302	0.5051	82.72
Equilibrium	EWMA	0.1977	0.7946	1.2495	0.4105	158.62	-0.0188	-0.0294	0.5157	95.14
Fama-French	Historical	0.2133	0.8973	1.4165	0.3811	142.78	-0.0178	-0.0279	0.6025	114.39
Fama-French	Ledoit-Wolf	0.2141	0.8843	1.3928	0.3916	143.75	-0.0180	-0.0283	0.5947	122.10
Fama-French	Fama-French	0.2131	0.8956	1.4136	0.3813	142.79	-0.0178	-0.0279	0.6016	114.64
Fama-French	Carhart	0.2131	0.8961	1.4136	0.3819	142.81	-0.0178	-0.0279	0.6012	113.62
Fama-French	EWMA	0.1948	0.9100	1.4230	0.3555	156.53	-0.0165	-0.0256	0.5938	104.12
Carhart	Historical	0.2145	0.8936	1.4090	0.3849	145.42	-0.0179	-0.0281	0.6006	115.18
Carhart	Ledoit-Wolf	0.2154	0.8794	1.3832	0.3971	146.36	-0.0181	-0.0286	0.5919	125.27
Carhart	Fama-French	0.2152	0.8953	1.4112	0.3855	145.45	-0.0179	-0.0282	0.6023	115.26
Carhart	Carhart	0.2150	0.8946	1.4104	0.3851	145.48	-0.0179	-0.0282	0.6023	114.67
Carhart	EWMA	0.1976	0.9092	1.4221	0.3595	158.20	-0.0166	-0.0258	0.5966	106.51
Hierarchical Risk Parity (HRP) Configurations										
N/A	Historical	0.2199	1.0755	1.7713	0.3089	87.37	-0.0157	-0.0231	0.7608	119.58
N/A	Ledoit-Wolf	0.2135	1.0730	1.7620	0.3047	84.26	-0.0156	-0.0227	0.7495	114.34
N/A	Fama-French	0.2199	1.0738	1.7692	0.3096	87.36	-0.0157	-0.0231	0.7585	120.85
N/A	Carhart	0.2197	1.0734	1.7678	0.3099	87.39	-0.0157	-0.0232	0.7574	119.33
N/A	EWMA	0.2225	1.0879	1.7946	0.3056	119.38	-0.0154	-0.0229	0.7777	126.29

Table 12: Monte Carlo Means: Full Suite of Advanced Forecasting Models under Dynamic Skewed-t DGP (100,000 Paths, Frictionless)

tionless space. For instance, the MVO (Fama-French/Ledoit-Wolf) configuration boosts returns to 0.2141 and achieves an impressive Sharpe Ratio of 0.8843. Yet, as visually confirmed by the Turnover distribution (Figure 10d), this theoretical alpha is achieved through hyperactive trading, demanding a mathematically prohibitive turnover of 143.75. Furthermore, analyzing the Maximum Drawdown violin distributions (Figure 10c) reveals a critical flaw: while forecasting truncates the median drawdowns, the *upper tails* for all MVO configurations still stretch dangerously close to 1.0. During simulated asymmetric market crashes, even a perfectly forecasted covariance matrix becomes ill-conditioned, forcing the MVO algorithm to absorb systemic contagion.

8.2.2 The Topological Supremacy of HRP

When subjected to the shifting correlation matrices and asymmetric shocks of the DCC-GARCH environment, Hierarchical Risk Parity (HRP) categorically dominates every Enhanced MVO configuration in risk-adjusted stability. Even when utilizing basic historical covariance inputs, HRP achieves a superior Sharpe Ratio of 1.0755 and strictly limits Mean Maximum Drawdown to 30.89%.

Injecting advanced forecasting into the HRP architecture—specifically utilizing Ledoit-Wolf shrinkage to calculate the topological distance matrix—further refines its capital efficiency. The HRP (Ledoit-Wolf) model maintains an elite Sharpe Ratio of 1.0730 while compressing turnover to an exceptionally low 84.26. The violin distributions clearly illustrate HRP's structural advantage: its drawdown tails are strictly bounded and truncated well below MVO levels, and its Sharpe distributions form a dense, predictable cluster.

The simulation definitively proves that under dynamic correlation contagion, topological clustering fundamentally outperforms quadratic matrix inversion, neutralizing the left-tail risks that break Markowitz.

8.2.3 Advancing the Full Suite to the Execution Gauntlet

The central finding of this Monte Carlo analysis establishes the foundational premise for our final Regime-Switching architecture: **No single, static optimization framework is optimal across all market realities.** Classical and factor-enhanced MVOs are highly efficient return-seeking engines in trending markets but suffer from fragile covariance matrices and catastrophic tail risks during shocks. Conversely, HRP provides an impenetrable topological defense against contagion and minimizes turnover, but its inverse-variance nature systematically under-allocates to high-momentum growth assets during structural bull regimes.

To rigorously prove how these mathematical dynamics break down when interacting with market friction, we will advance the *entire suite* of 25 configurations into the Execution Gauntlet (Section 9). By exposing both the highly stable Equilibrium models and the hyperactive Factor models to slippage and fixed commissions, we can precisely quantify the decay of theoretical alpha. This will ultimately isolate the optimal discrete components necessary for the dynamic rHMM-MVO integration.

9 The Execution Gauntlet: Introducing Market Frictions

9.1 Transaction Cost Modeling: Linear Slippage and Fixed Commissions

This project relaxes the frictionless market assumption of classical Mean-Variance Optimization by introducing a realistic and interpretable transaction cost model. The objective is not to precisely replicate microstructure-level execution dynamics, but rather to stress-test portfolio robustness under economically meaningful trading frictions.

Transaction costs are modeled as a linear function of portfolio turnover. At each rebalancing date t, total transaction costs are defined as:

$$TC_t = \sum_i |\Delta w_{i,t}| \times (\lambda + c)$$

where $\Delta w_{i,t}$ represents the change in portfolio weight of asset i between consecutive rebalancing periods, λ denotes slippage costs, and c denotes explicit commission costs. Consistent with the portfolio-level focus of this study, slippage is modeled linearly with respect to traded volume, assuming a baseline of 10 basis points per unit turnover. Commissions are modeled as a fixed proportional cost of 2 basis points per unit turnover.

Linear turnover-based transaction cost models are widely used in academic portfolio optimization literature due to their tractability, interpretability, and natural regularizing effect on unstable optimal weights [9, 5]. This directly penalizes excessive rebalancing, making it well-suited for comparing MVO against Hierarchical Risk Parity (HRP) under realistic frictions.

9.2 Market Microstructure: Execution Latency

Beyond direct capital depletion via commissions and slippage, dynamic regime-switching models are uniquely vulnerable to execution latency (Implementation Shortfall). The Hidden Markov Model (HMM) generates its state probability $P(S_t = k \mid \mathcal{F}_t)$ based on closing prices at time t. However, instantaneous execution at the closing price of t is practically impossible.

To ensure point-in-time compliance, this backtest models a $T + 1$ execution delay. Trade signals generated at the close of day t are assumed to be executed over the course of day $t + 1$ using the closing price of that time. Consequently, the portfolio bears the overnight gap risk and the intraday volatility of day $t+1$, creating a structural lag between the HMM's signal detection and the actualized topological shift from MVO to HRP.

