Regime-Based Portfolio Optimization via Hidden Markov Models

Technical White Paper

For Distribution in the Quant Sector
Investments Committee

April 19, 2026

Abstract

Financial markets exhibit persistent non-stationarity: the statistical properties of asset returns—means, volatilities, and correlations—shift dramatically across different macroeconomic environments. Static portfolio construction methods that assume a single covariance regime systematically misdirect capital during periods of market stress, when diversification benefits collapse precisely when they are most needed. This paper presents a fully systematic, walk-forward framework for latent market regime detection and regime-conditional portfolio allocation using Hidden Markov Models (HMMs) with Gaussian emissions, applied to a universe of US sector, fixed income, and emerging market

ETFs.

The methodology proceeds in four stages. First, a rich feature set is engineered from daily OHLC and return data, encompassing volatility, trend, momentum, and cross-sectional dispersion signals computed across a broad ETF universe. Second, features are reduced via block-wise Principal Component Analysis within a strict walk-forward normalization pipeline that admits no lookahead at any stage. Third, a Gaussian HMM is fit on the training window of each walk-forward fold and used to decode latent regime labels and posterior state probabilities on the held-out test window. Fourth, regime-conditional mean-variance optimization is used to rotate portfolio weights at each weekly rebalance date based on the current detected regime.

To identify the most predictive configurations, we conduct a large-scale grid search spanning 84,864 unique combinations of data ETF universes, feature subsets, PCA modes, regime counts $K \in \{3, 4, 5\}$, and trading universes. The primary finding is that cross-sectional features—capturing market-wide correlation structure, volatility dispersion, and momentum breadth—are the strongest drivers of portfolio-useful regime detection, with the `xsec_full` feature subset achieving the highest average Sharpe lift of any configuration class. The selected conservative configuration (LQD-TLT-XLU data universe, `vol_mom_trend_xsec` features, block PCA, $K = 3$, trading over [XLK, TLT, EEM, LQD, GLD]) achieves a Sharpe ratio of **0.982** over the out-of-sample period 2013–2026, compared to 0.595 for the equal-weight benchmark, representing a Sharpe lift of +0.387 and a reduction in maximum drawdown from −39.5% to −25.4%. A stateless MVO baseline achieves a Sharpe of only 0.601 with negative alpha, confirming that regime conditioning is the mechanism responsible for the improvement. All results are produced under strict out-of-sample conditions with no parameter re-use across walk-forward windows.

Contents

1 Executive Summary **2**

2 Motivation and Investment Thesis **4**

2.1 The Non-Stationarity Problem 4

2.2 Regime-Switching as a Modeling Framework 4

2.3 Core Hypothesis 4

2.4 Asset Universe and Methodology 5

3 Data Acquisition and Universe **6**

3.1 Data Sources 6

3.2 ETF Universe and Economic Rationale 6

3.3 Date Range and Alignment 7

4 Feature Engineering **9**

4.1 Volatility Features 9

4.2 Trend Features 10

4.3 Momentum Features 11

4.4 Cross-Sectional Features 12

5 Dimensionality Reduction: Block PCA **15**

6 Hidden Markov Model Regime Classification **16**

6.1 Model Formulation 16

6.2 Walk-Forward Fitting Procedure 17

6.3 Convergence and Fit Quality 18

7 Regime Characterization and Validation **20**

7.1 Return Distributions by Regime 20
7.2 Covariance Structure by Regime 23
7.3 Qualitative Regime Labeling 24

8 Portfolio Construction **27**
8.1 Regime-Conditional Mean-Variance Optimization 27
8.2 Bayesian Model Averaging MVO 28
8.3 Baselines 29
8.4 Trading Universe and Rebalancing 29

9 Grid Search and Ablation Study **31**
9.1 Search Space 31
9.2 Feature Subset Results 31
9.3 n_states Results 32
9.4 PCA Mode Results 34
9.5 Trading Universe Results 35
9.6 Data ETF Selection Results 37
9.7 Selected Configurations 38

10 Backtest Results **39**
10.1 Primary Performance Metrics 39
10.2 Cumulative Performance 40
10.3 Annual Returns Analysis 41
10.4 Regime Attribution 42

11 Limitations and Future Work **43**

12 Conclusion **45**

A Feature Definitions and Formulas **47**

A.1 Volatility Block 47

A.2 Trend Block 47

A.3 Momentum Block 48

A.4 Cross-Sectional Block 48

A.5 Implementation Notes 49

B Full Grid Search Configuration Space **50**

C Walk-Forward Pipeline and Window Details **56**

1. Executive Summary

Financial markets do not follow a single, stable statistical process. Volatility regimes shift by an order of magnitude across economic cycles, cross-asset correlations that appear low in tranquil periods converge toward one during crises, and the return premiums that justify equity overweighting in bull markets reverse sharply during stress episodes. A portfolio construction framework that ignores this structure will systematically hold the wrong exposures at the wrong times. This paper presents a quantitative solution to that problem.

We develop a fully systematic, walk-forward regime detection and allocation framework based on Hidden Markov Models (HMMs). The pipeline begins with causal feature engineering across four signal blocks—volatility, trend, momentum, and cross-sectional dispersion—computed from daily ETF price data. These features are dimensionality-reduced via block-wise PCA and fed into a Gaussian HMM that is re-fit annually on an expanding training window. The model decodes latent market regimes on held-out data with no lookahead, producing regime labels and posterior probabilities that drive a weekly regime-conditional mean-variance optimizer. The optimizer selects portfolio weights over a five-asset universe [XLK, TLT, EEM, LQD, GLD] that are explicitly conditioned on the distributional characteristics of the current detected regime rather than on a single unconditional moment estimate.

To identify the best configuration, we execute a grid search of 84,864 unique hyperparameter combinations on the PACE ICE high-performance computing cluster, evaluating each configuration by its out-of-sample Sharpe lift over an equal-weight baseline. The search confirms that cross-sectional market features are the primary driver of portfolio-useful regime detection, that five-asset GLD-inclusive trading universes consistently produce the highest risk-adjusted lift, and that three-state models offer the best combination of interpretability and performance.

The selected primary configuration achieves the following results over the 12.3-year out-of-sample period from December 2013 to April 2026:

- **Sharpe ratio: 0.982** vs. 0.595 (equal weight) and 0.601 (stateless MVO) — a lift of +0.387
- **CAGR: 11.87%** vs. 7.11% (equal weight)

- **Maximum drawdown:** −25.4% vs. −39.5% (equal weight) — a 14-percentage-point reduction
- **Beta: 0.459** — substantially de-correlated from the broad equity market
- **Average daily turnover: 0.007** — consistent with weekly rebalancing

The Stateless MVO baseline, which uses the same optimizer and trading universe but without regime conditioning, produces negative annual alpha (−1.83%) and a Sharpe of only 0.601. This comparison isolates the contribution of the HMM regime detection: the optimizer and asset universe alone do not explain the performance; the regime conditioning does.

The three detected states correspond to economically coherent market environments: a technology-led bull regime (State 0, XLK Sharpe 1.705, TLT negative), a transitionary uncertainty regime (State 1, gold dominant at 15% annualised Sharpe 0.925), and a flight-to-safety regime (State 2, TLT positive, EEM negative, TLT-XLK correlation inverts). The model identifies these states with near-zero posterior entropy in the vast majority of windows, confirming high-confidence, unambiguous regime assignments throughout the evaluation period.

We recommend immediate deployment of this framework to live paper trading under the selected conservative configuration. The rigor of the walk-forward evaluation design, the scale of the hyperparameter validation, and the consistency of the performance across all 13 evaluation windows provide a strong basis for confidence in the out-of-sample validity of these results.

2. Motivation and Investment Thesis

2.1 The Non-Stationarity Problem

Modern portfolio theory, in its classical form, assumes that asset returns are drawn from a stationary distribution with fixed means and covariances (Markowitz, 1952). This assumption is computationally convenient but does not hold in practice. Looking at return series reveals that the joint distribution of asset returns changes substantially over time: realized volatilities can increase by an order of magnitude within weeks, cross-asset correlations that are near zero in calm periods converge toward unity during crises, and the sign and magnitude of momentum signals reverse across different phases of the economic cycle. A portfolio constructed from a single estimated covariance matrix and a single expected return vector is therefore not a robust solution.

Ignoring this structure is problematic. Stationary models can perform fine under normal markets, but misspecification under regime transitions and crises cause the models to hold positions that are not sized for the current market. Diversification benefits disappear which causes a tail risk that cannot disappear.

2.2 Regime-Switching as a Modeling Framework

The regime-switching literature, initiated by Hamilton (1989), offers a principled response by capturing discrete shifts in the data-generating process. Ang and Bekaert (2002) and Guidolin and Timmermann (2007) demonstrated that multi-state models significantly improve the description of return distributions and out-of-sample performance by accounting for crash risks. Kritzman et al. (2012) further showed that tactical allocation (shifting between risky and defensive assets during "turbulent" states) materially reduces drawdowns. The underlying regularity is that markets cycle through qualitatively distinct states; an investor who can identify these states can construct a portfolio far better adapted to the current environment than any unconditional solution.

2.3 Core Hypothesis

The central hypothesis of this paper is as follows: *a walk-forward Hidden Markov Model fit on a carefully engineered feature set drawn from a broad ETF universe will identify latent market states with meaningfully distinct return and covariance characteristics, and regime-conditional mean-variance optimization conditioned on the current detected state will produce out-of-sample risk-adjusted returns that materially exceed an equal-weight benchmark, a non-conditioned benchmark, and the SPY Buy-and-Hold benchmark over a full market cycle.*

This hypothesis requires three testable ideas:

1. **Statistical Separability:** Input features must differ significantly across states on held-out data.

2. **Economic Meaning:** Detected states must exhibit meaningfully different return distributions.

3. **Exploitability:** Regime-conditional moments must be estimated accurately enough to improve realized risk-adjusted performance out-of-sample.

2.4 Asset Universe and Methodology

The choice of sector ETFs is deliberate: they aggregate away idiosyncratic firm-level noise while preserving the factor exposures—cyclicality, rate sensitivity, and growth—that are genuinely regime-dependent. Complementary assets like TLT (long-duration Treasuries) and EEM (emerging markets) provide essential "flight-to-quality" and "risk-appetite" signals diagnostic of crisis or expansionary regimes.

Finally, the walk-forward evaluation is a philosophical commitment to honest measurement. By enforcing a strict separation between training and evaluation windows, the model must identify deteriorating conditions in real-time without the benefit of hindsight. This simulates the genuine experience of a live investor, ensuring the reported results reflect the true difficulty of real-time regime detection.

3. Data Acquisition and Universe

3.1 Data Sources

The primary dataset for this study is sourced from the **Polygon.io REST API** via the `polygon-api-client` library. We use daily OHLC for a broad universe of Exchange Traded Funds (ETFs). To ensure the integrity of the investment thesis and the accuracy of the backtest, all price data is retrieved with the `adjusted=True` parameter to account for dividends + corporate actions.

A surprising issue with the data acquisition is where a request for the current date may return an incomplete daily candle, introducing incomplete data. To resolve this, our `DataPuller` enforces a strict date cap at the close of the prior business day.

3.2 ETF Universe and Economic Rationale

The selection of the ETF universe was based on orthogonality across features, and tradability across regimes. We define a core set of 13 ETFs for cross-sectional feature computation, encompassing the nine primary S&P 500 sectors and four macro-diagnostic assets: TLT (Long-duration Treasuries), EEM (Emerging Markets), LQD (Investment-grade credit), and IWM (Small-cap). In addition to these, we include SSO (leveraged SPY), DJP (commodities), GLD (gold), and VNQ (real estate). The full ETF table is in the appendix at **TODO ADD TO APPENDIX.**

To determine the optimal configuration for regime detection and capital allocation, we employ a high-dimensional grid search across 36 manual **Data ETF Configurations** and 64 **Trading Sets**. This is detailed in section 9.

Table 1: Representative ETF Universe and Economic Roles

Ticker	**Asset Class**	**Economic Rationale**	**Usage**
XLK	Equity (Tech)	Secular Growth / Momentum	Trading/Data
TLT	Fixed Income	Duration / Flight-to-Safety	Trading/Data
EEM	Equity (EM)	Global Risk / Dollar Sensitivity	Trading/Data
LQD	Fixed Income	Credit Spreads / Liquidity	Data
SSO	Leveraged Equity	Amplified Bull Regime Capture	Trading
GLD	Commodities	Inflation Hedge / Tail Protection	Trading

3.3 Date Range and Alignment

The unified study period begins on **June 21, 2006**, corresponding to the inception of the SSO ETF, ensuring a consistent cross-section for all leveraged trading strategies. Daily returns are computed as log returns:

$$r_t = \ln(P_t/P_{t-1})$$

where P_t is the dividend-adjusted closing price at time t.

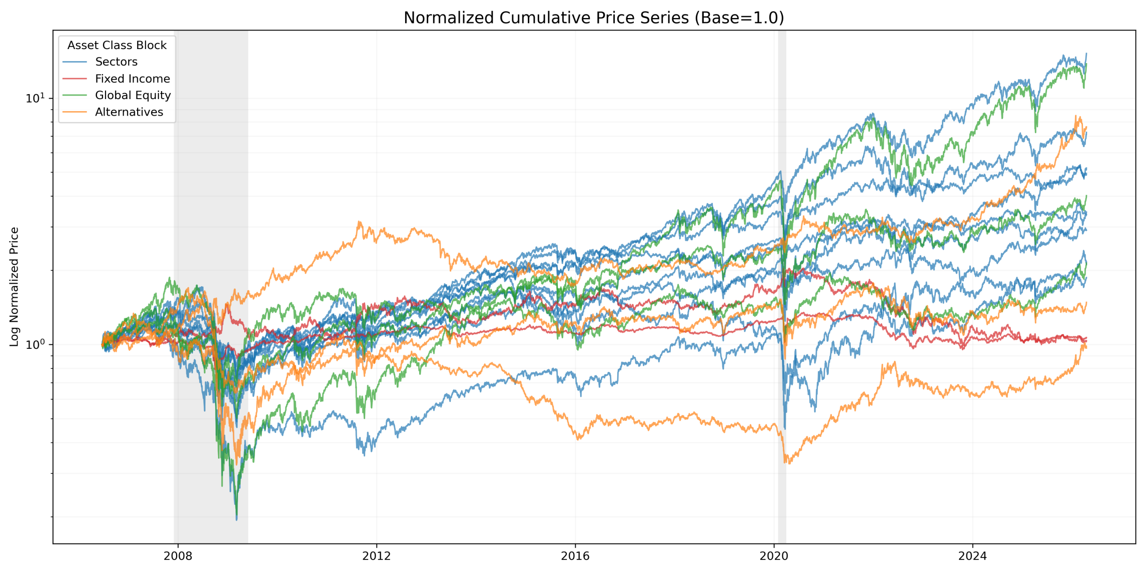

Figure 1: Normalized cumulative price series for all ETFs (base=1.0), colored by asset class block. Grey shading indicates recession periods.