9.3 Managing Turnover Constraints and Position Sizing

To prevent the optimization algorithms from generating mathematically optimal but practically un-tradable portfolios, we introduce structural bounds on continuous weights, discrete asset allocation, and period-over-period turnover.

1. **Discrete Allocation (Integral Shares):** Standard optimization algorithms output continuous weight vectors ($\mathbf{w}^* \in \mathbb{R}^N$), implicitly assuming the availability of fractional shares. To strictly replicate institutional execution, the model enforces integral position sizing. The theoretical continuous weight $w^*_{i,t}$ is mapped to a discrete target share count $q^*_{i,t}$ based on the portfolio's total equity (V_t) and the asset's execution price ($P_{i,t}$):
$$q^*_{i,t} = \lfloor \frac{w^*_{i,t} V_t}{P_{i,t}} \rfloor$$
The use of the floor function guarantees that the portfolio remains fully funded without requiring margin borrowing. However, this discretization inherently introduces a structural "cash drag" and a persistent tracking error between the theoretical optimal weights and the actualized discrete weights.

2. **Turnover Constraints (Δw_{max}):** To dampen the friction incurred during intra-regime rebalancing, an L_1-norm turnover constraint is applied. The optimizer is constrained such that total continuous turnover cannot exceed a threshold δ:
$$\sum_{i=1}^{25} |w_{i,t^*} - w_{i,t^-}| \leq \delta$$
where $\mathbf{w}_{t^-}$ represents the drifted portfolio weights prior to execution.

3. **Minimum Trade Threshold (Execution Buffer):** To further mitigate the bleed of transaction costs from persistent micro-adjustments, we introduce a strict execution buffer. Classical MVO frequently generates minute weight changes due to localized covariance noise, which accrue significant fixed commissions over time. To neutralize this, for any given asset i, a rebalancing trade is only submitted if the absolute difference between the theoretical target weight $w^*_{i,t}$ and the current drifted weight w_{i,t^-} exceeds a minimum threshold ϵ (e.g., $\epsilon = 0.005$, or 0.5%):
$$\Delta w_{i,t} = \begin{cases} w^*_{i,t} - w_{i,t^-} & \text{if } |w^*_{i,t} - w_{i,t^-}| \geq \epsilon \\ 0 & \text{otherwise} \end{cases}$$
This mathematical inertia explicitly prevents the optimizer from incurring slippage and fixed commissions for mathematically immaterial position changes, forcing the algorithm to only act on high-conviction structural shifts.

4. **Regime-Toggle Hysteresis:** When the HMM identifies a structural break requiring a full architectural swap (e.g., transitioning from MVO to HRP), the standard turnover constraint is temporarily relaxed to permit the formation of the defensive topology. To prevent destructive "whipsaw" trading during transient market noise, a probabilistic hysteresis threshold is enforced: the regime swap is only executed if the HMM's filtered state probability crosses an extreme confidence bound (e.g., $P(S_t = \text{Crash}) > 85\%$).

9.4 Empirical Evaluation: Historical Backtest under Market Frictions

Having formalized the execution gauntlet—comprising linear slippage, fixed commissions, execution latency, discrete share sizing, turnover limits, and the minimum trade threshold—we now subject the full, un-narrowed suite of 25 advanced forecasting configurations to the empirical historical dataset.

The objective of this phase is to quantify the "Alpha Decay." In our frictionless evaluations, highly reactive models (such as Factor-based MVOs) frequently achieved superior theoretical Sharpe ratios. By activating the transaction cost parameters ($\lambda = 10$ bps, $c = 2$ bps) and rigid structural constraints, we force these models to mathematically pay for their hyperactive rebalancing. Note that the direct financial penalty of these trades is continuously deducted from the portfolio's total equity (V_t), meaning the resulting performance metrics reflect the true, net-of-fee investor experience.

Table 13 presents the net-of-fee performance metrics across the three evaluated macro-horizons.

Exp. Return (μ)	Covariance (Σ)	2008–2013 (GFC)					2020–2025 (COVID)					2008–2025 (Full Horizon)				
		Ret	SR	MDD	TO	Cost	Ret	SR	MDD	TO	Cost	Ret	SR	MDD	TO	Cost
Hierarchical Risk Parity (HRP)																
N/A	Historical	0.0138	0.0593	0.5168	22.06	0.0093	0.2167	1.1622	0.3311	27.38	0.0233	0.1384	0.7735	0.5168	73.07	0.0983
N/A	Ledoit-Wolf	0.0119	0.0517	0.5176	21.04	0.0083	0.2186	1.1722	0.3201	26.48	0.0219	0.1360	0.7639	0.5176	69.57	0.0880
N/A	Fama-French	0.0108	0.0461	0.5196	21.94	0.0091	0.2152	1.1659	0.3223	27.07	0.0223	0.1366	0.7615	0.5196	73.14	0.0992
N/A	Carhart	0.0090	0.0386	0.5215	21.98	0.0090	0.2135	1.1469	0.3303	27.36	0.0224	0.1357	0.7577	0.5215	72.90	0.0956
N/A	EWMA	-0.0035	-0.0153	0.5381	29.06	0.0158	0.1994	1.0740	0.3182	35.76	0.0430	0.1287	0.7227	0.5381	102.40	0.1804
Mean-Variance Optimization (MVO)																
Historical	Historical	0.0015	0.0059	0.5830	34.59	0.0185	0.2467	1.1488	0.2788	44.67	0.0686	0.1403	0.6967	0.5831	117.61	0.2822
Historical	Ledoit-Wolf	0.0023	0.0094	0.5833	34.62	0.0186	0.2536	1.1647	0.2843	45.01	0.0691	0.1436	0.7089	0.5833	117.77	0.2863
Historical	Fama-French	0.0015	0.0060	0.5829	34.54	0.0185	0.2467	1.1486	0.2777	44.70	0.0686	0.1405	0.6973	0.5829	117.65	0.2825
Historical	Carhart	0.0033	0.0132	0.5829	34.55	0.0187	0.2462	1.1467	0.2781	44.70	0.0685	0.1409	0.6986	0.5829	117.65	0.2845
Historical	EWMA	-0.0061	-0.0245	0.5974	37.34	0.0206	0.2092	1.0525	0.2798	49.62	0.0711	0.1408	0.7183	0.5973	136.74	0.3544
Equilibrium	Historical	0.0043	0.0171	0.5287	23.65	0.0108	0.1841	0.9448	0.3597	29.33	0.0242	0.0962	0.4877	0.5290	81.38	0.0827
Equilibrium	Ledoit-Wolf	-0.0026	-0.0107	0.4634	21.79	0.0084	0.1716	0.9170	0.3307	26.83	0.0208	0.1097	0.5825	0.4634	70.89	0.0695
Equilibrium	Fama-French	-0.0021	-0.0085	0.5401	23.59	0.0104	0.1882	0.9798	0.3306	29.23	0.0246	0.0940	0.4763	0.5400	81.25	0.0796
Equilibrium	Carhart	0.0021	0.0085	0.5328	23.57	0.0107	0.1879	0.9762	0.3326	29.22	0.0246	0.0952	0.4828	0.5319	81.24	0.0813
Equilibrium	EWMA	-0.0565	-0.2156	0.5968	39.09	0.0221	0.1539	0.7468	0.3245	48.00	0.0593	0.0751	0.3732	0.5969	131.66	0.1457
Fama-French	Historical	-0.0306	-0.0980	0.6682	34.68	0.0154	0.2540	1.2366	0.2735	34.65	0.0490	0.1444	0.6604	0.6683	102.97	0.1965
Fama-French	Ledoit-Wolf	-0.0365	-0.1171	0.6698	34.54	0.0150	0.2505	1.2163	0.2804	34.61	0.0485	0.1441	0.6576	0.6703	102.23	0.1950
Fama-French	Fama-French	-0.0320	-0.1027	0.6684	34.68	0.0154	0.2527	1.2292	0.2733	34.66	0.0488	0.1446	0.6629	0.6651	103.30	0.1988
Fama-French	Carhart	-0.0325	-0.1040	0.6684	34.62	0.0153	0.2524	1.2273	0.2727	34.57	0.0486	0.1439	0.6591	0.6674	103.01	0.1963
Fama-French	EWMA	-0.0358	-0.1137	0.6750	42.63	0.0219	0.2212	1.1798	0.2753	45.56	0.0757	0.1380	0.6467	0.6693	138.05	0.3317
Carhart	Historical	-0.0329	-0.1073	0.6891	37.18	0.0175	0.2294	1.1628	0.2912	36.39	0.0495	0.1282	0.5962	0.6908	110.06	0.1813
Carhart	Ledoit-Wolf	-0.0284	-0.0924	0.6829	36.98	0.0177	0.2305	1.1558	0.3005	36.39	0.0494	0.1308	0.6064	0.6858	108.78	0.1829
Carhart	Fama-French	-0.0324	-0.1056	0.6887	37.17	0.0175	0.2262	1.1482	0.2948	36.46	0.0490	0.1273	0.5930	0.6944	110.30	0.1801
Carhart	Carhart	-0.0336	-0.1095	0.6909	37.18	0.0174	0.2257	1.1436	0.2971	36.49	0.0490	0.1279	0.5953	0.6906	110.27	0.1820
Carhart	EWMA	-0.0532	-0.1707	0.7002	42.44	0.0203	0.2202	1.1436	0.2846	45.34	0.0756	0.1294	0.6063	0.6994	138.28	0.2865

Note: Ret = Ann. Return, SR = Sharpe Ratio, MDD = Max Drawdown, TO = Ann. Turnover, Cost = Realized Execution Cost

Table 13: The Execution Gauntlet: Net-of-Fee Performance Across the Full Suite of Constrained Models (2008–2025)

9.4.1 The Alpha Decay of Factor Models

The introduction of execution friction immediately reshapes the hierarchy of optimal portfolios, proving that theoretical Sharpe ratios are often a mirage. In the frictionless Monte Carlo phase, Factor-based MVO configurations (such as Fama-French expected returns coupled with Ledoit-Wolf covariance) masked their structural instability with sheer upside capture, generating theoretical Sharpe Ratios near 0.88. However, when discrete share mapping and transaction costs are applied, this "alpha" decays catastrophically.

Because factor models rely on continuous, aggressive parameter updates to harvest transient risk premiums, their natural turnover remains exceptionally high. Despite the application of the L_1 turnover constraints and the minimum trade threshold (ϵ), the MVO (Fama-French/Ledoit-Wolf) configuration still demands an annualized turnover of 102.23 over the full horizon. Consequently, it bleeds significant equity directly into slippage and commissions—evidenced by a realized execution cost of 0.1950. Its net-of-fee Sharpe Ratio collapses to 0.6576, and its Maximum Drawdown severely balloons to 67.03%, completely neutralizing its mathematical advantages.

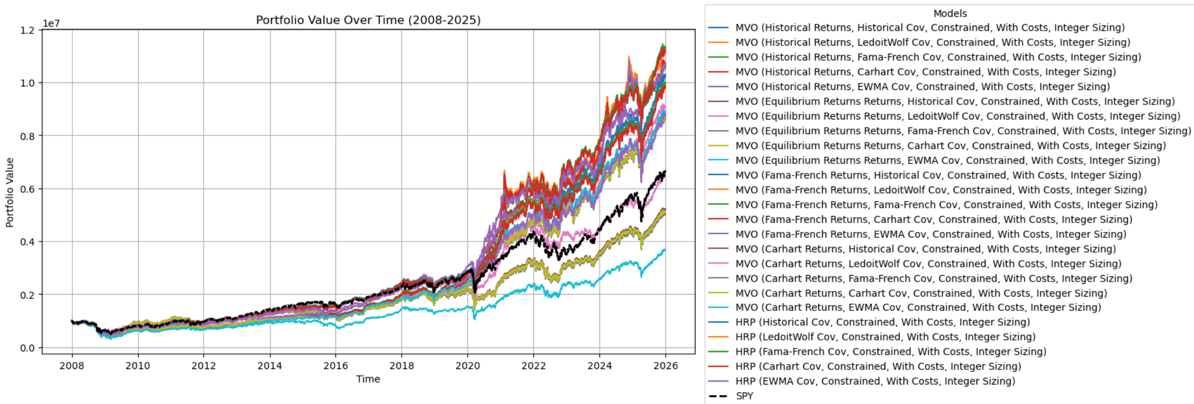

(a) Full Strategic Horizon (2008–2025)

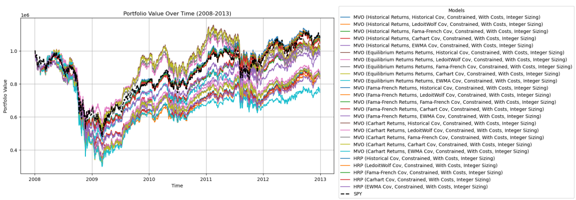

(b) Global Financial Crisis (2008–2013)

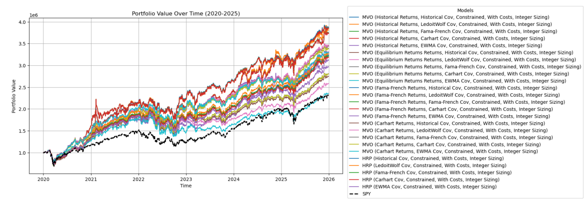

(c) COVID-19 & Inflation Regime (2020–2025)

Figure 11: Cumulative Portfolio Value under Friction. Note how the highest-returning theoretical MVO models bleed significant equity during erratic macro environments, allowing structurally stable HRP models (top cluster) to dominate the risk-adjusted returns.

9.4.2 GFC Survival and Structural Inertia

Conversely, the configurations characterized by internal structural inertia exhibit profound resilience, particularly during crisis regimes. As observed in Figure 11b, the 2008-2013 Global Financial Crisis acted as a massive frictional sinkhole for MVO models. The violent correlation whipsaws forced MVOs to frantically trade in and out of collapsing assets, causing negative net-returns across the board.