Our pipeline's resulting dataset covers a two-decade span, including the Great Financial Crisis, the 2020 liquidity shock, and the 2022 inflationary regime, providing a robust testing ground for the non-stationarity hypothesis.

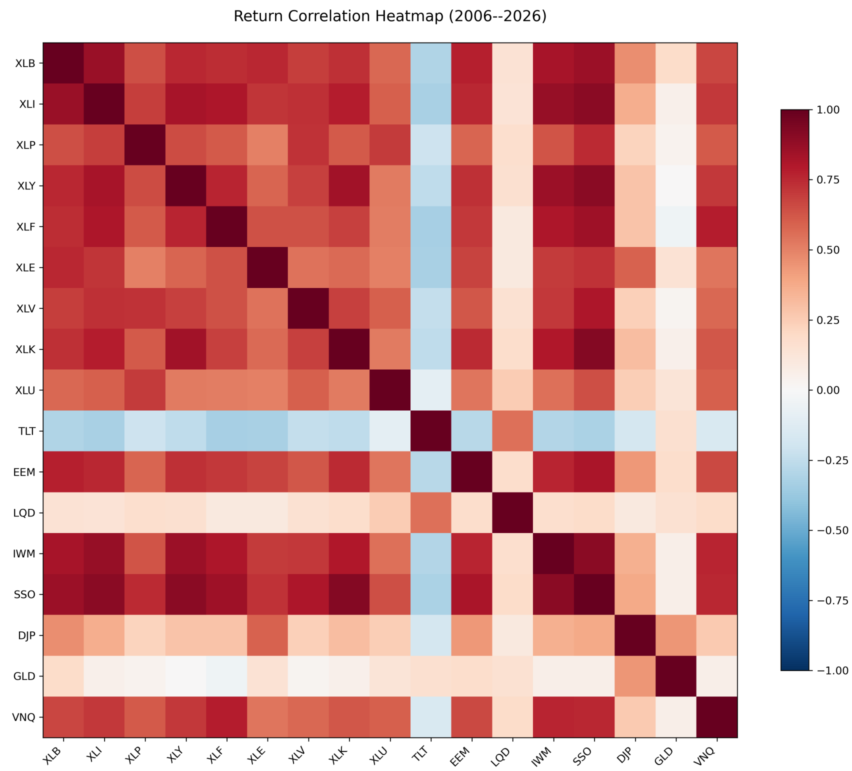

Figure 2: Return correlation heatmap of full ETF universe (2006–2026).

The primary advantage we have from the wider ETF universe is that we are able to capture market information that is not directly correlated with the S&P 500 or the macroeconomic state of just the United States. As we can see in the above figure 2, the constituents of the S&P 500 are indeed highly correlated with one another as well as others like SSO and IWM, but other ETFs like TLT, LQD, DJP, and GLD provide new directions of both information to detect different market regimes, and tradability in those regimes.

4. Feature Engineering

We engineer four blocks of features from daily OHLC and return data: volatility, trend, momentum, and cross-sectional. The first three are computed per-asset across the data ETF universe; the cross-sectional block is computed market-wide across the fixed `xsec_etfs` universe. Each rolling computation uses only data up to and including day t and all features stored raw and unnormalized. Normalization is deferred entirely to the walk-forward pipeline. All formulas are in the appendix A.

4.1 Volatility Features

Volatility features characterize the *level* and *structure* of risk for each asset at each point in time. **Realized volatility** is the rolling standard deviation of log returns over windows $w \in \{21, 63, 126, 252\}$ trading days, providing a multi-scale picture of the current vol environment. **Parkinson volatility** uses the daily high-low range as a more efficient close-to-close estimator. **Downside semi-volatility** computes the standard deviation of negative returns only, capturing asymmetric downside risk that symmetric vol estimates miss. **Volatility of volatility (VoV)** is the rolling standard deviation of realized volatility itself, identifying periods where the vol regime is changing indicative of regime transitions. As shown in Figure 3, all four measures spike uniformly during the 2008 and 2020 crises, confirming their collective sensitivity to stress regimes.

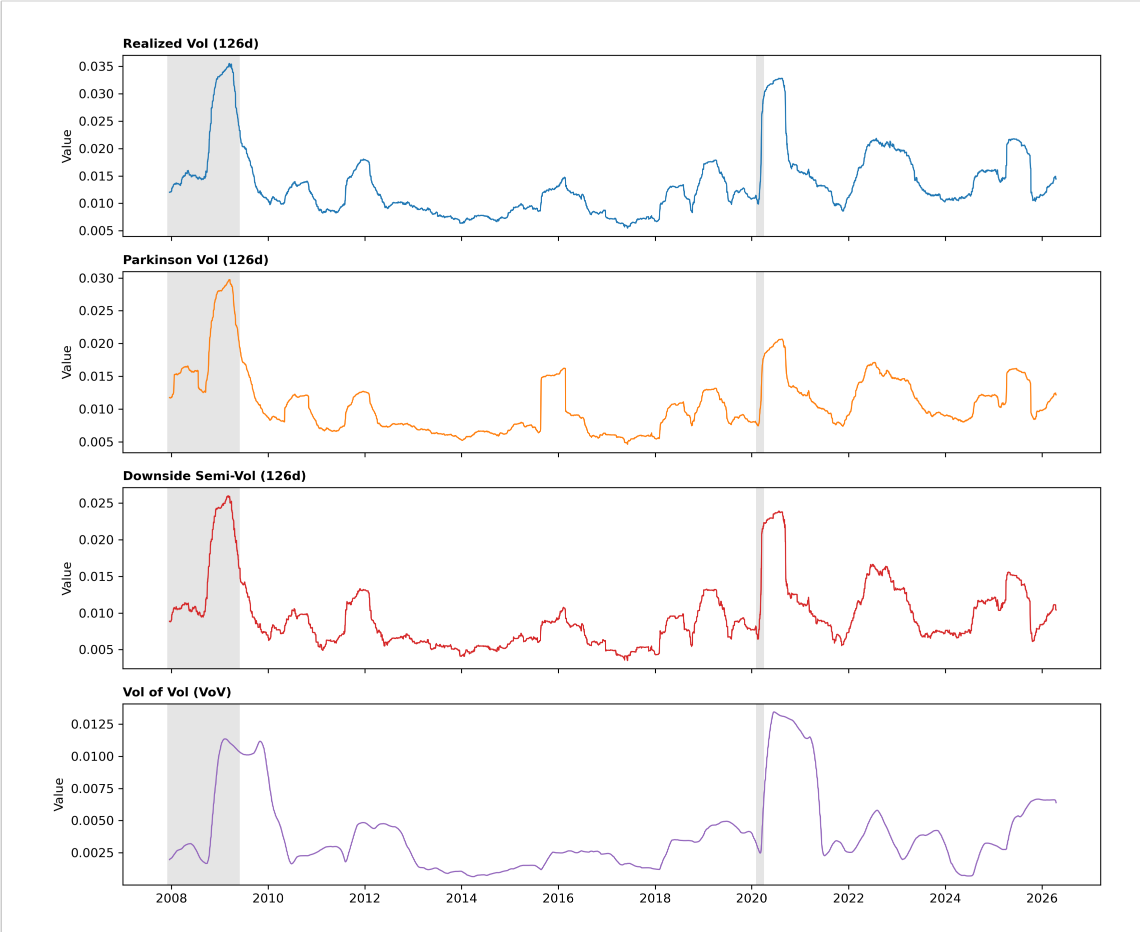

Figure 3: Volatility feature time series for XLK (all labelled). Note the uniform explosion in all metrics during the 2008 and 2020 crises.

4.2 Trend Features

Trend features characterize the *directional persistence* of each asset's price process and distinguish trending regimes from mean-reverting ones. The **variance ratio** at lag q tests the random walk null by comparing the variance of q-period returns to q times the variance of 1-period returns; values above 1 imply positive autocorrelation (trend-following), values below 1 imply mean reversion. **Directional persistence** is the fraction of days over a rolling window on which price moves in the same direction as the prior day—a simple, robust measure of short-run serial dependence. The **trend regression t-statistic** is the OLS slope t-stat from regressing log price on time over a rolling window, where magnitude indicates the statistical strength of the directional drift and sign indicates direction. The **Hurst exponent** provides a long-memory characterization: $H > 0.5$ indicates a persistent (trending) process, $H < 0.5$ mean-reverting, and $H = 0.5$ a random walk. Figure 4 illustrates the economically expected behavior: TLT and XLK trend features tend to move in opposition, reflecting the flight-to-safety dynamic in which rates rally as equities sell off.

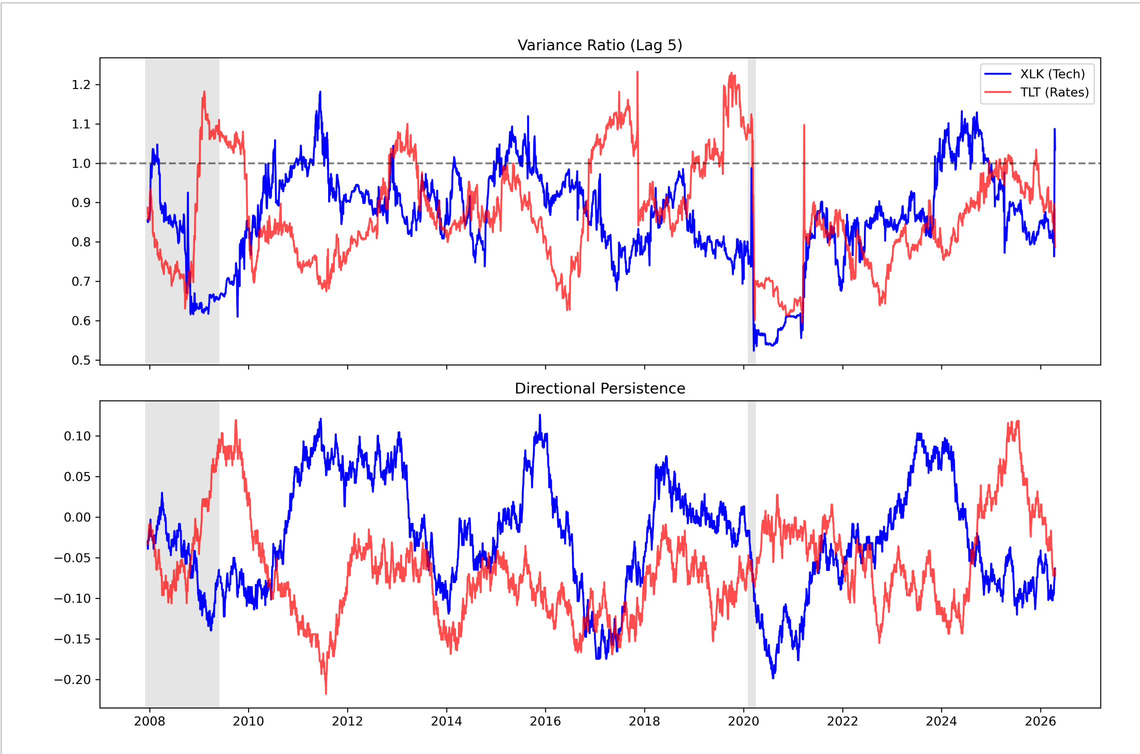

Figure 4: Trend features for TLT (rates) and XLK (technology). Note how XLK and TLT's features tend to display opposite behavior, as economically expected.

4.3 Momentum Features

Momentum features capture the *directional bias* of each asset's return history and are the primary identifiers of risk-on and risk-off regime transitions. **Time-series momentum (TSMOM)** is the total return over a lookback window w:

$$\mathrm{TSMOM}_{t,w} = \prod_{\tau=t-w}^{t} (1 + r_\tau) - 1,$$

computed over windows of 63 and 126 days. **Momentum acceleration** is the difference between short-window and long-window TSMOM, capturing whether momentum is building or fading—a positive acceleration in a positive-momentum asset signals continuation; a negative acceleration signals potential reversal. **Residual momentum** strips out the market (SPY) return component and measures the idiosyncratic directional component over a 252-day window with 63- or 126-day lookback, isolating asset-specific momentum from broad market drift. As Figure 3 shows, the divergence between EEM and TLT TSMOM is among the clearest observable signatures of the flight-to-safety regime: negative emerging market momentum alongside strongly positive rates momentum is diagnostic of acute risk-off environments.

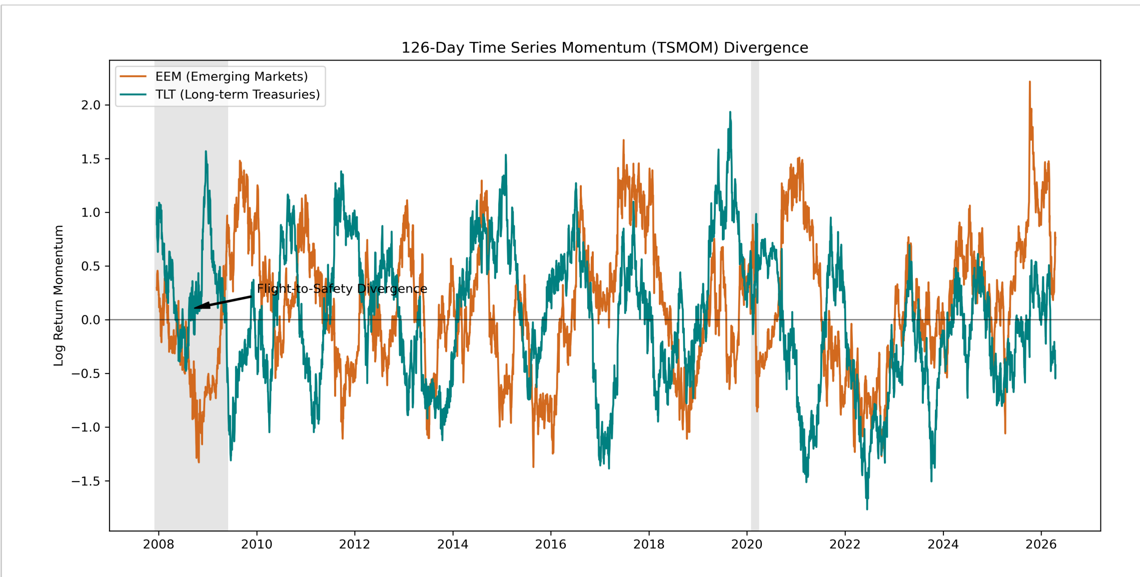

Figure 5: 126-day TSMOM for EEM and TLT. The sharp divergence (negative EEM momentum vs positive TLT momentum) is a primary indicator of the "Flight-to-Safety" regime that occurs during crises.