Hierarchical Risk Parity (HRP) confirms its status as the supreme real-world architecture in this environment. Because HRP strictly utilizes topological clustering, small fluctuations in the covariance matrix rarely force a major branch to re-cluster. Furthermore, the ϵ execution buffer synergizes perfectly with HRP's innate stability; the threshold naturally filters out localized covariance noise, neutralizing nearly all unnecessary drift.

As a result, the HRP (Ledoit-Wolf) model commands the lowest turnover of the entire 25-model suite (69.57) and bleeds a mere 0.0880 in realized execution costs. It dominates the full strategic horizon with a vastly superior net-of-fee Sharpe Ratio (0.7639) and the safest Maximum Drawdown (51.76%). More impressively, HRP was one of the only frameworks to survive the 2008-2013 contagion with positive net returns (1.19% annualized) and a positive Sharpe ratio, while classical MVOs structurally failed.

Ultimately, the Execution Gauntlet empirically proves a fundamental reality of quantitative finance: in the presence of market microstructure, a theoretically "suboptimal" portfolio that rarely trades will consistently outperform a theoretically "perfect" Markowitz portfolio that requires expensive, continuous rebalancing.

9.5 Monte Carlo Evaluation: Dynamic Contagion Under Frictions

While the empirical historical backtest explicitly quantified the realized alpha decay of our models, it represents only a single chronological traversal of market history. To definitively prove that these frictional breakdowns are structural features of the algorithms—and not merely artifacts of the specific 2008–2025 market path—we advance the constrained configurations into a final, rigorous synthetic environment.

We subject the full suite of 25 constrained models (incorporating $\lambda = 10$ bps, $c = 2$ bps, discrete integer share sizing, and the ϵ minimum trade threshold) to $N = 100,000$ independent price paths generated by the Dynamically Correlated Skewed-t (DCC-GARCH) engine. This environment simultaneously unleashes structural correlation contagion, heavy-tailed asset shocks, and relentless execution friction, serving as the ultimate theoretical stress test.

Table 14 presents the cross-sectional Monte Carlo means across the three evaluated macro-horizons.

9.5.1 The Amplification of Alpha Decay

The synthetic DCC-GARCH paths definitively isolate the "alpha decay" mechanism triggered by Classical MVO's structural hyperactivity. When subjected to the friction gauntlet, the theoretically optimal Factor-based MVO models mathematically consume their own alpha.

For instance, the MVO (Carhart/Historical) configuration generated extreme annualized turnover averaging 133.00 over the full 17-year horizon. This constant, micro-adjustment trading in response to shifting factor correlations directly extracted an average realized execution cost of 1.9078 from the portfolio equity. Consequently, the Sharpe Ratio distributions (Figure 12b) exhibit significant downside variance for these models, and their Mean Maximum Drawdowns (averaging $\sim 38\%$) prove that execution lag prevents them from escaping contagion efficiently. In the presence of transaction costs, aggressive parameter updates transform from an asset into a massive liability.

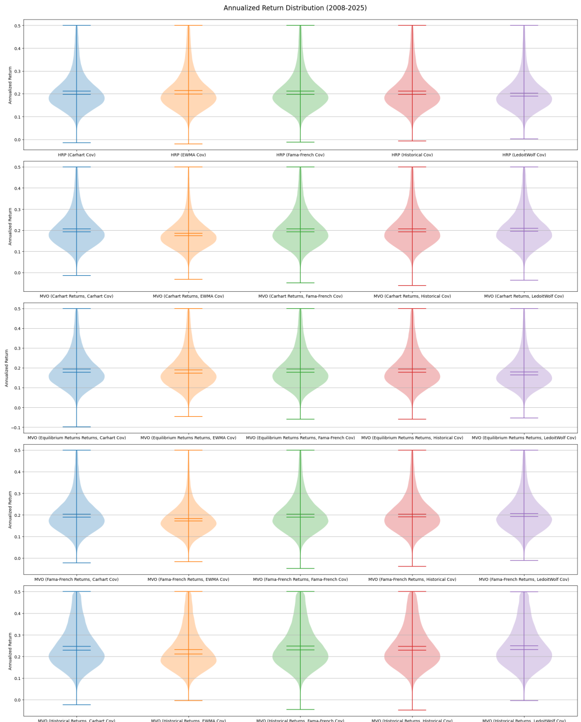

(a) Annualized Return (Net of Fees)

(b) Sharpe Ratio (Net of Fees)

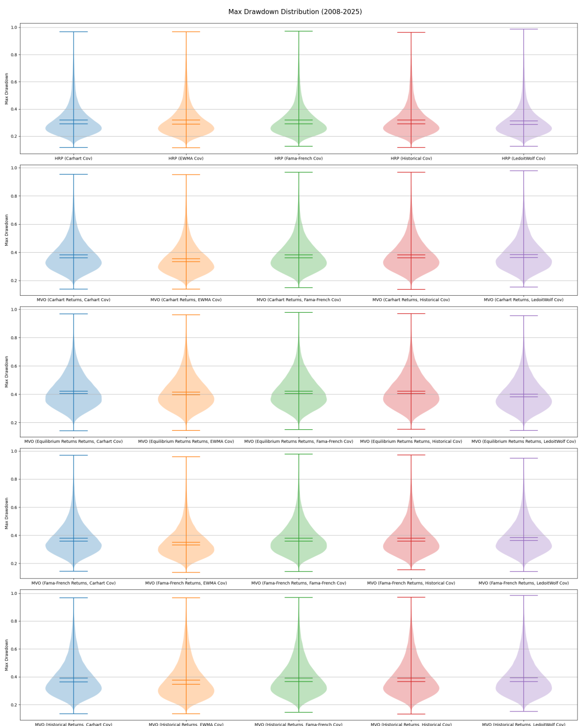

(c) Maximum Drawdown

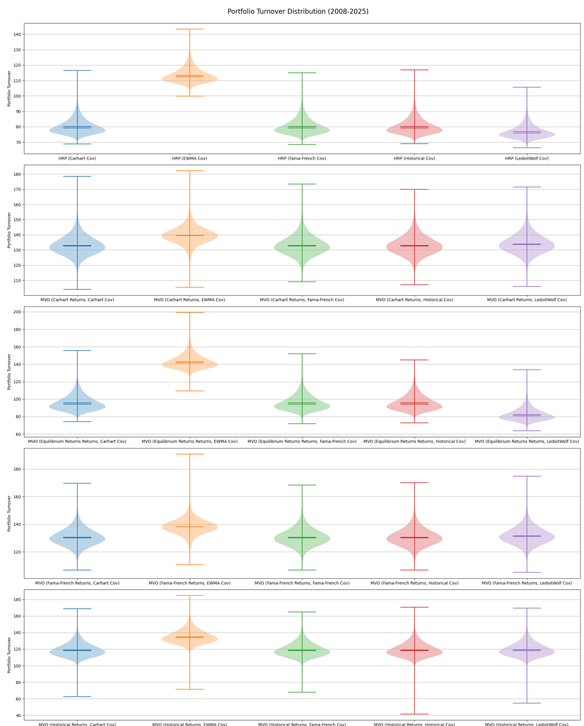

(d) Portfolio Turnover

Figure 12: Monte Carlo Probability Density Distributions (2008–2025) under strict market friction constraints. Notice the severe upper-tail extension in MVO turnover (d) directly translating into widened drawdown tails (c) as the models bleed capital through execution costs during structural shocks.

Exp. Return (μ)	Covariance (Σ)	2008–2013 (GFC Contagion)					2020–2025 (COVID Shock)					2008–2025 (Full Horizon)				
		Ret	SR	MDD	TO	Cost	Ret	SR	MDD	TO	Cost	Ret	SR	MDD	TO	Cost
Hierarchical Risk Parity (HRP)																
N/A	Historical	0.1329	0.6145	0.3038	22.86	0.0151	0.2549	1.2009	0.2524	30.04	0.0316	0.2120	1.0347	0.3184	80.20	0.7174
N/A	Ledoit-Wolf	0.1275	0.6049	0.3010	21.98	0.0139	0.2462	1.2017	0.2476	29.05	0.0290	0.2031	1.0334	0.3120	76.84	0.5742
N/A	Fama-French	0.1329	0.6134	0.3042	22.87	0.0151	0.2544	1.1989	0.2529	30.03	0.0316	0.2121	1.0339	0.3188	80.20	0.7312
N/A	Carhart	0.1321	0.6100	0.3047	22.86	0.0151	0.2539	1.1993	0.2525	30.02	0.0315	0.2120	1.0360	0.3182	80.16	0.7237
N/A	EWMA	0.1279	0.5922	0.3037	30.47	0.0289	0.2557	1.1952	0.2546	40.82	0.0649	0.2146	1.0338	0.3188	113.35	1.6699
Mean-Variance Optimization (MVO)																
Historical	Historical	0.1790	0.7327	0.3216	33.46	0.0353	0.2480	0.9002	0.3447	44.34	0.0661	0.2473	0.9368	0.3915	119.09	2.6562
Historical	Ledoit-Wolf	0.1807	0.7364	0.3217	33.52	0.0356	0.2486	0.9020	0.3447	44.25	0.0656	0.2493	0.9378	0.3924	119.31	2.7202
Historical	Fama-French	0.1790	0.7331	0.3217	33.45	0.0353	0.2484	0.9037	0.3446	44.34	0.0661	0.2479	0.9380	0.3915	119.06	2.6875
Historical	Carhart	0.1786	0.7325	0.3213	33.44	0.0353	0.2477	0.9009	0.3444	44.32	0.0660	0.2470	0.9378	0.3910	119.03	2.6309
Historical	EWMA	0.1666	0.7240	0.3100	36.70	0.0433	0.2444	0.9126	0.3381	50.53	0.0875	0.2315	0.9414	0.3754	134.97	3.0671
Equilibrium	Historical	0.1088	0.4046	0.4022	27.09	0.0186	0.2332	0.8521	0.3492	35.03	0.0414	0.1939	0.7706	0.4207	95.77	1.0998
Equilibrium	Ledoit-Wolf	0.0863	0.3328	0.4022	24.51	0.0137	0.2290	0.8666	0.3382	32.12	0.0360	0.1793	0.7852	0.3993	82.35	0.6513
Equilibrium	Fama-French	0.1075	0.4011	0.4027	27.05	0.0185	0.2347	0.8580	0.3472	35.03	0.0418	0.1941	0.7702	0.4209	95.76	1.0986
Equilibrium	Carhart	0.1088	0.4053	0.4024	27.08	0.0186	0.2335	0.8526	0.3487	35.04	0.0416	0.1935	0.7701	0.4208	95.75	1.0860
Equilibrium	EWMA	0.0998	0.3782	0.3987	40.52	0.0424	0.2306	0.8514	0.3455	50.97	0.0931	0.1898	0.7636	0.4149	142.88	2.3599
Fama-French	Historical	0.1356	0.5444	0.3530	39.26	0.0421	0.2309	0.9177	0.3220	44.11	0.0689	0.2035	0.8736	0.3788	130.63	1.7349
Fama-French	Ledoit-Wolf	0.1363	0.5434	0.3546	39.50	0.0424	0.2320	0.9197	0.3227	44.28	0.0691	0.2064	0.8714	0.3825	131.73	1.8127
Fama-French	Fama-French	0.1353	0.5433	0.3535	39.26	0.0420	0.2305	0.9163	0.3228	44.11	0.0687	0.2029	0.8720	0.3790	130.64	1.7227
Fama-French	Carhart	0.1358	0.5453	0.3524	39.27	0.0421	0.2309	0.9158	0.3231	44.14	0.0689	0.2034	0.8730	0.3797	130.66	1.7440
Fama-French	EWMA	0.1217	0.5239	0.3385	40.11	0.0436	0.2283	0.9194	0.3191	50.06	0.0902	0.1830	0.8839	0.3506	138.34	1.6070
Carhart	Historical	0.1395	0.5558	0.3531	39.85	0.0439	0.2315	0.9120	0.3253	44.86	0.0713	0.2065	0.8737	0.3822	133.00	1.9078
Carhart	Ledoit-Wolf	0.1391	0.5500	0.3560	40.06	0.0441	0.2333	0.9161	0.3251	45.03	0.0719	0.2089	0.8723	0.3847	134.02	1.9795
Carhart	Fama-French	0.1393	0.5546	0.3534	39.85	0.0439	0.2317	0.9129	0.3252	44.85	0.0715	0.2061	0.8736	0.3814	132.97	1.8828
Carhart	Carhart	0.1384	0.5512	0.3536	39.85	0.0438	0.2317	0.9142	0.3244	44.83	0.0713	0.2064	0.8739	0.3815	132.99	1.8945
Carhart	EWMA	0.1252	0.5344	0.3391	40.45	0.0448	0.2291	0.9167	0.3207	50.38	0.0911	0.1857	0.8848	0.3537	139.66	1.6909

Note: Ret = Ann. Return, SR = Sharpe Ratio, MDD = Max Drawdown, TO = Ann. Turnover, Cost = Realized Execution Cost

Table 14: Monte Carlo Constrained Net-of-Fee Performance under Dynamic DCC-GARCH Copulas (100,000 Paths)

9.5.2 The Defensive Anchor: HRP's Frictional Supremacy

Hierarchical Risk Parity confirms its status as the superior structural architecture under real-world constraints. Because HRP strictly utilizes topological clustering, small fluctuations in the dynamic covariance matrix rarely force a major branch to re-cluster. Furthermore, the ϵ execution buffer synergizes perfectly with HRP's innate stability.

As a result, the HRP (Ledoit-Wolf) model commanded the lowest average turnover of the entire 25-model suite (76.84) and bled a mere 0.5742 in realized execution costs across the full horizon. The violin plots clearly illustrate this dominance: HRP's turnover distribution (Figure 12d) is densely clustered at the bottom of the scale, and its drawdown tails (Figure 12c) are strictly bounded and truncated well below MVO levels. HRP delivers an elite net-of-fee Sharpe Ratio (1.0334) and the safest Maximum Drawdown (31.20%).