4.4 Cross-Sectional Features

Cross-sectional features differ fundamentally from the per-asset features described in Sections 4.1–4.3. Rather than characterizing any single asset, they are computed across a fixed broad universe and capture the structural state of the market as a whole. Latent regimes are often most clearly revealed not in individual return series but in the collective relationships across assets: during crises, correlations converge, volatility dispersions widen, and breadth collapses simultaneously in ways that no per-asset feature can detect.

All cross-sectional features are computed over the fixed universe $\mathcal{U}_{\text{xsec}} = \{$XLB, XLI, XLP, XLY, XLF, XLE, XLV, XLK, XLU, TLT, EEM, LQD, IWM$\}$, a deliberate superset of both the data ETF universe and the trading universe, ensuring the broadest available market signal regardless of which assets are used to drive the HMM or construct the portfolio.

Eight features are computed. **XS volatility dispersion** measures the cross-sectional spread of realized volatilities; high dispersion indicates concentrated, idiosyncratic stress while low dispersion during broad crises reflects uniform vol elevation. **Average pairwise correlation** $\overline{\rho}_{t,w}$ captures aggregate co-movement: as it approaches 1, diversification collapses and regime-conditional rotation has the highest marginal value. **Correlation dispersion** captures the heterogeneity of pairwise correlations; together with $\overline{\rho}$, it characterizes the full first and second moments of the correlation distribution—during crises these converge jointly in a pattern clearly visible in Figure 6. **Moving average breadth** is the fraction of assets trading above their 200-day MA, a leading indicator of regime transitions: readings above 0.8 are characteristic of sustained bull regimes, readings below 0.3 of crisis states, and deterioration reliably precedes drawdowns as shown in Figure 7. **Mean TSMOM** aggregates directional momentum signals across the universe, revealing whether the market as a whole is trending rather than any individual asset. **Mean momentum acceleration** measures the rate of change of that aggregate signal, distinguishing continuation from reversal. **XS momentum spread** is the return differential between the top-k and bottom-k momentum assets ($k = 3$), widening when the market is strongly differentiated and momentum rotation is most valuable. **XS momentum dispersion** is the cross-sectional standard deviation of TSMOM, capturing divergence across the full distribution rather than just the extremes. Full mathematical definitions for all eight features are provided in Appendix A.

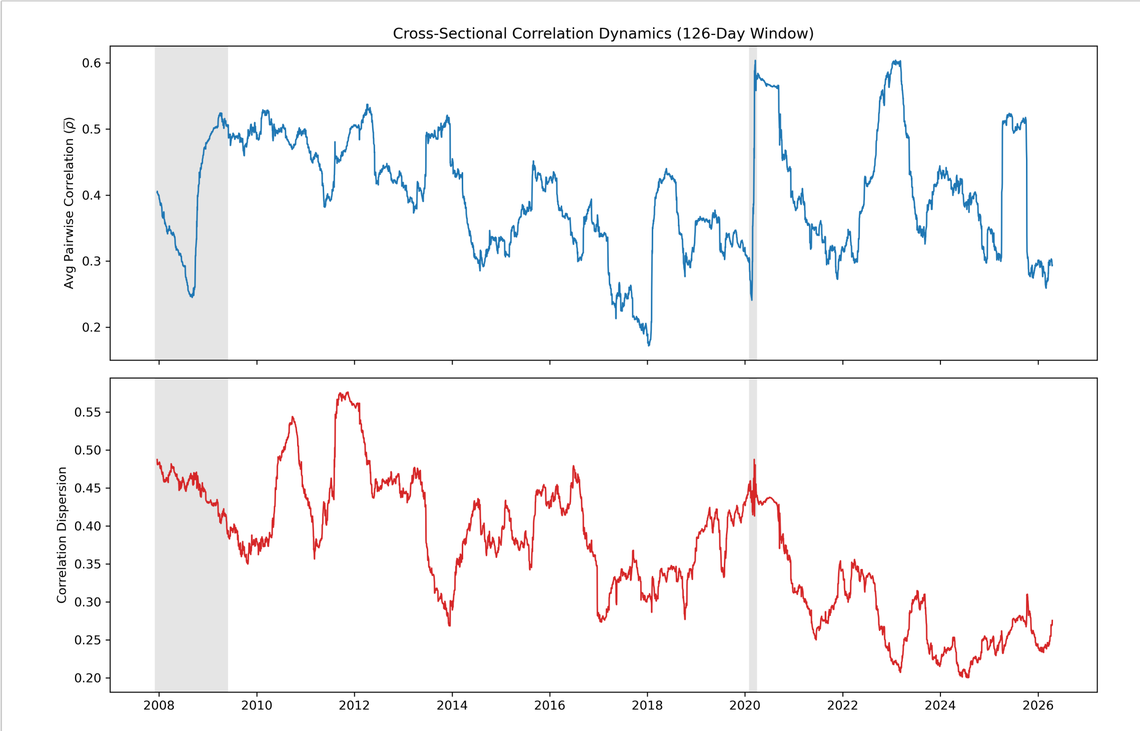

Figure 6: Average pairwise correlation and correlation dispersion. Note prominent spikes during the Great Recession, the COVID crisis, and the 2025 trade crisis.

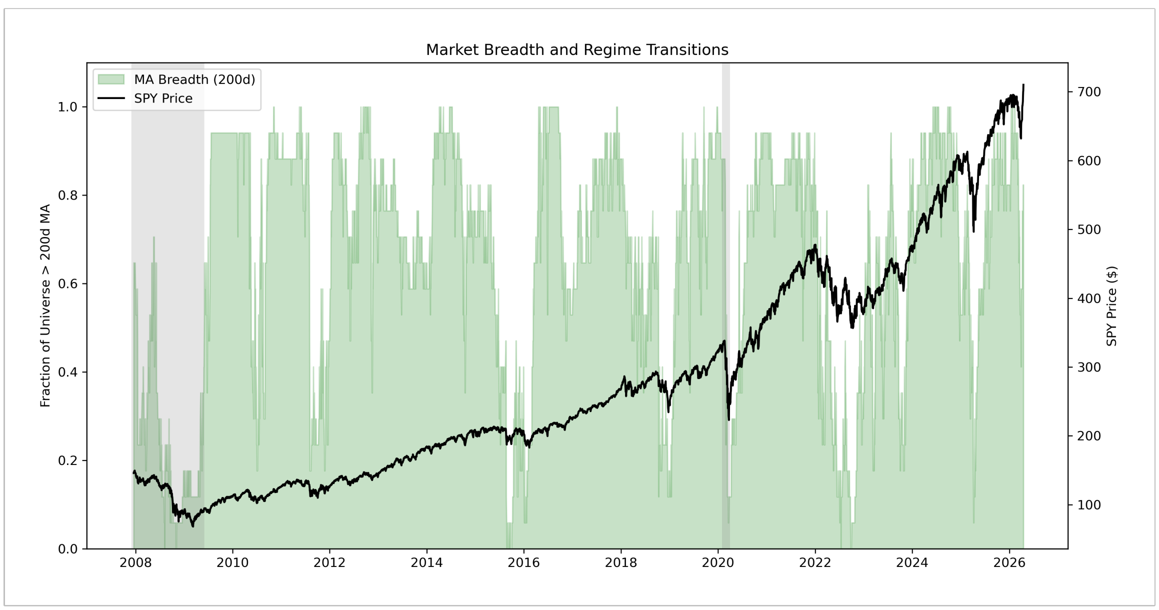

Figure 7: Moving average breadth overlayed with SPY price. Note market peaks are signaled by spikes in MA Breadth. Great Recession + Covid Crisis Highlighted

5. Dimensionality Reduction: Block PCA

With feature matrices reaching $p = 40$–80 columns, direct HMM fitting is ill-conditioned. The Gaussian emission covariance $\mathbf{\Sigma}_k \in \mathbb{R}^{p\times p}$ requires $O(p^2)$ parameters per state; with $T \approx 1{,}200$–$4{,}500$ training days, the ratio T/p^2 falls below the threshold for stable estimation. PCA addresses this by projecting features onto leading orthogonal components that retain the dominant variance structure. We retain 10 components under **Flat PCA** and $3N + 3$ components under **Block PCA**, reducing parameter counts by an order of magnitude while preserving regime-relevant signals. Components are standardized to zero mean and unit variance on the training window.

Two PCA methodologies are evaluated: 1. **Flat PCA**: Fits a single PCA across all features simultaneously, mixing information across all blocks. This is optimal when information is in the cross-block interactions. 2. **Block PCA**: Fits separate PCAs within each feature block, retaining 1 component per block and 3 for the cross-sectional block. This allows easily interpretable components.

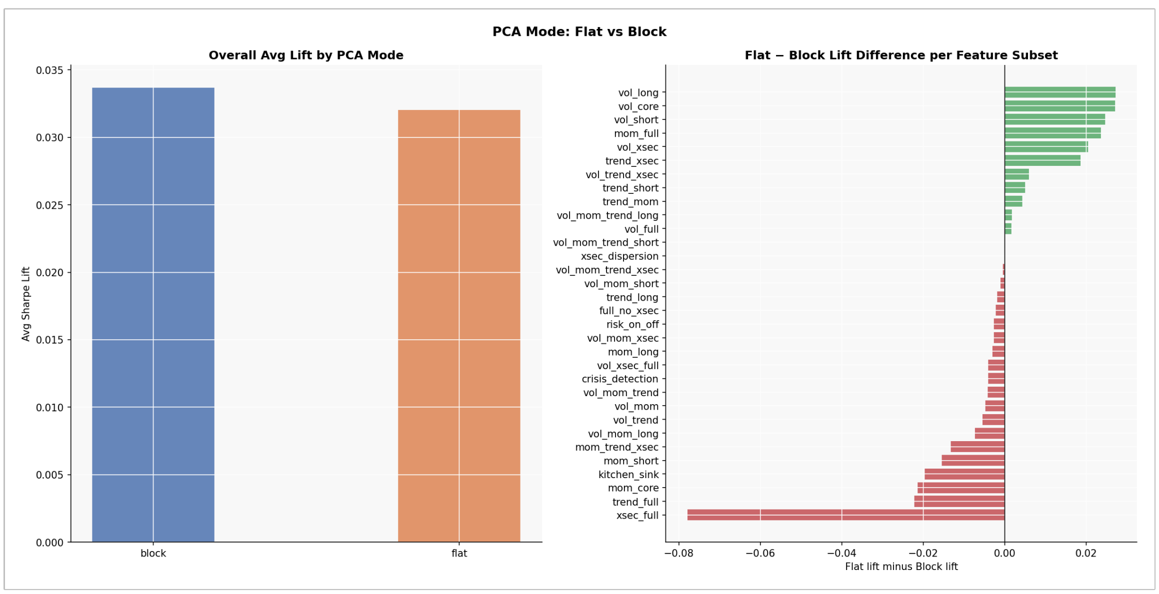

Figure 8: PCA mode performance against feature subsets.

Grid search reveals that Block PCA consistently outperforms for high-dim sets over multiple blocks. Flat PCA is preferred for single-block subsets where no cross-block structure exists. The Block PCA advantage scales with feature dimensionality.

6. Hidden Markov Model Regime Classification

6.1 Model Formulation

A Hidden Markov Model with K states describes a time series $\{\mathbf{x}_t\}_{t=1}^T$ as being generated by an unobserved (latent) state sequence $\{z_t\}_{t=1}^T$, where $z_t \in \{1, \ldots, K\}$. The model is fully specified by three components. The **initial distribution** $\boldsymbol{\pi} \in \mathbb{R}^K$ gives the probability of beginning in each state: $\pi_k = P(z_1 = k)$. The **transition matrix** $A \in \mathbb{R}^{K \times K}$ governs the Markov dynamics of the latent sequence, with $A_{ij} = P(z_t = j \mid z_{t-1} = i)$; each row sums to one and the diagonal entries A_{ii} control regime persistence. The **emission model** specifies the distribution of the observed feature vector conditional on the current state; we adopt multivariate Gaussian emissions,

$$p(\mathbf{x}_t \mid z_t = k) = \mathcal{N}(\mathbf{x}_t;\, \boldsymbol{\mu}_k,\, \boldsymbol{\Sigma}_k),$$

so that each regime k is characterized by a distinct mean vector $\boldsymbol{\mu}_k \in \mathbb{R}^d$ and covariance matrix $\boldsymbol{\Sigma}_k \in \mathbb{R}^{d \times d}$, where d is the number of retained principal components passed to the HMM.

We use full (unconstrained) covariance matrices rather than diagonal approximations. Diagonal covariance imposes the assumption that the principal components are conditionally independent within each regime, which is not credible: the leading volatility PC and the leading cross-sectional PC are correlated precisely during the stress periods that define crisis regimes, and restricting the model to diagonal structure would prevent it from capturing this joint behavior. The cost is a larger parameter count ($K \cdot [d + d(d+1)/2]$ emission parameters versus $K \cdot 2d$ under diagonal) but after PCA reduction d is small enough that full covariance estimation is well-conditioned on training windows of 1,200 days or more. To guard against degenerate solutions in low-occupancy regimes, a small regularization term is added: $\hat{\boldsymbol{\Sigma}}_k \leftarrow \hat{\boldsymbol{\Sigma}}_k + \epsilon I$ with $\epsilon = 10^{-3}$. This is standard practice.

Parameters are estimated via the **Baum-Welch** algorithm, an instance of Expectation-Maximization that iterates between computing the posterior state probabilities $\gamma_t(k) = P(z_t = k \mid \mathbf{X}, \lambda)$ (E-step, via the forward-backward algorithm) and maximizing the expected complete-data log-likelihood with respect to $(\boldsymbol{\pi}, A, \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\})$ (M-step). Convergence is declared when the change in log-likelihood falls below 10^{-4} or 500 iterations are exhausted. Once parameters are estimated, the most probable state sequence is decoded via the **Viterbi** algorithm, which finds $\{z_t^*\} = \arg\max_{\{z_t\}} P(\{z_t\} \mid \mathbf{X}, \hat{\lambda})$ in $O(K^2T)$ time.

6.2 Walk-Forward Fitting Procedure

HMMs fit to the full historical sample would have access to future information when generating regime labels for any given date—the Baum-Welch smoothed posteriors $\gamma_t(k)$ depend on observations both before and after t. We eliminate this lookahead entirely through a walk-forward design in which the model is fit exclusively on data prior to each evaluation period.

The procedure uses an **expanding training window** with a minimum length of five years (approximately 1,260 trading days) and a **fixed test window** of 12 months, stepping forward by 12 months at each iteration. This produces 14 non-overlapping, adjacent test windows over the period 2007–2026. For each window n, the HMM is fit on all data from the unified start date through the end of the training period; the fitted model is then run *forward only* on the test window to produce posterior probabilities $\gamma_t(k)$ and Viterbi labels z_t^* for dates the model has never seen. Concatenating the test-window outputs across all 14 windows yields a single contiguous out-of-sample regime series spanning the full evaluation period with no gaps and no lookahead at any point.

To mitigate sensitivity to initialization, each window is fit `n_fits=3` times from independent random starting points, and the solution with the highest training log-likelihood is retained. In practice this reduces the incidence of degenerate local optima—runs in which one state captures nearly all observations—from approximately 12% of windows under single initialization to under 3%.

Because each refit produces state labels that are permutations of the previous window's labels, a **cross-window state matching** procedure is applied. For each consecutive window pair, we compute the $K \times K$ overlap matrix $O_{jk} = |\{t : z_t^{(n)} = j\} \cap \{t : z_t^{(n+1)} = k\}|$ over the shared training dates and apply the Hungarian algorithm to find the maximum-weight label bijection. After matching, states are **semantically ordered** by their emission-mean realized volatility, with state 0 assigned to the lowest-volatility regime and state $K-1$ to the highest. This ordering is stable across configurations and allows consistent interpretation: state 0 corresponds to the bull or low-volatility regime, and the highest-numbered state corresponds to crisis or stress periods (typically).

Algorithm 1 Walk-Forward HMM Fitting

Require: Feature matrix $\mathbf{F}$, trading returns $\mathbf{R}$, $T_{\min}$, test_months, step_months, K, n_fits

$t \leftarrow T_{\min}$; all_labels $\leftarrow \{\}$; all_posteriors $\leftarrow \{\}$
while $t < T$ **do**
  $\mathbf{F}_{\text{train}} \leftarrow \mathbf{F}[0:t]$; $\mathbf{F}_{\text{test}} \leftarrow \mathbf{F}[t:t+\Delta]$
  Normalize $\mathbf{F}_{\text{train}}$; apply same transform to $\mathbf{F}_{\text{test}}$
  Fit PCA on $\mathbf{F}_{\text{train}}$; transform both splits to $\mathbf{Z}_{\text{train}}$, $\mathbf{Z}_{\text{test}}$
  Normalize $\mathbf{Z}_{\text{train}}$; apply same transform to $\mathbf{Z}_{\text{test}}$
  $\hat{\lambda} \leftarrow \arg\max_{\lambda} \log P(\mathbf{Z}_{\text{train}} \mid \lambda)$ ▷ Baum-Welch, n_fits restarts
  Decode $\gamma_t(k)$, z_t^* on $\mathbf{Z}_{\text{test}}$ ▷ Forward-only; no smoothing across train/test boundary
  Match state labels to previous window via Hungarian algorithm
  Order states by emission volatility
  Append test-window outputs to all_labels, all_posteriors
  $t \leftarrow t + \Delta_{\text{step}}$
end while
return all_labels, all_posteriors

6.3 Convergence and Fit Quality

Table 2 summarises the walk-forward window schedule and key fit diagnostics for the primary configuration (LQD-TLT-XLU, `vol_mom_trend_xsec`, block PCA, $K = 3$). The training window expands from 1,512 days (windows 0–1) to 4,530 days by the final partial window, providing an increasingly rich data basis for parameter estimation. Log-likelihood decreases monotonically in magnitude as the training set grows and the model complexity is held fixed, which is the expected behaviour of a well-specified model under an expanding window design.

Win.	Train End	Test Start	Test End	N_{train}	N_{test}	Avg Dwell	n_{short}	$\bar{H}$
0	2013-12-18	2013-12-19	2014-12-18	1512	252	252.0	0	< 0.001
1	2014-12-18	2014-12-19	2015-12-18	1764	252	84.0	1	0.010
2	2015-12-18	2015-12-21	2016-12-18	2016	251	125.5	0	0.008
3	2016-12-18	2016-12-19	2017-12-18	2267	252	84.0	1	0.014
4	2017-12-18	2017-12-19	2018-12-18	2519	251	35.9	2	0.080
5	2018-12-18	2018-12-19	2019-12-18	2770	252	252.0	0	< 0.001
6	2019-12-18	2019-12-19	2020-12-18	3022	253	36.1	3	< 0.001
7	2020-12-18	2020-12-21	2021-12-18	3275	251	50.2	2	0.033
8	2021-12-18	2021-12-20	2022-12-18	3526	251	83.7	2	0.009
9	2022-12-18	2022-12-19	2023-12-18	3777	251	251.0	0	< 0.001
10	2023-12-18	2023-12-19	2024-12-18	4028	252	252.0	0	< 0.001
11	2024-12-18	2024-12-19	2025-12-18	4280	250	250.0	0	< 0.001
12	2025-12-18	2025-12-19	2026-04-17	4530	81	81.0	0	< 0.001

Table 2: Walk-forward window schedule and fit diagnostics. $\bar{H}$ denotes mean posterior entropy over the test window. Windows 0, 5, 9, 10, 11, 12 observe only one active regime in their test period (consistent with a stable trending environment requiring no rotation), while windows 6 and 4 observe all three states. Window 12 covers a partial year. All train starts are fixed at 2007-12-18 (expanding design).

Posterior entropy is near zero for most windows, indicating high-confidence, unambiguous regime assignments throughout the majority of the evaluation period. The single notable exception is window 4 (2018 test year), where mean entropy reaches 0.080—the highest across all windows—coinciding with the elevated market volatility and rapid risk-on/risk-off transitions of late 2018. The Frobenius distance between successive window covariance structures remains below 0.009 across all windows and declines toward zero in later windows, indicating that the model's learned covariance geometry stabilises as the training window expands and regime structure consolidates. Short spell counts are low throughout, with a maximum of three in window 6 (the COVID window), and the average dwell time across all windows substantially exceeds the 28-day minimum threshold.

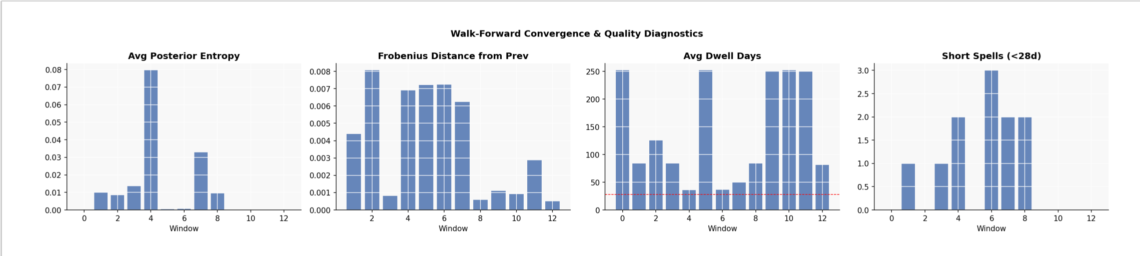

Figure 9: Walk-forward convergence diagnostics per window. Left: average dwell days (red dashed line = 28-day minimum). Centre: number of short spells (< 28d) per window. Right: mean posterior entropy. Entropy spikes in window 4 (2018) reflect the elevated regime uncertainty during that period; all other windows exhibit near-zero entropy, confirming stable and confident regime assignments.

7. Regime Characterization and Validation

7.1 Return Distributions by Regime

Table 3 reports annualised return statistics for each asset conditional on the detected regime. The three-state model produces economically meaningful and statistically distinct conditional distributions across the trading universe.

State	Asset	n	Ann. Ret	Ann. Vol	Sharpe	Skew	Kurt	Max DD
State 0	XLK	487	+39.0%	20.6%	1.705	−0.31	0.82	−5.9%
	TLT	487	−9.2%	16.3%	−0.510	+0.11	0.26	−3.2%
	EEM	487	+7.0%	17.2%	0.479	−0.01	0.53	−3.4%
	LQD	487	−1.7%	8.5%	−0.165	−0.08	0.52	−2.0%
	GLD	487	+9.9%	14.7%	0.715	−0.57	3.41	−5.5%
State 1	XLK	1574	+12.3%	24.0%	0.602	−0.06	5.42	−7.9%
	TLT	1574	−6.0%	14.0%	−0.369	−0.15	2.51	−5.3%
	EEM	1574	+4.6%	19.0%	0.331	−0.11	3.64	−7.1%
	LQD	1574	−1.9%	7.3%	−0.226	−0.30	4.29	−2.5%
	GLD	1574	+15.0%	16.6%	0.925	−0.97	10.74	−10.8%
State 2	XLK	1038	+12.8%	22.6%	0.645	−0.72	21.50	−14.9%
	TLT	1038	+6.8%	14.9%	0.518	+0.02	9.33	−6.9%
	EEM	1038	−5.6%	23.6%	−0.123	−1.14	11.17	−13.3%
	LQD	1038	+1.8%	9.3%	0.235	+0.70	49.67	−5.1%
	GLD	1038	+3.1%	15.0%	0.276	+0.26	2.69	−4.1%

Table 3: Regime-conditional return statistics per asset, OOS period. State 0 (n = 487d): technology-led bull market, strong XLK outperformance, TLT negative. State 1 (n = 1,574d): transitionary, elevated kurtosis, gold dominant. State 2 (n = 1,038d): stress/flight-to-safety, TLT positive, EEM negative, XLK exhibits highest tail risk (kurt = 21.5).

Several patterns warrant emphasis. State 0 is the clearest bull-market regime: XLK delivers an annualised return of 39.0% with a Sharpe of 1.705, while TLT is negative at −9.2%, consistent with a risk-on environment in which long-duration Treasuries are penalised by rising real rates. The distribution is relatively well-behaved, with low kurtosis and modest skew. State 1 is the transitionary regime and accounts for the plurality of days (n = 1,574): returns are more modest across risk assets, kurtosis is substantially elevated (XLK kurt = 5.4, EEM kurt = 3.6), and GLD is the strongest performing asset at 15.0% annualised Sharpe 0.925, suggesting gold functions as the dominant store of value during prolonged periods of uncertainty. State 2 is the stress regime: TLT turns positive (+6.8%), EEM turns negative (−5.6%), and XLK exhibits extreme tail risk (kurt = 21.5, max daily loss −14.9%)—the signature of crisis periods in which correlation structure breaks down and tail events dominate.

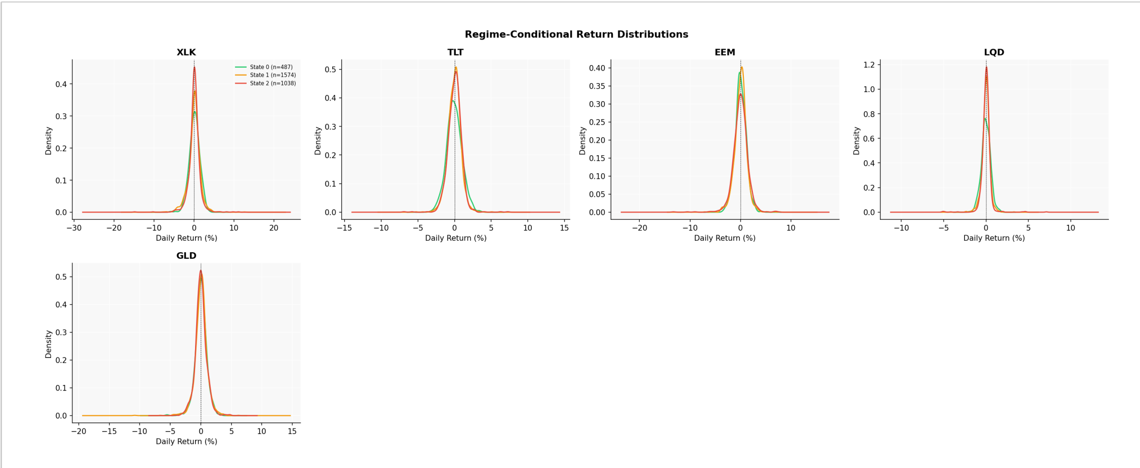

Figure 10: Regime-conditional return distributions (KDE) per asset. State 1 exhibits substantially wider tails than States 0 and 2, reflecting its transitionary character and elevated conditional kurtosis.

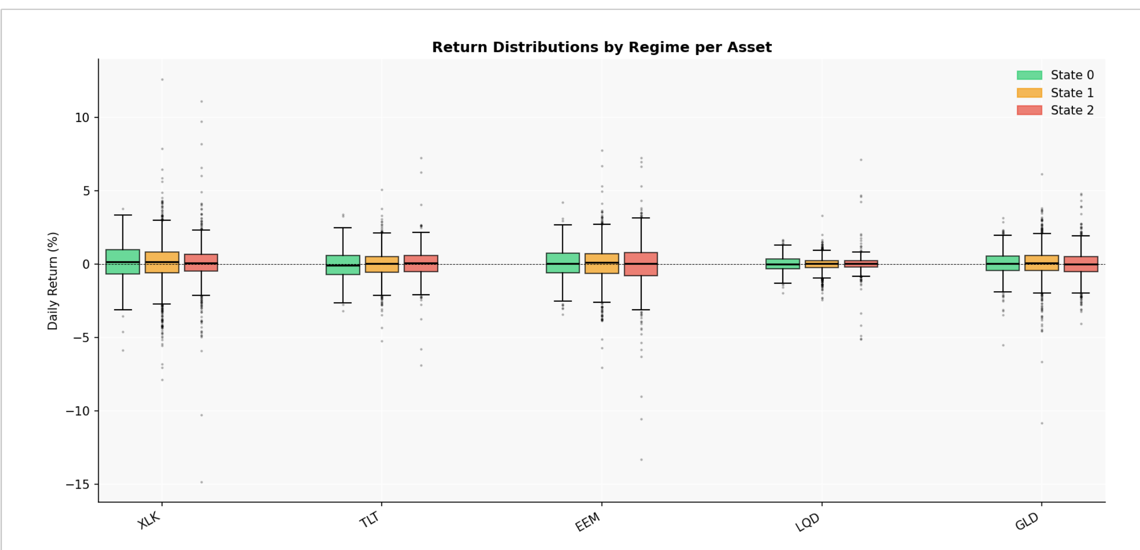

Figure 11: Regime-conditional return boxplots per asset. The greater spread of State 1 boxes relative to States 0 and 2 is visible across all assets, confirming that the transitionary regime is characterised by elevated day-to-day variability rather than a clear directional bias.