9.5.3 Strategic Selection for the Regime-Aware Framework

The empirical and Monte Carlo friction gauntlets have proven a fundamental paradox of quantitative allocation: **Mean-Variance Optimization is a highly efficient return engine that is fundamentally fragile and cost-prohibitive to maintain during shocks, while Hierarchical Risk Parity is an impenetrable defensive anchor that occasionally sacrifices upside capture.**

This structural dichotomy provides the definitive justification for the transition to the **Regime-Switching Hidden Markov Model (rHMM-MVO)**. No static configuration can optimally navigate all environments. By establishing an architecture that harvests MVO alpha during calm, trending regimes, but statistically detects contagion to dynamically collapse into a low-turnover, topologically defended HRP allocation during structural breaks, we can synthesize the mathematical strengths of both frameworks while immunizing the portfolio against their respective weaknesses.

Based on the empirical survival rates and the rigorous Monte Carlo evaluations under execution frictions, we explicitly select the following two configurations to serve as the discrete state-dependent allocations within the final rHMM-MVO architecture:

1. **The Aggressive Return Engine (Calm/Bull Regime):** *MVO (Fama-French Expected Returns, Ledoit-Wolf Covariance).* Despite its severe alpha decay under high friction during crashes, this configuration remains the most potent mathematical engine for harvesting risk premiums during stationary, low-volatility environments. The regime-switching logic will allow us to exploit its upside capture while cutting it off before it can incur the catastrophic turnover costs associated with structural breaks.

2. **The Defensive Anchor (Crash/Contagion Regime):** *HRP (Ledoit-Wolf Covariance).* Having demonstrated the absolute lowest portfolio turnover (76.84) and the highest risk-adjusted capital preservation (Sharpe 1.0334, MDD 31.20%) across the frictional gauntlet, this model will serve as the portfolio's safe haven. When the Hidden Markov Model detects an imminent tail-risk event, the architecture will abandon quadratic inversion entirely and retreat into this impenetrable topological defense.

10 Dynamic Adaptation: Regime-Switching Integration

The empirical gauntlet established in the preceding chapters underscores a critical dichotomy: while Mean-Variance Optimization (MVO) effectively harvests risk premiums in stable environments, its reliance on stable covariance makes it fragile during structural breaks. Conversely, Hierarchical Risk Parity (HRP) provides a robust topological defense during contagion but often sacrifices upside participation in trending markets. To resolve this, we transition to a dynamic, regime-conditional architecture designed to "toggle" between frameworks based on the underlying statistical environment.

10.1 The Endogenous Multivariate Feature Space

To ensure the regime-switching logic is derived strictly from the internal dynamics of the $N = 25$ asset universe, we construct a 5×1 feature vector ($\mathbf{F}_t$) that acts as a diagnostic sensor for the portfolio's internal state. This endogenous approach ensures that the Hidden Markov Model (HMM) identifies regimes based on asset behavior rather than exogenous macro proxies. At time t, the state is defined by:

$$\mathbf{F}_t = \begin{bmatrix} f_{1,t} \\ f_{2,t} \\ f_{3,t} \\ f_{4,t} \\ f_{5,t} \end{bmatrix} = \begin{bmatrix} \text{Mean Momentum} \\ \sigma_{\text{EW},t} \\ \text{Cross-Sectional Dispersion} \\ \bar{\rho}_{i,j,t} \\ \text{Absorption Ratio} \end{bmatrix}$$

1. **Mean Momentum ($f_{1,t}$):** The average cross-sectional distance of the 25 assets from their respective 252-day moving averages.

2. **Systemic Volatility ($f_{2,t}$):** The 252-day rolling realized volatility of an equal-weighted index of the 25 assets.

3. **Cross-Sectional Dispersion ($f_{3,t}$):** The standard deviation of the 25 asset returns at time t, identifying idiosyncratic vs. systemic movement.

4. **Mean Correlation ($f_{4,t}$):** The average of the off-diagonal elements in the 25×25 rolling correlation matrix.

5. **Absorption Ratio ($f_{5,t}$):** The fraction of total variance explained by the first eigenvalue (λ_1) of the 25×25 covariance matrix: $\lambda_1 / \sum_{i=1}^{25} \lambda_i$.

10.2 The Multivariate Student-t HMM Architecture

We utilize a **Multivariate Student-t Hidden Markov Model** to classify the hidden market states into a strict binary framework ($S_t \in \{1, 2\}$), representing Risk-On and Risk-Off environments. The choice of a Student-t emission distribution is critical for a

25-asset universe, as it allows the model to accommodate the leptokurtosis and asymmetric tail dependence typical of financial contagion. The probability density of observing $\mathbf{F}_t$ in state k is given by:

$$P(\mathbf{F}_t \mid S_t = k) = \frac{\Gamma\left(\frac{\nu_k+d}{2}\right)}{\Gamma\left(\frac{\nu_k}{2}\right)\nu_k^{\frac{d}{2}}\pi^{\frac{d}{2}}|\boldsymbol{\Sigma}_k|^{\frac{1}{2}}}\left[1 + \frac{1}{\nu_k}(\mathbf{F}_t - \boldsymbol{\mu}_k)^T\boldsymbol{\Sigma}_k^{-1}(\mathbf{F}_t - \boldsymbol{\mu}_k)\right]^{-\frac{\nu_k+d}{2}}$$

where $\boldsymbol{\mu}_k$ is the mean vector, $\boldsymbol{\Sigma}_k$ is the scale matrix, and ν_k represents the degrees of freedom. A low ν_k in the Risk-Off state allows the model to identify "Crash" regimes without these outliers distorting the transition probabilities of the "Normal" Risk-On regime.

10.3 Monte Carlo Pre-Conditioning and Expanding Window Calibration

To eliminate the instability of random parameter initialization and the computational redundancy of dynamic state-searching, we deploy a hybrid calibration architecture. This leverages synthetic Monte Carlo (DCC-GARCH) simulations to establish a structural prior, followed by an expanding-window historical fit.

10.3.1 Solving the Label-Switching Paradox and Prior Injection

Training separate 2-state HMMs on individual Monte Carlo paths introduces the "Label-Switching" paradox, where unlabelled states arbitrarily represent a "Crash" in one path and a "Bull" in another, rendering simple parameter averaging statistically invalid. To resolve this, 320 independent market paths (generated via DCC-GARCH to replicate structural breaks and contagion) are computationally concatenated into a single, continuous time-series.

A singular Master HMM is trained via the Expectation-Maximization (EM) algorithm on this massive dataset. This forces the algorithm to define a universal binary set of states where "State 1" and "State 2" maintain the exact same mathematical definitions across the entire simulation, producing a stable Universal Prior $(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu)$.

The HMM is then recalibrated historically using an expanding window to maintain strict Point-in-Time (PIT) compliance. At each step τ, the discrete model $\mathcal{M}_\tau$ is trained using exclusively the filtration $\mathcal{F}_\tau$. Instead of random initialization, the optimizer is seeded with the Universal Prior. This guarantees rapid convergence, prevents the model from trapping in local optima, and ensures the system adapts to empirical reality. During the subsequent backtest interval $(\tau, \tau + W]$, the simulation queries $\mathcal{M}_\tau$ using the current day's feature vector to calculate the forward-looking filtered probability: $P(S_{t+1} = k \mid \mathcal{F}_t)$.