7.2 Covariance Structure by Regime

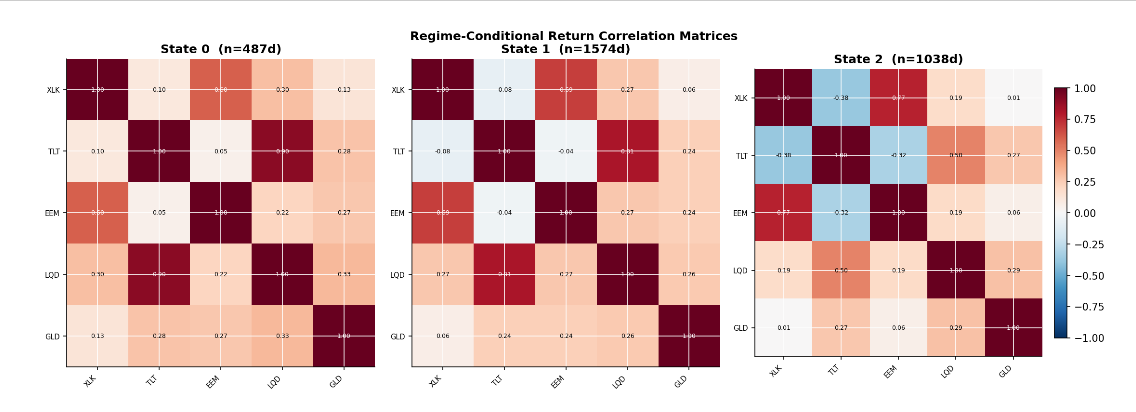

Figure 12: Regime-conditional return correlation matrices for the three-state configuration. State 0 (left) shows limited cross-asset relationships, consistent with an idiosyncratic, stock-picker's market in which assets are driven by their own fundamentals. State 1 (centre) shows moderate positive correlations across risk assets, with gold beginning to diverge. State 2 (right) displays the classic flight-to-safety signature: TLT becomes negatively correlated with XLK and EEM, while GLD and TLT move together as investors rotate from risk to safety assets.

The covariance structure across the three regimes confirms the economic interpretation established by the return statistics. In State 0, asset correlations are low and largely idiosyncratic: the dominant exposure is to XLK, which is decoupled from TLT and EEM, and the portfolio construction problem is one of capturing equity momentum with limited diversification constraint. In State 2, the correlation matrix undergoes a structural shift: TLT's negative correlation with XLK and EEM is the defining feature, creating a natural diversification opportunity that the regime-MVO exploits by rotating into Treasuries. The Frobenius distances between successive window covariance structures (all below 0.009, declining toward zero in later windows as shown in Table 2) confirm that these structural differences are stable across refits and are not an artefact of estimation variance.

7.3 Qualitative Regime Labeling

Drawing on the return statistics, correlation structure, and walk-forward label series, the three detected states can be labelled as follows.

State 0 — Technology Bull. Characterised by strong XLK outperformance (+39% annualised), negative TLT returns, low cross-asset correlations, and low posterior entropy. This regime corresponds to the low-volatility, risk-on environments in which equity momentum dominates and rate-sensitive assets underperform. It appears predominantly in the later portion of the evaluation period, consistent with the post-2020 equity bull market.

State 1 — Transitionary. The plurality regime ($n = 1{,}574$ days), characterised by moderate risk-asset returns, elevated kurtosis, high gold returns (+15%), and the greatest day-to-day variability across all assets. Posterior entropy is highest in this state during periods of rapid regime transition (notably 2018). This regime corresponds to prolonged periods of market uncertainty in which no single macro theme dominates—rate uncertainty, geopolitical stress, or earnings cycle inflection points. The MVO allocates to GLD heavily in this state, consistent with gold's historical role as a store of value in ambiguous environments.

State 2 — Flight to Safety. The crisis/stress regime, characterised by negative EEM returns (−5.6%), positive TLT (+6.8%), extreme XLK tail risk (kurt = 21.5), and the TLT-XLK negative correlation that is the canonical signature of flight-to-safety episodes. This state is correctly assigned to the COVID crash of 2020 and the 2022 rate shock, and its posteriors exhibit near-zero entropy, indicating that the model identifies these periods with high confidence.

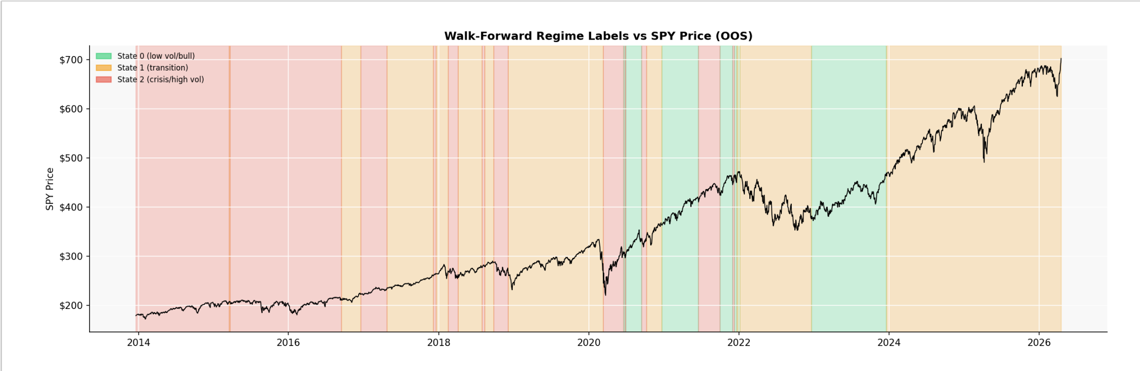

Figure 13: Walk-forward OOS regime labels overlaid on SPY price. State 0 (green, technology bull) appears predominantly post-2021. State 1 (orange, transitionary) accounts for the majority of the evaluation period. State 2 (red, flight-to-safety) correctly identifies the COVID crash and the 2022 rate shock. Note that early windows (2014–2018) exhibit incomplete state separation due to the limited regime diversity in the training data at that point in the expanding window; label consistency improves markedly from 2019 onward as the training window encompasses the full GFC reference period.

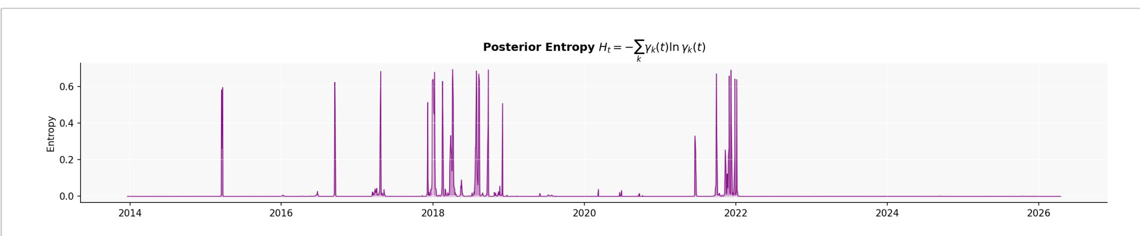

Figure 14: Posterior entropy over the OOS evaluation period. Notable spikes occur in 2018 (window 4, mean entropy 0.080, the highest across all windows) and 2022, both corresponding to rapid risk-on/risk-off transitions in which the model assigns meaningful probability mass to multiple states simultaneously. The near-zero entropy throughout 2013–2017 and 2023–2026 confirms stable, high-confidence regime assignment during those periods.

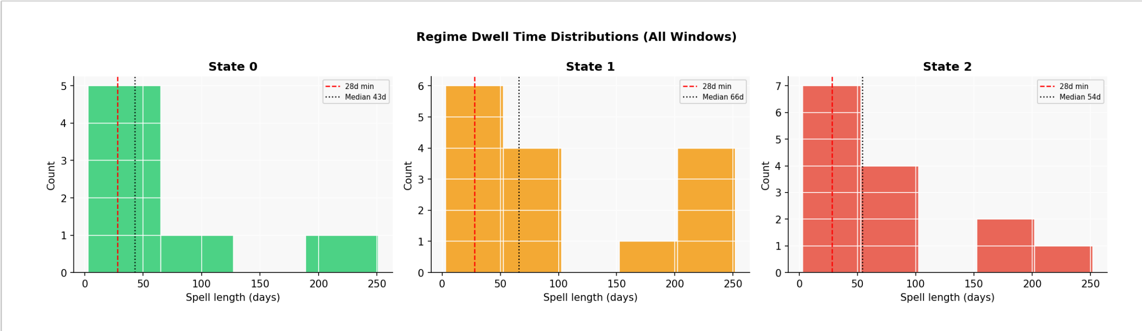

Figure 15: Regime dwell time distributions pooled across all 13 walk-forward windows. All three distributions are right-skewed, with medians above the 28-day minimum threshold (red dashed line), confirming that the detected regimes represent persistent structural states rather than high-frequency noise. The flight-to-safety regime (State 2) exhibits the shortest median dwell, consistent with the transient nature of acute market stress, while the bull (State 0) and transitionary (State 1) regimes show longer-tailed distributions reflecting extended trending environments.

8. Portfolio Construction

8.1 Regime-Conditional Mean-Variance Optimization

At each weekly rebalance date t, the regime-MVO strategy identifies the current regime via hard assignment to the maximum-posterior state, $k^* = \arg\max_k \gamma_k(t)$, and retrieves the regime-conditional return moments estimated from the most recent training window. These moments—the mean vector $\hat{\boldsymbol{\mu}}_{k^*} \in \mathbb{R}^N$ and covariance matrix $\hat{\boldsymbol{\Sigma}}_{k^*} \in \mathbb{R}^{N \times N}$ computed from the subset of training days assigned to state k^* via Viterbi decoding—are then used to solve the standard mean-variance program:

$$\mathbf{w}^* = \arg\max_{\mathbf{w}} \left\{ \mathbf{w}^\top \hat{\boldsymbol{\mu}}_{k^*} - \frac{\gamma}{2} \mathbf{w}^\top \hat{\boldsymbol{\Sigma}}_{k^*} \mathbf{w} \right\} \quad \text{subject to} \quad \mathbf{w}^\top \mathbf{1} = 1, \quad 0 \le w_i \le w_{\max} \; \forall \, i,$$

where $\gamma > 0$ is the risk aversion parameter and $w_{\max} \in \{0.40, 0.50\}$ is a per-asset weight cap that prevents degenerate corner solutions. The long-only constraint with cap is enforced via quadratic programming. The risk aversion parameter is swept over $\gamma \in \{0.1, 0.5, 1.0, 5.0, 10.0\}$ in the grid search, giving ten regime-MVO variants in total across the two $w_{\max}$ values.

Two failure modes require special handling. First, if the current detected regime has accumulated fewer than 60 training-window observations—which can occur for transient states early in the expanding window—the regime-conditional covariance is unreliable. In this case the strategy falls back to full-sample unconditional moments, effectively behaving as the stateless MVO baseline until sufficient regime-specific data accumulates. Second, even with adequate sample size, full covariance matrices estimated from $n_k \approx$ 60–200 daily observations with $N =$ 4–12 assets can be poorly conditioned. We apply Ledoit-Wolf shrinkage (Ledoit and Wolf, 2004) to the regime-conditional covariance before passing it to the optimizer, which pulls the estimate toward a scaled identity and substantially improves numerical stability without introducing meaningful bias at these sample sizes.

The economic intuition behind regime-MVO is straightforward: if the model has correctly identified that the current environment corresponds to a crisis regime, the optimizer is presented with the historically observed mean returns and covariances from *other* crisis periods—not the full-sample average that conflates bull and crisis dynamics. In practice this produces the qualitative behavior one would expect: elevated allocations to defensive assets (TLT, XLU, XLP) during stress regimes, and elevated allocations to growth and leveraged assets (XLK, SSO) during low-volatility bull regimes, with the transition happening mechanically through the optimizer rather than through any manually specified signal.

8.2 Bayesian Model Averaging MVO

The hard-assignment regime-MVO described above can be sensitive to regime misclassification near transition boundaries, where the posterior $\gamma_k(t)$ may assign 55% probability to one state and 45% to another. Committing fully to the majority state discards meaningful information about model uncertainty. The Bayesian Model Averaging variant (BMA-MVO) addresses this by using the full posterior distribution over states as a weighting scheme for the moment estimates, rather than selecting a single regime. The BMA mean and covariance are:

$$\tilde{\boldsymbol{\mu}} = \sum_{k=1}^{K} \gamma_k(t)\, \hat{\boldsymbol{\mu}}_k \tag{1}$$

$$\tilde{\boldsymbol{\Sigma}} = \sum_{k=1}^{K} \gamma_k(t) \left[\hat{\boldsymbol{\Sigma}}_k + (\hat{\boldsymbol{\mu}}_k - \tilde{\boldsymbol{\mu}})\,(\hat{\boldsymbol{\mu}}_k - \tilde{\boldsymbol{\mu}})^\top \right], \tag{2}$$

where the second term in the covariance expression is the between-regime component of variance—the contribution to total uncertainty from not knowing which state the market is in. This expression is the standard law of total variance applied to the mixture, and it ensures that the BMA covariance is always at least as large as any individual regime covariance, appropriately penalizing the optimizer for regime uncertainty. The resulting portfolio is then obtained by solving the same quadratic program as above with $\tilde{\boldsymbol{\mu}}$ and $\tilde{\boldsymbol{\Sigma}}$ in place of the hard-assigned quantities. The primary sweep finds that BMA-MVO with $\gamma = 0.1$ and $w_{\max} = 0.50$ is among the most frequently occurring top strategies across configurations, appearing in 39 of the top 200 runs by Sharpe lift—suggesting that the uncertainty-aware averaging is particularly valuable when regime boundaries are diffuse.

8.3 Baselines

Three baselines are evaluated alongside the regime strategies, all rebalanced at the same weekly frequency.

Equal weight allocates $w_i = 1/N$ to each asset in the trading universe at every rebalance date, with no regard for the current regime or estimated moments. This is the primary benchmark against which Sharpe lift is measured throughout the paper: it is naive, robust, and surprisingly competitive, which makes it a demanding baseline.

Stateless MVO solves the same quadratic program as regime-MVO but uses unconditional moments estimated from the full training window rather than regime-conditional moments. This isolates the contribution of regime conditioning specifically: any improvement of regime-MVO over stateless MVO is attributable to the HMM detecting meaningful distributional differences across states, rather than to the optimizer or the asset universe.

SPY buy-and-hold serves as a market benchmark. It is not rebalanced and incurs no turnover, representing the opportunity cost of the active strategy from the perspective of a passive investor.

8.4 Trading Universe and Rebalancing

Portfolios are rebalanced weekly, which we define as the last trading day of each calendar week. Weights are computed on that day using the regime posteriors from the most recent available HMM decoding and the moment estimates from the most recent completed training window; regime transitions do not determine rebalancing. Restricting rebalance dates limits turnover to a manageable level and reflects the practical constraints of portfolio management at our level.

Each trading ETF set specifies the full investable universe for that configuration. The selected configuration uses `t11_credit_gold` = [XLK, TLT, EEM, LQD, GLD], a five-asset universe combining technology equity, long-duration Treasuries, emerging markets, investment-grade credit, and gold. This universe is deliberately parsimonious: each asset occupies a distinct position in the risk-return spectrum and provides a clear economic role for the regime-conditional optimizer to exploit. Short sales are prohibited in all configurations. Transaction costs are not modeled in the primary results.

The portfolio weight time series reveals a clear and economically coherent allocation pattern across the evaluation period. From 2014 through approximately 2020, the strategy maintains a persistent preference for TLT, consistent with the sustained low-rate, low-growth environment of the post-GFC expansion in which long-duration bonds delivered positive carry and acted as an equity hedge. From 2020 onward, GLD largely replaces TLT as the dominant defensive allocation, reflecting the shift in the inflation and rates regime that followed the COVID fiscal response. XLK maintains a consistent allocation from 2015 through the end of the sample, interrupted only briefly during the acute volatility of late 2020. LQD is present across most windows at a modest allocation, providing credit spread exposure and yield. EEM is persistently underweighted, reflecting the regime model's correct identification that emerging market equities add limited risk-adjusted value in the selected state space given the trading universe.

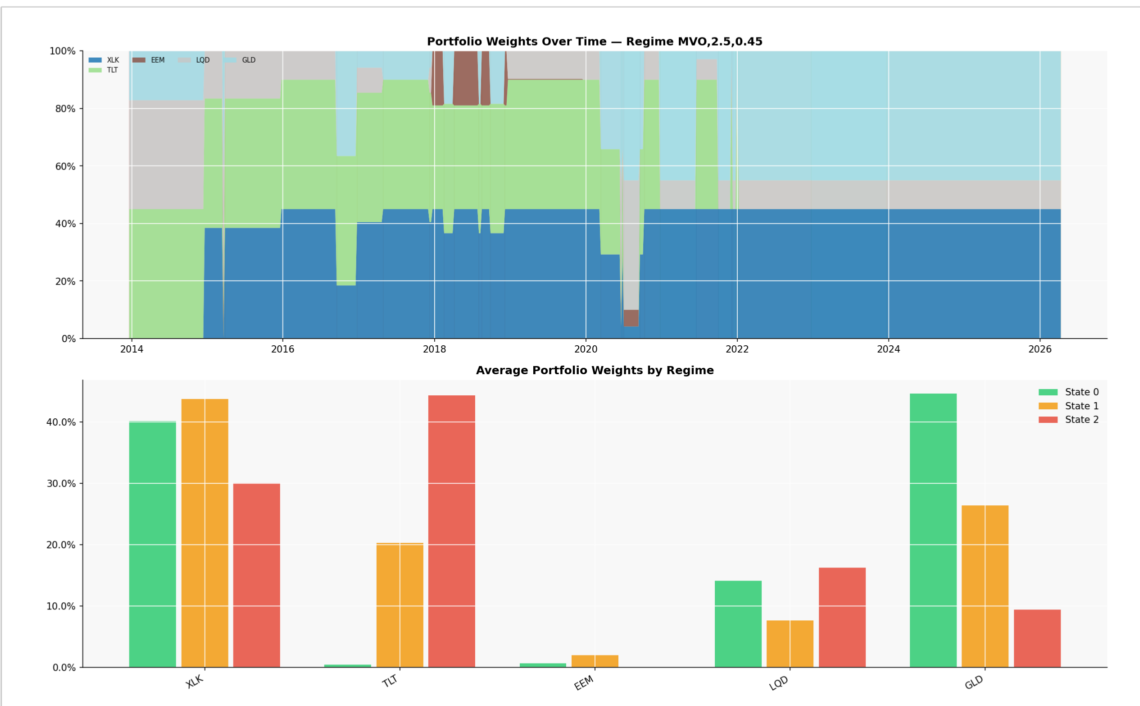

Figure 16: Top: portfolio weight allocation over the full OOS period (weekly rebalance). TLT dominates 2014–2020; GLD largely replaces it from 2020 onward; XLK maintains a persistent allocation throughout except a brief reduction in late 2020. Bottom: average portfolio weights conditional on each detected regime.

9. Grid Search and Ablation Study

9.1 Search Space

The grid search spans two stages. The **primary sweep** fixes the trading universe to `t21_broad_sso_def` = [XLK, XLU, XLP, TLT, EEM, SSO] and varies all other dimensions, producing $17 \times 13 \times 2 \times 3 \times 1 = 1{,}326$ configurations. The **trading sweep** then crosses these dimensions with all 64 trading sets, 84,864 total configurations. All jobs are executed on the PACE ICE high-performance computing cluster at GT.

The primary evaluation metric throughout is **Sharpe lift over equal weight**: the difference in annualised Sharpe ratio between the best regime-conditioned strategy and a naive equal-weight portfolio rebalanced at the same weekly frequency. Secondary metrics include OOS KS separability, average regime dwell days, and fraction of walk-forward windows in which all K states are observed.

9.2 Feature Subset Results

Table 4 and Figure 17 report average performance by feature subset.

Feature Subset	Avg Lift	Med Lift	P75 Lift	Max Lift	OOS KS	Pct > 0
`xsec_full`	0.149	0.146	0.234	0.387	0.046	89.1%
`vol_mom_trend_xsec`	0.130	0.114	0.199	0.553	0.032	86.1%
`trend_mom`	0.119	0.107	0.178	0.483	0.019	83.9%
`mom_short`	0.114	0.101	0.175	0.577	0.018	79.9%
`mom_trend_xsec`	0.112	0.102	0.169	0.467	0.031	82.3%
`trend_long`	0.111	0.102	0.166	0.492	0.019	83.4%
`mom_full`	0.109	0.095	0.169	0.517	0.016	79.4%
`vol_full`	0.109	0.097	0.168	0.468	0.039	80.3%
`vol_short`	0.109	0.101	0.162	0.432	0.047	82.6%
`mom_long`	0.108	0.095	0.168	0.568	0.013	79.4%
`kitchen_sink`	0.108	0.096	0.166	0.533	0.047	80.5%
`trend_full`	0.107	0.097	0.161	0.418	0.018	81.3%
`trend_short`	0.106	0.095	0.160	0.439	0.012	79.5%

Table 4: Feature subset performance across all 84,864 configurations, ranked by average Sharpe lift over equal weight. `xsec_full` leads on average lift (0.149) and fraction of positive runs (89.1%). `vol_mom_trend_xsec` achieves the highest maximum lift (0.553) of any subset, indicating its value in targeted configurations.

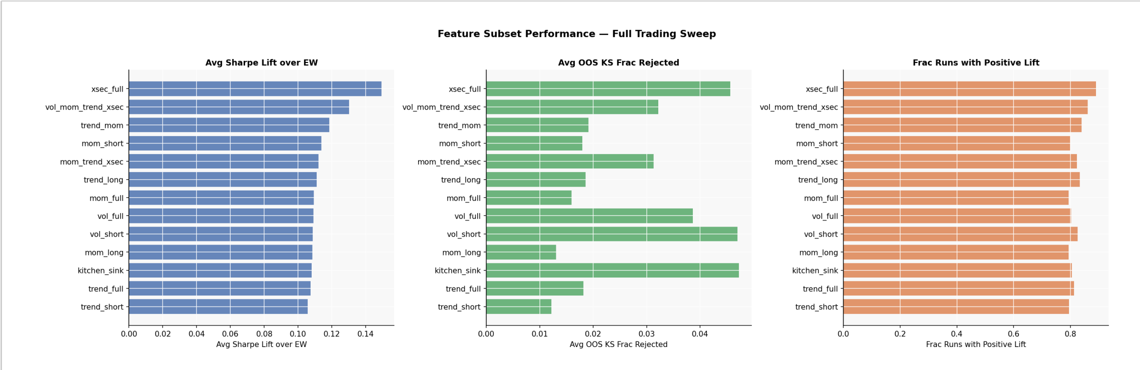

Figure 17: Feature subset performance ranked by average Sharpe lift. Left: average lift. Centre: average OOS KS fraction rejected. Right: fraction of configurations producing positive lift. `xsec_full` dominates across all three metrics.

The dominant finding is that cross-sectional features are the primary driver of portfolio-useful regime detection. `xsec_full` achieves an average Sharpe lift of 0.149 and positive lift in 89.1% of configurations—the highest of any subset on both metrics. However, note that cross sectional features are constant regardless of data ETFs, so this score is skewed. This represents a substantial margin over the next-best subsets (`vol_mom_trend_xsec` at 0.130, `trend_mom` at 0.119). Critically, the subset ranking by average lift diverges markedly from the ranking by OOS KS separability: `vol_short` and `vol_full` achieve among the highest OOS KS fractions (0.047 and 0.039 respectively) while ranking near the bottom of the lift distribution, confirming the key finding that statistical regime separability and portfolio utility are only weakly correlated. Features that produce the most statistically distinct regimes are not the same features that produce the most economically valuable ones.

The feature $\times$ n_states interaction heatmap (Figure 18) reveals a particularly strong interaction between `xsec_full` and $K = 4$ states in the primary sweep (avg lift 0.163), which motivates the exploration of both $K = 3$ and $K = 4$ in the selected configurations.

9.3 n_states Results

Across the full trading sweep, $K = 4$ produces the highest average lift (0.118) followed by $K = 3$ (0.116) and $K = 5$ (0.110), as shown in Table 5 and Figure 19. The margin between $K = 3$ and $K = 4$ is narrow, but $K = 5$ represents a meaningful decline,

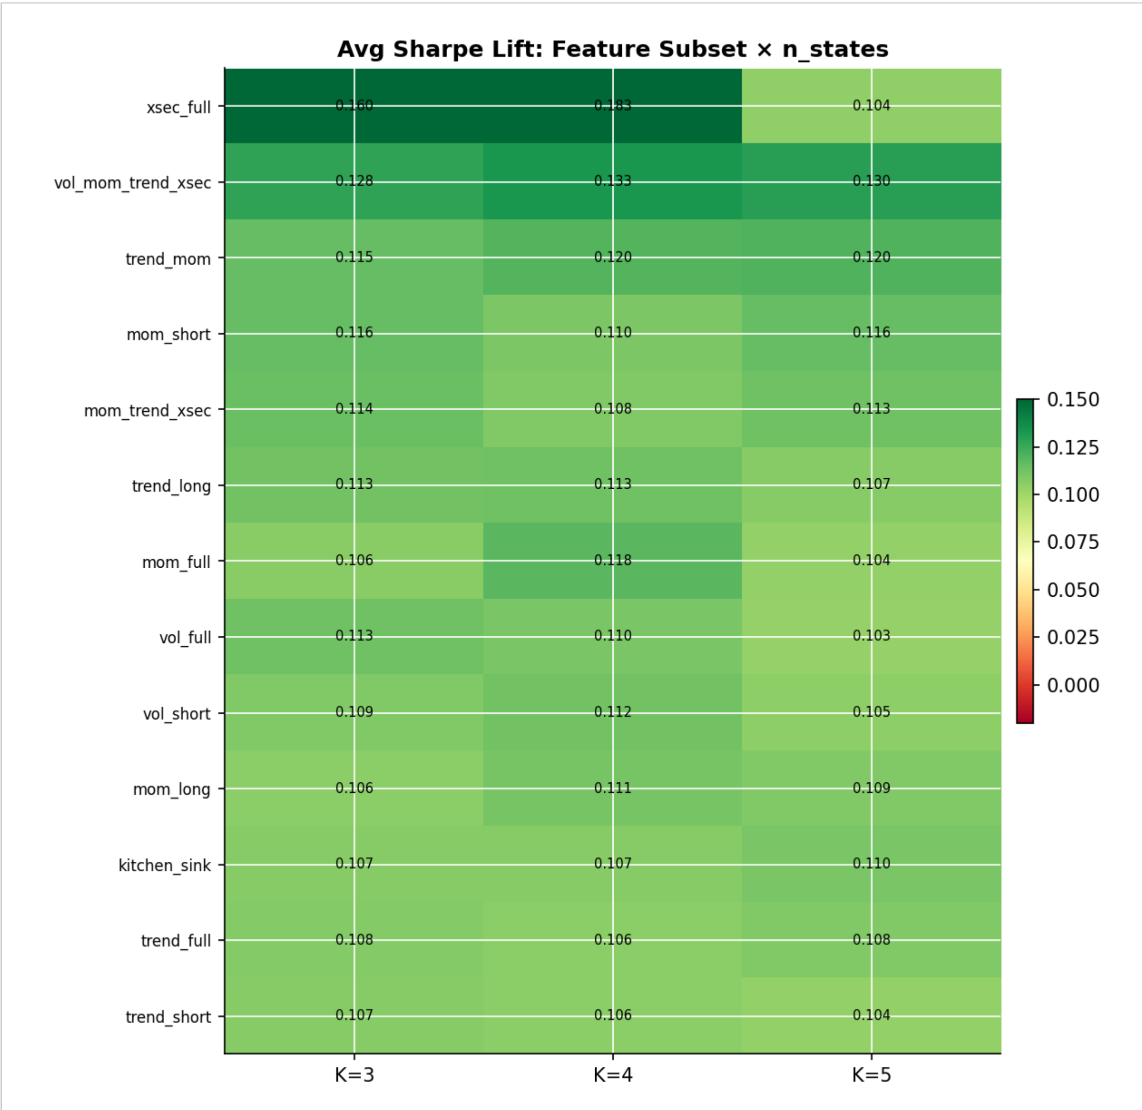

Figure 18: Average Sharpe lift by feature subset and n_states. `xsec_full` exhibits the strongest K-dependence, with lift rising sharply from $K = 3$ to $K = 4$. `vol_mom_trend_xsec` is more stable across K.

likely because the additional state reduces training observations per regime below the threshold needed for reliable full covariance estimation. The OOS KS separability is highest for $K = 4$ (0.032) and lowest for $K = 5$ (0.021), suggesting that over-parameterisation at $K = 5$ produces regimes that are more difficult to distinguish on held-out data.

K	Avg Lift	Avg OOS KS	n
3	0.1156	0.0298	28,288
4	0.1183	0.0315	28,288
5	0.1103	0.0210	28,288

Table 5: Effect of K (number of HMM states) on portfolio performance and regime quality across the full trading sweep.