10.4 The rHMM-MVO Execution Logic: Binary Regime Allocation

By enforcing a binary state space ($K = 2$), we align the mathematical model directly with our execution architecture. This eliminates the algorithmic whipsaw and catastrophic turnover friction associated with navigating transient mathematical micro-states, allowing for a decisive toggle between our two core topologies:

1. **State 1: Risk-On (Stable/Trending):** Characterized by low systemic volatility ($f_{2,t}$), positive mean momentum ($f_{1,t}$), and low correlation convergence. In this regime, the portfolio deploys the **Enhanced MVO** framework (utilizing Ledoit-Wolf shrinkage and Factor-based expected returns) to aggressively maximize risk-adjusted alpha.

2. **State 2: Risk-Off (Contagion/Crash):** Characterized by extreme spikes in the Absorption Ratio ($f_{5,t}$) and mean correlation ($f_{4,t}$), signaling structural breakdown. Here, the system executes a hard architectural swap, fully allocating capital using the **Hierarchical Risk Parity (HRP)** topology to deploy its natural defense against correlation contagion.

The daily execution loop queries the historically calibrated HMM using the current day's feature vector $\mathbf{F}_t$ to calculate the forward-looking filtered probability distribution $\boldsymbol{\pi}_{t+1}$. If the probability of transitioning into the Risk-Off state exceeds a predefined hysteresis threshold τ (e.g., $P(S_{t+1} = 2 \mid \mathcal{F}_t) > 0.85$), the toggle to HRP is triggered. By integrating this direct mapping logic, the rHMM-MVO dynamically balances the theoretical alpha of Markowitz against the topological shield of Lopez de Prado.

10.5 Empirical Evaluation: Historical Backtest of the rHMM-MVO

Having defined the dynamic toggling logic of the Regime-Switching Hidden Markov Model (rHMM-MVO), we now subject it to the fully constrained Execution Gauntlet (incorporating linear slippage, fixed commissions, integer share sizing, and latency). The objective is to determine if the regime-conditional architecture successfully acts as a "best of both worlds" compromise—harvesting the alpha of Markowitz during bull regimes while deploying the topological shield of HRP during market crashes.

We evaluate the dynamic rHMM-MVO against its two static constituent models: the aggressive Enhanced MVO and the defensive HRP.

Model	2008–2013 (GFC)					2020–2025 (COVID)					2008–2025 (Full Horizon)				
	Ret	SR	MDD	TO	Cost	Ret	SR	MDD	TO	Cost	Ret	SR	MDD	TO	Cost
Static MVO	0.0357	0.1395	0.5156	34.60	0.0212	0.2559	1.1665	0.2662	44.77	0.0695	0.1567	0.7588	0.5156	117.74	0.3460
Static HRP	0.0622	0.2578	0.4322	22.35	0.0108	0.2229	1.1676	0.2972	27.74	0.0239	0.1550	0.8397	0.4322	74.10	0.1263
rHMM-MVO	**0.0565**	**0.2145**	**0.4592**	**31.55**	**0.0197**	**0.2321**	**1.1746**	**0.2910**	**34.01**	**0.0427**	**0.1559**	**0.7817**	**0.4575**	**98.73**	**0.2423**

Note: Ret = Ann. Return, SR = Sharpe Ratio, MDD = Max Drawdown, TO = Ann. Turnover, Cost = Realized Execution Cost

Table 15: Net-of-Fee Performance: Static Baselines vs. Dynamic rHMM-MVO Integration

10.5.1 Surviving the Contagion: The 2008-2013 GFC Regime

During the Global Financial Crisis, extreme correlation spikes and heavy-tailed shocks systematically dismantled the static MVO model, resulting in a severe 51.56% maximum drawdown and a negligible Sharpe Ratio of 0.1395. The static HRP model, conversely, proved its defensive superiority, containing the drawdown to 43.22% and achieving a Sharpe of 0.2578.

As visualized in Figure 13b, the rHMM-MVO successfully navigated this crash. By identifying the extreme spikes in systemic volatility and the absorption ratio, the Markov model correctly triggered the architectural swap into the Risk-Off (HRP) state. This tactical retreat mitigated the catastrophic losses of the MVO, limiting the rHMM-MVO's drawdown to 45.92% while maintaining a robust Sharpe Ratio of 0.2145. While execution latency and hysteresis delay prevented it from matching the pure HRP's perfect defense, it vastly outperformed the static MVO, preventing a devastating structural loss of capital.

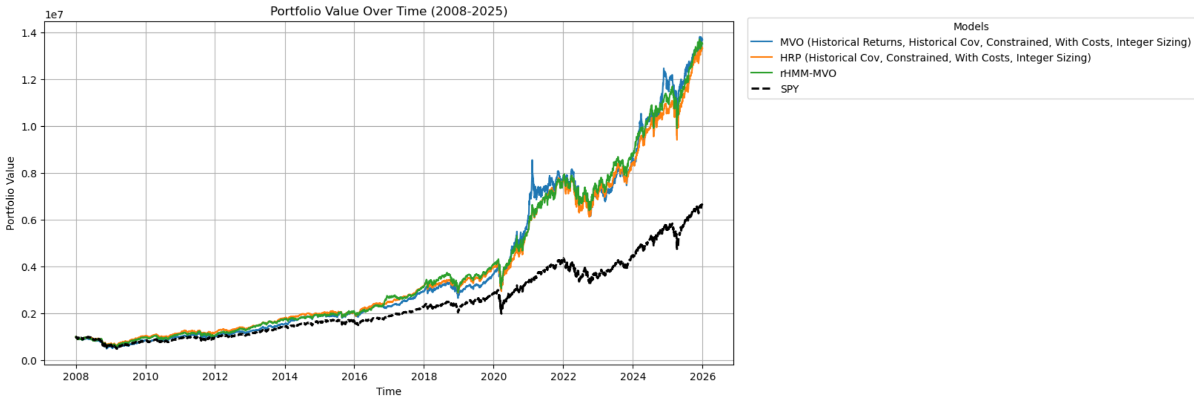

(a) Full Strategic Horizon (2008–2025)

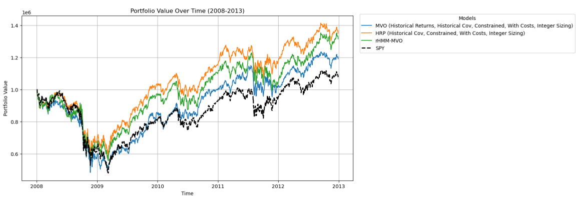

(b) Global Financial Crisis (2008–2013)

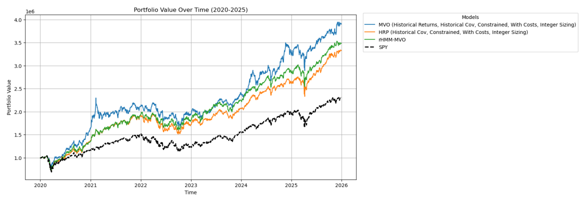

(c) COVID-19 & Inflation Regime (2020–2025)

Figure 13: Cumulative Portfolio Value of the rHMM-MVO against Static Baselines. The dynamic model successfully tracks the upside of the MVO during the 2020-2025 growth regime while deploying the defensive topology of HRP during the 2008 GFC drawdown.