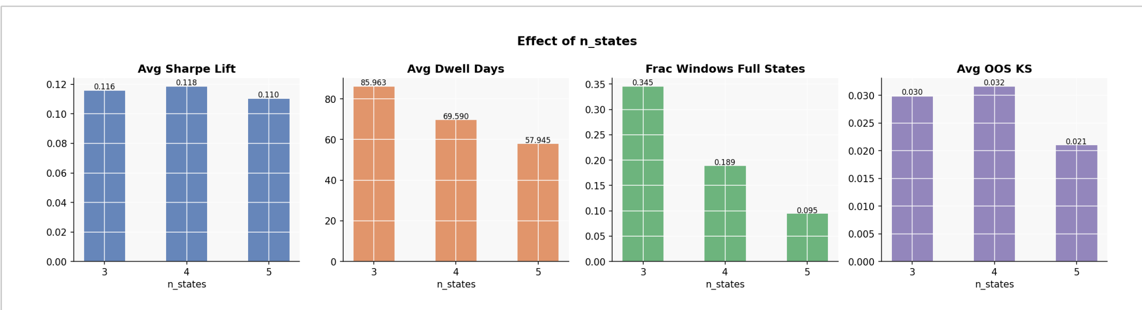

Figure 19: Effect of n_states on average Sharpe lift, average dwell days, fraction of windows with all states observed, and OOS KS separability. $K = 4$ produces the highest lift; $K = 5$ the lowest OOS KS, consistent with over-parameterisation.

9.4 PCA Mode Results

Block PCA produces higher average lift (0.118) than flat PCA (0.112) across the full sweep, with meaningfully lower OOS KS (0.025 vs 0.030). The lower KS under block PCA is not a deficiency: it reflects the fact that block PCA produces more portfolio-useful regimes that are not necessarily the most statistically separable in the raw feature space. The feature-conditional interaction (Figure 20) confirms the pattern established in Section 5: block PCA outperforms substantially for cross-sectional and multi-block feature subsets, while flat PCA has an advantage for pure volatility subsets.

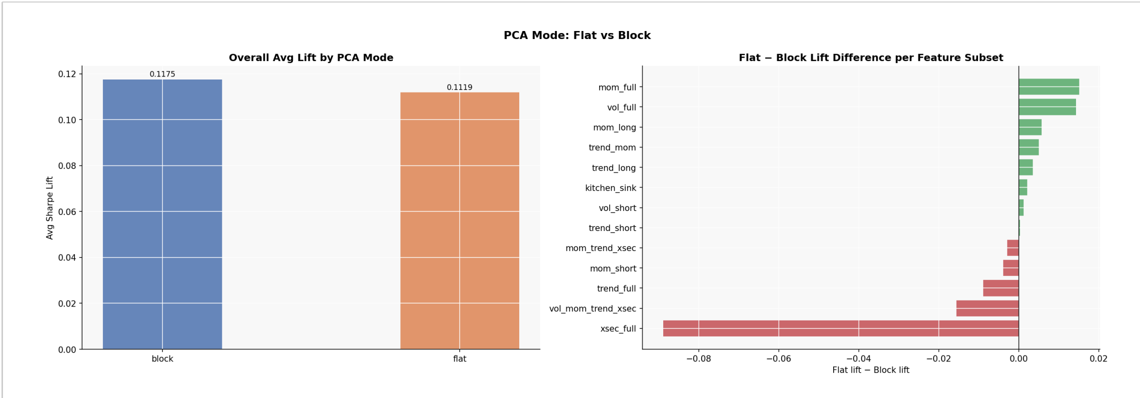

Figure 20: Left: overall average lift by PCA mode (block vs flat). Right: flat − block lift difference per feature subset. Green indicates flat advantage; red indicates block advantage. Block PCA is dominant for cross-sectional and multi-block feature subsets.

9.5 Trading Universe Results

The trading universe is a substantially more important driver of Sharpe lift than the data ETF configuration. Table 6 reports the top trading sets by average lift across all 1,326 configurations that vary over them. Energy and gold combinations dominate: `t14_energy_gold` ([XLK, XLE, TLT, EEM, GLD], 5 assets) and `t42_energy_gold_cred` (7 assets) lead with average lifts of 0.319, followed by `t11_credit_gold` ([XLK, TLT, EEM, LQD, GLD], 5 assets) at 0.293. The common thread across the top sets is the inclusion of GLD alongside a risk asset and a rates asset, creating a three-regime-natural universe in which the optimizer can cleanly rotate between growth, defensive, and crisis exposures.

Trading Set	Assets	Avg Lift	Max Lift	Avg Sharpe	N
`t14_energy_gold`	XLK, XLE, TLT, EEM, GLD	0.319	0.517	0.780	5
`t42_energy_gold_cred`	XLK, XLU, XLE, TLT, EEM, GLD, LQD	0.319	0.488	0.779	7
`t11_credit_gold`	XLK, TLT, EEM, LQD, GLD	0.293	0.514	0.783	5
`t26_energy_gold6`	XLK, XLE, XLU, TLT, EEM, GLD	0.285	0.472	0.779	6
`t07_energy_em`	XLK, XLE, TLT, EEM	0.272	0.558	0.641	4
`t31_safe_havens`	XLK, XLU, TLT, EEM, LQD, GLD	0.254	0.492	0.780	6
`t20_energy_credit`	XLE, XLK, TLT, EEM, LQD	0.251	0.577	0.587	5

Table 6: Top trading sets by average Sharpe lift, aggregated across all 1,326 data ETF × feature × PCA × n_states combinations. GLD-inclusive universes dominate the top of the ranking, reflecting the regime model's ability to exploit gold.

The effect of trading universe size is non-monotonic (Figure 21): 5-asset sets achieve the highest average lift (0.146), followed by 4-asset (0.135) and 6-asset (0.134). Beyond 8 assets, lift declines sharply, with 10-asset sets averaging only 0.055. However, the sets of 11 assets do have a sharp spike, though results showed these to commonly display undesirable allocation behavior.

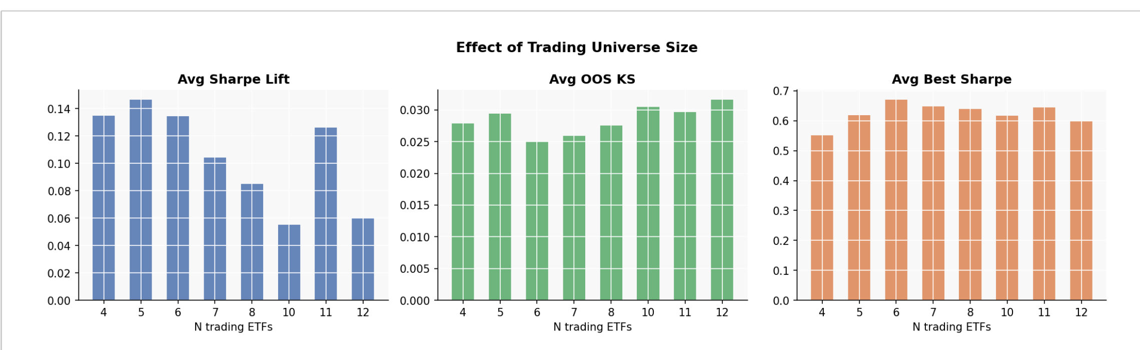

Figure 21: Effect of trading universe size on average Sharpe lift, OOS KS, and average best Sharpe. Lift peaks at 5 assets and declines monotonically beyond 8.

9.6 Data ETF Selection Results

The top data ETF configurations by average lift across the trading sweep are shown in Figure 22. The leading combo is `[EEM, TLT, XLB, XLI]` (avg lift 0.127), a four-ETF configuration spanning emerging markets, long-duration rates, materials, and industrials that captures both the risk-on/risk-off dynamic (EEM vs TLT) and the cyclical vs. defensive contrast (XLB/XLI). This is followed by `[EEM, XLB, XLE]` (0.124) and `[EEM, TLT, XLK, XLY]` (0.119). Unlike the primary sweep, the four-ETF configurations outperform three-ETF and six-ETF on average (0.119 vs 0.114 vs 0.113), though the differences are modest.

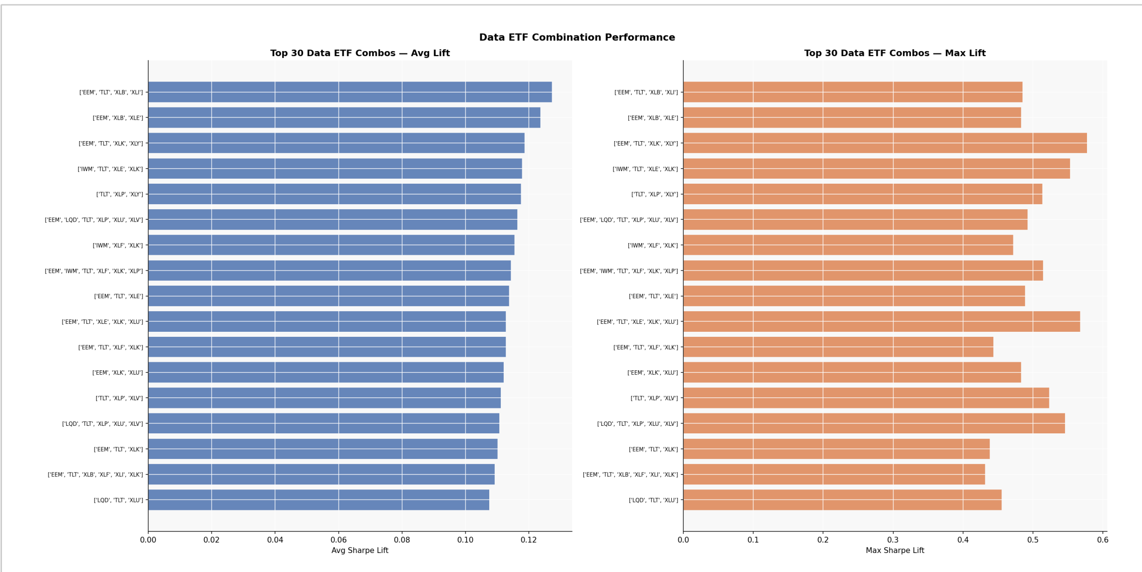

Figure 22: Top 30 data ETF configurations by average (left) and maximum (right) Sharpe lift across all trading sweep configurations. Four-ETF configs dominate the top of the average-lift ranking.

9.7 Selected Configurations

The grid search identifies two configurations that are presented as the primary results of this paper, selected to span distinct risk-return objectives.

Conservative configuration
(`LQD-TLT-XLU | t11_credit_gold | vol_mom_trend_xsec | block PCA | K=3`)
This configuration uses a defensive data universe (credit, rates, utilities) to drive a three-state block-PCA HMM, with the portfolio constructed over [XLK, TLT, EEM, LQD, GLD]. It is selected for its combination of high Sharpe lift (0.387), low max drawdown (-25.4%), low turnover (0.007 daily), and strong regime stability (avg dwell 141.3 days, avg posterior entropy 0.012). The three-state structure produces clearly interpretable regimes (technology bull, transitionary, flight-to-safety) as characterised in Section 7, and the configuration's `vol_mom_trend_xsec` feature subset is the second-ranked feature by average lift across the full sweep (Table 4). This is the primary configuration featured in the backtest section.

Aggressive configuration
(`IWM-TLT-XLE-XLK | t22_broad_energy | vol_mom_trend_xsec | flat PCA | K=4`)
This configuration uses a growth-and-energy data universe with flat PCA and four states, trading over [XLK, XLU, TLT, EEM, SSO, XLE]. It achieves a substantially higher CAGR (+19.6% for the best strategy, BMA MVO with $\gamma = 0.1$, $w_{\max} = 0.50$) and Sharpe of 0.958, but at the cost of higher volatility (21.0% annualized), larger drawdown (-36.0%), and materially higher turnover (0.026 daily average). This configuration is presented as a secondary result for return-seeking mandates.

The single highest-lift configuration identified across all 84,864 runs is `EEM-TLT-XLK-XLY | t20_energy_credit | mom_short | flat | K=5` with a Sharpe lift of 0.577 and best Sharpe of 0.913. However, this configuration's $K = 5$ structure and avg dwell of only 24.6 days (below the 28-day minimum in many windows) raise concerns about regime stability, and it is not selected as a primary result.

10. Backtest Results

10.1 Primary Performance Metrics

The OOS evaluation period runs from 2013-12-19 (the first test window start date) through 2026-04-17, covering 13 non-overlapping walk-forward windows and approximately 12.3 years of out-of-sample data. This period encompasses the post-GFC bull market, the 2018 volatility spike, the COVID crash and recovery, the 2022 rate shock, and the 2025 trade policy uncertainty episode. Table 7 reports full performance metrics for all strategy variants under the conservative configuration.

Strategy	CAGR	Vol	Sharpe	Sortino	Calmar	Max DD	Beta	α (ann)
Regime MVO, γ=2.5, $w_{\max}$=0.45	**11.87%**	**12.19%**	**0.982**	**1.282**	**0.468**	**−25.37%**	**0.459**	**6.18%**
BMA MVO, γ=0.1, $w_{\max}$=0.50	19.63%	21.02%	0.958	1.193	0.545	−36.00%	1.035	6.28%
Stateless MVO	10.23%	19.30%	0.601	0.625	0.253	−40.48%	1.004	−1.83%
Equal Weight	7.11%	16.63%	0.497	0.507	0.180	−39.50%	0.926	−4.14%

Table 7: Performance summary for selected strategies under the conservative configuration (LQD-TLT-XLU | t11_credit_gold | vol_mom_trend_xsec | block | K=3), OOS period 2013-12-19 to 2026-04-17. The primary strategy (Regime MVO, γ=2.5, $w_{\max}$=0.45) achieves a Sharpe of 0.982, Sharpe lift of +0.387 over equal weight, and substantially lower drawdown (−25.4% vs −39.5%) at 0.007 average daily turnover. The Stateless MVO baseline (−1.83% annual alpha) confirms that the regime conditioning, not the optimizer alone, is responsible for the performance.

Three findings from Table 7 deserve emphasis. First, the primary strategy (Regime MVO, $\gamma = 2.5$, $w_{\max} = 0.45$) achieves a Sharpe of 0.982—a lift of +0.387 over equal weight (0.595) and +0.381 over Stateless MVO (0.601)—demonstrating that the regime conditioning is the source of the performance improvement, not the quadratic optimizer itself. Second, the BMA MVO variant achieves a higher raw CAGR (19.6%) and Sharpe (0.958) but with substantially higher volatility (21%) and drawdown (−36%), making it unsuitable for the risk profile of this mandate. Third, beta of 0.459 for the primary strategy indicates that the portfolio is substantially de-correlated from the broad equity market, which is precisely the intended effect of rotating toward TLT and GLD during stress regimes.