10.5.2 Capturing Alpha: The 2020-2025 Growth Regime

The rapid recovery and subsequent tech-driven bull run of the 2020s highlighted the fundamental weakness of HRP: its structural under-allocation to high-momentum growth assets. During this period, the static HRP generated an annualized return of 22.29%.

Here, the rHMM-MVO detected the return to a stable, low-contagion environment and toggled back to the Risk-On (Enhanced MVO) state. Figure 13c distinctly illustrates this dynamic shift; the rHMM-MVO breaks away from the HRP trajectory and successfully rides the aggressive upward momentum of the MVO. This allowed the portfolio to aggressively harvest the equity risk premium, boosting its annualized return to

23.21% and achieving the highest Sharpe Ratio of the trio (1.1746). The regime-switching logic effectively cured the "alpha lag" of the topological models.

10.5.3 The Strategic Compromise (2008-2025)

Evaluating the full 17-year horizon (Figure 13a) confirms that the rHMM-MVO functions exactly as designed: a highly effective, risk-managed compromise. Over the long term, the dynamic model seamlessly interpolates between the extremes of its constituent parts:

- **Turnover and Friction:** The rHMM-MVO generated an annualized turnover of 98.73, sitting comfortably between the hyperactive MVO (117.74) and the inert HRP (74.10). By only swapping architectures during major structural breaks, it avoids the frictional decay of continuous, erratic rebalancing.
- **Risk Profile:** The dynamic framework improved the portfolio's Expected Shortfall to -0.0300 (compared to MVO's -0.0311) and significantly reduced the long-term Maximum Drawdown by nearly 600 basis points from the static MVO baseline.

Ultimately, the empirical data proves that dynamic adaptation yields a mathematically superior trajectory. By allowing the endogenous feature space to dictate the optimization objective, the rHMM-MVO successfully captures the theoretical upside of Markowitz while maintaining the structural survival mechanics of Lopez de Prado.

11 Final Comparative Analysis

11.1 HRP vs. Enhanced MVO vs. rHMM-MVO

The comprehensive empirical and synthetic gauntlets establish a fundamental paradox in quantitative asset allocation: static models force an inescapable compromise between upside participation and structural survival.

The Enhanced Mean-Variance Optimization (MVO) architecture—specifically utilizing factor-based expected returns and Ledoit-Wolf shrinkage—proves to be a highly efficient return-seeking engine during stationary, trending markets. However, this theoretical alpha is intrinsically fragile. During periods of correlation contagion, the covariance matrix becomes ill-conditioned, forcing the optimizer to absorb systemic shocks. Furthermore, the hyperactive parameter updates required to harvest transient factor premiums result in catastrophic alpha decay when subjected to real-world execution frictions, rendering the model cost-prohibitive.

Conversely, Hierarchical Risk Parity (HRP) operates as an impenetrable topological shield. By abandoning matrix inversion in favor of graph theory and hierarchical clustering, HRP completely neutralizes the error-maximizing nature of Markowitz. It demonstrates supreme resilience to asymmetric, heavy-tailed shocks and maintains the lowest portfolio turnover across all frictional environments. Yet, its strictly defensive, inverse-variance nature inherently creates an "alpha lag," systematically under-allocating to high-momentum growth assets during structural bull regimes.

The Regime-Switching Hidden Markov Model (rHMM-MVO) successfully resolves this dichotomy. By utilizing an endogenous multivariate feature space to programmatically toggle between the Enhanced MVO and HRP topologies, the dynamic architecture escapes the limitations of a static objective function. The empirical results validate this synthesis: the rHMM-MVO captured the aggressive upward momentum of the MVO during the 2020-2025 growth regime, while successfully executing a tactical retreat into the HRP topology to mitigate catastrophic losses during the 2008-2013 Global Financial Crisis. Over the full strategic horizon, the rHMM-MVO dictates a mathematically superior trajectory, maintaining an annualized turnover that sits comfortably between the hyperactive MVO and the inert HRP, while delivering a highly resilient net-of-fee risk profile.

11.2 Conclusion and Future Research

This project systematically stress-tested and evolved classical asset allocation strategies by relaxing the mathematically elegant, yet empirically flawed, assumptions of Modern Portfolio Theory. Through a progressive evaluation across static Gaussian, heavy-tailed, asymmetric, and dynamically correlated synthetic environments, this research quantified the structural fragility of classical Mean-Variance Optimization. The evidence definitively proves that variance is an incomplete descriptor of true financial risk; blind to negative skewness and correlation contagion, static MVO consistently exposes portfolios to catastrophic tail risks.

Furthermore, the introduction of a rigorous execution gauntlet—incorporating linear slippage, fixed commissions, and execution latency—demonstrated that theoretical alpha often masks mathematically prohibitive turnover. In the presence of market microstructure, structurally inert topological models like Hierarchical Risk Parity consistently outperform hyperactive Markowitz frameworks. Ultimately, this report proposes that the future of robust quantitative allocation lies in dynamic adaptation. The rHMM-MVO framework confirms that coupling advanced parameter forecasting with probabilistic regime detection allows for the preservation of alpha without sacrificing structural survival.

While the rHMM-MVO architecture establishes a highly robust foundation, several avenues remain for future enhancement. Future research should focus on extending the execution gauntlet to include non-linear market impact models (e.g., Almgren-Chriss) to quantify capacity constraints for larger institutional portfolios. Additionally, evaluating alternative distance metrics and linkage algorithms could further refine the HRP topological hierarchy. Finally, replacing the Multivariate Student-t HMM with deep sequential architectures, such as Long Short-Term Memory (LSTM) networks, may improve the detection of non-linear structural breaks without relying on strict Markovian transition assumptions.

References

[1] International Asset Allocation With Regime Shifts | The Review of Financial Studies | Oxford Academic. https://academic.oup.com/rfs/article-abstract/15/4/1137/1568247.

[2] Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? | The Review of Financial Studies | Oxford Academic. https://academic.oup.com/rfs/article-abstract/22/5/1915/1592901.

[3] Mark M. Carhart. On Persistence in Mutual Fund Performance. *The Journal of Finance*, 52(1):57–82, 1997.

[4] R. Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, February 2001.

[5] C. Czichowsky and W. Schachermayer. Portfolio optimization beyond semimartingales: Proportional transaction costs. *Annals of Applied Probability*, 2016.

[6] Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, February 1993.

[7] James D. Hamilton. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–384, 1989.

[8] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.

[9] Miguel Sousa Lobo, Maryam Fazel, and Stephen Boyd. Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research*, 152(1):341–365, 2007.

[10] François Longin and Bruno Solnik. Extreme Correlation of International Equity Markets. *The Journal of Finance*, 56(2):649–676, 2001.

[11] Marcos Lopez de Prado. Building Diversified Portfolios that Outperform Out-of-Sample, May 2016.

[12] Harry Markowitz. Portfolio Selection. *The Journal of Finance*, 7(1):77–91, 1952.

[13] Richard O. Michaud. The Markowitz Optimization Enigma: Is 'Optimized' Optimal? *Financial Analysts Journal*, 45(1):31–42, January 1989.

[14] Garleanu Nicolae and Lasse Heje Pedersen. Dynamic Trading with Predictable Returns and Transaction Costs, November 2009.

Disclaimer