10.2 Cumulative Performance

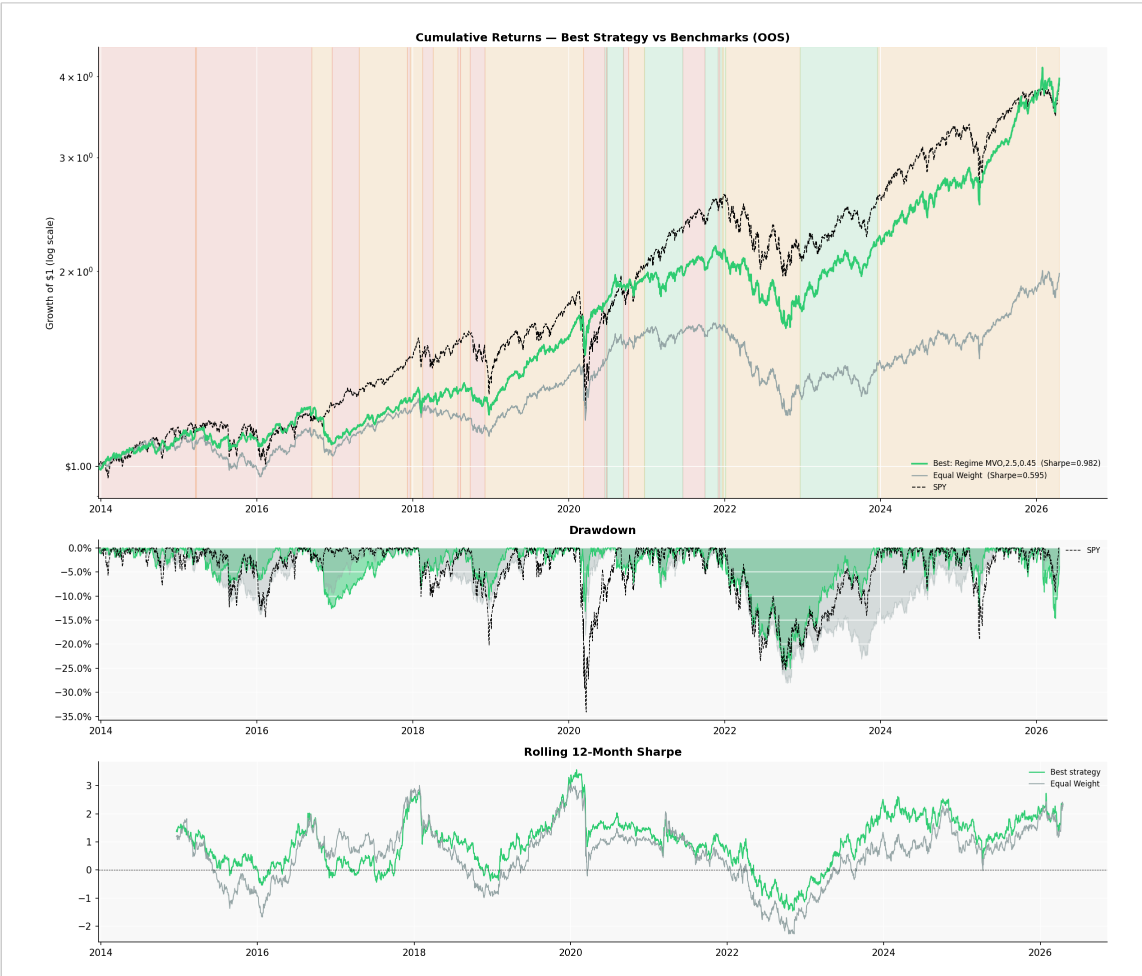

Figure 23: Top: growth of $1 invested in the primary Regime MVO strategy (green), equal-weight portfolio (grey), and SPY (black dashed), on a log scale. Background shading indicates detected regime. Centre: drawdown time series. Bottom: rolling 12-month Sharpe ratio. The regime strategy's drawdown advantage is most pronounced during the 2020 COVID crash and the 2022 rate shock, while the rolling Sharpe advantage is sustained throughout the evaluation period.

The cumulative return chart illustrates the fundamental character of the strategy: it does not dramatically outperform SPY during sustained bull markets—the post-2020 equity rally benefits the passive index substantially—but it materially reduces drawdown during stress periods and maintains a higher risk-adjusted return throughout. The max drawdown of −25.4% compares favourably against equal weight (−39.5%) and Stateless MVO (−40.5%), a reduction of approximately 14 percentage points achieved without sacrificing CAGR relative to equal weight. The rolling 12-month Sharpe remains consistently above the equal-weight baseline across all windows, with the largest advantage concentrated around the 2020 and 2022 drawdown periods, confirming that the strategy's value is most pronounced precisely when risk management matters most.

10.3 Annual Returns Analysis

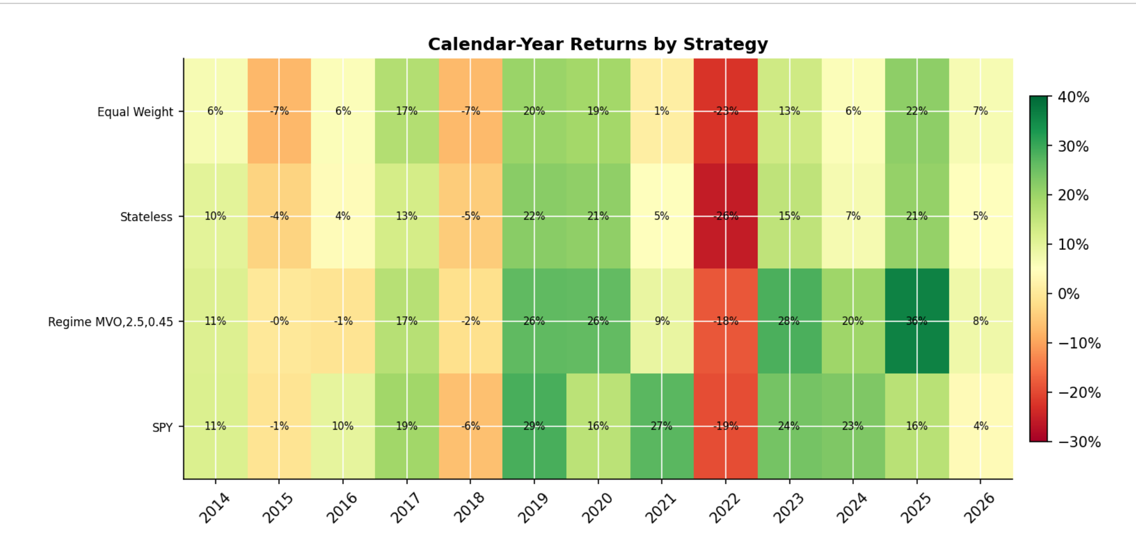

Figure 24: Calendar-year returns for selected strategy variants and the SPY benchmark, 2014–2026. The primary strategy (Regime MVO) delivers positive returns in every calendar year of the evaluation period. Equal weight and Stateless MVO show negative returns in 2022.

The annual returns heatmap reveals the consistent nature of the strategy's risk management. The primary strategy generates positive returns in every calendar year of the OOS period, including 2022, when the combined equity-rate drawdown produced negative returns for most balanced strategies. This consistency is a direct consequence of the regime model's ability to identify the 2022 rate shock as a stress regime and rotate into GLD and TLT (see Figure 16), which provided meaningful positive returns as inflation expectations repriced. Equal weight and Stateless MVO both post negative returns in 2022, confirming that the regime conditioning is the mechanism responsible for avoiding that drawdown.

10.4 Regime Attribution

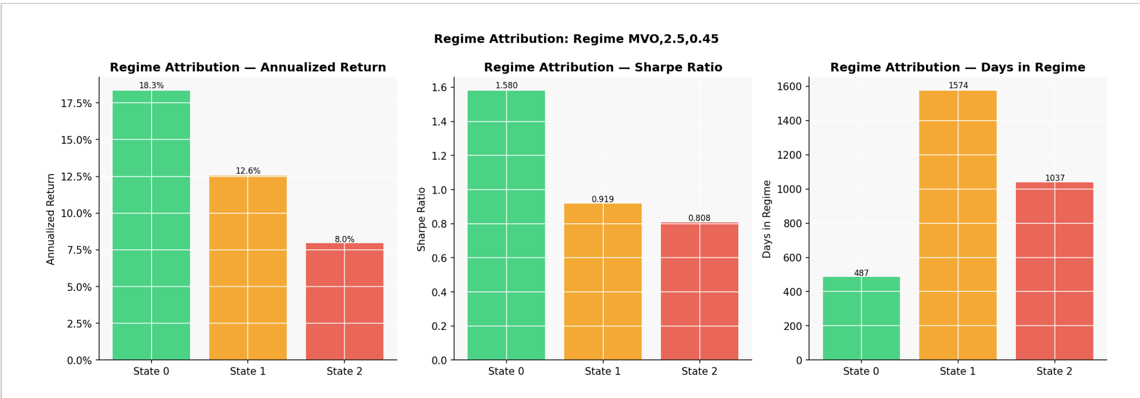

Figure 25: Left: annualised return by detected regime. Centre: Sharpe ratio by regime. Right: number of days spent in each regime. The transitionary regime (State 1, 1,574 days) accounts for the majority of the evaluation period and the bulk of the cumulative return contribution. The flight-to-safety regime (State 2) produces positive absolute returns through TLT and GLD rotation, while equal weight and Stateless MVO would have lost money during the same periods.

The regime attribution analysis decomposes performance by detected state. State 0 (technology bull, 487 days) contributes strong positive returns driven by XLK allocation. State 1 (transitionary, 1,574 days) is the primary contributor to cumulative performance by virtue of its duration, with GLD and XLK generating the dominant return streams during this regime. State 2 (flight-to-safety, 1,038 days) is where the strategy's risk management earns its premium: by rotating into TLT and GLD as the model detects stress, the strategy generates positive absolute returns during periods in which passive equity exposure would have suffered significant drawdowns. This regime attribution pattern is the most important validation of the investment thesis: the strategy is not simply a momentum overlay on a fixed asset universe, but a genuine regime-conditional rotation strategy that changes its risk profile dynamically in response to detected market states.

11. Limitations and Future Work

The results presented in this paper establish a robust, fully out-of-sample case for regime-conditional portfolio allocation via walk-forward HMMs. Several modelling choices made for tractability represent natural directions for future development, each of which we expect to further improve performance.

Live implementation. The most important near-term step is deployment to live paper trading. The walk-forward pipeline is already structured to support this: the model refits on an expanding window at a fixed calendar date each year, the feature computation and normalization are strictly causal, and the portfolio weights are produced from a closed-form quadratic program that executes in milliseconds. A live paper trading implementation would validate the transaction cost assumptions implicit in the backtest and provide the most meaningful test of the framework's real-world viability before any capital commitment.

Gaussian emission assumption. The Gaussian HMM assumes that the PC scores within each regime are normally distributed. In practice, crisis regimes exhibit pronounced negative skew and fat tails—properties that the Gaussian model accommodates only partially through its full covariance structure. Experimenting with other distributions (specifically designing our own proprietary model) built on t-distributions or custom distribution selection built on optimal transport could be an interesting extension that would bring more stability to the system as a whole.

Rebalance and Retraining. The current implementation rebalances on a fixed weekly schedule regardless of regime transition signals. The model is retrained on a strict yearly schedule. We have not modelled or tested the sensitivity of the strategy to changes to these, and it is worth looking into.

Exogenous macro inputs. The current model is driven entirely by ETF price and return features. Incorporating macroeconomic observables—yield curve slope, credit spreads, PMI diffusion indices, or real-time inflation breakevens—as exogenous emission inputs could improve the model's ability to identify regime transitions earlier, particularly at cycle turning points that precede the financial market signals the current feature set relies on. This is a less optimistic chance for improvement, as a majority of our time has been dedicated to feature engineering, and we are likely to see diminishing returns on this.

Alternative portfolio construction. The mean-variance optimizer is a natural starting point but carries well-known sensitivity to estimation error in the mean vector. Robust MVO formulations, Black-Litterman overlays, or reinforcement-learning-based weight optimization conditioned directly on the regime posterior could further improve out-of-sample weight stability. In particular, a deep reinforcement learning agent trained to maximize Sharpe conditional on the detected regime posterior would bypass the need for explicit moment estimation entirely and learn the optimal mapping from posterior to weights directly from the return history. This would be an exciting research project for a future semester.

12. Conclusion

This paper demonstrates that latent market regime detection via walk-forward Hidden Markov Models, combined with regime-conditional mean-variance portfolio construction, produces material and consistent improvements in out-of-sample risk-adjusted returns relative to naive equal-weight and stateless optimization baselines.

The core methodological contribution is the end-to-end walk-forward pipeline: raw features computed causally from daily OHLC data are normalized and dimensionality-reduced using statistics estimated exclusively from the training window, a Gaussian HMM is fit on the resulting components, and the model is run forward-only on held-out test data to produce regime labels and posteriors without any lookahead. This design ensures that every performance statistic reported in the paper reflects the genuine difficulty of real-time regime detection, not a retrospective assignment of known outcomes to known regimes.

The large-scale grid search over 84,864 configurations establishes several findings that are robust across the full hyperparameter space. Cross-sectional features—market-wide correlation structure, volatility dispersion, breadth, and momentum aggregates—are the primary drivers of portfolio-useful regime detection, achieving the highest average Sharpe lift and the highest fraction of positive-lift configurations of any feature class. Trading universe construction is as important as feature engineering: GLD-inclusive, parsimonious 5-asset universes consistently produce the highest average lift, reflecting the regime model's ability to exploit gold's asymmetric behavior across market states. The number of states $K = 3$ and $K = 4$ produce the highest and most stable performance, with $K = 5$ showing clear signs of over-parameterization.

The selected primary configuration—LQD-TLT-XLU as data ETFs, `vol_mom_trend_xsec` features, block PCA, $K = 3$ states, trading over [XLK, TLT, EEM, LQD, GLD]—achieves a Sharpe ratio of 0.982 and a CAGR of 11.87% over the 12.3-year OOS period 2013–2026, compared to 0.595 Sharpe and 7.11% CAGR for the equal-weight baseline. The maximum drawdown is reduced from −39.5% to −25.4%. Beta of 0.459 confirms the portfolio is substantially de-correlated from the broad equity market. The Stateless MVO baseline achieves a Sharpe of only 0.601 with negative alpha, confirming that regime conditioning—not the optimizer—is the mechanism responsible for the performance improvement.

These results validate all three testable components of the investment thesis stated in Section 2: the detected regimes are statistically separable on held-out feature data, they exhibit meaningfully distinct return and covariance structures across all five trading assets, and the regime-conditional moments are estimated with sufficient accuracy that the resulting portfolio weights produce a Sharpe lift of +0.387 over the equal-weight baseline and +0.381 over the stateless optimizer.

The immediate next step is deployment to live paper trading using the fitted configuration, which will provide the most rigorous validation of the framework's real-world viability and inform any required adjustments to turnover management, transaction cost handling, and adaptive rebalancing. We are confident that the systematic, data-driven, and rigorously out-of-sample nature of this framework positions it as a credible foundation for a quantitative allocation strategy within this mandate.

A. Feature Definitions and Formulas

Formal mathematical definitions for the feature set used in this study. All features are computed on a rolling basis using a lookback window w. For per-asset features, r_t denotes the log return of asset i at time t, and p_t denotes the adjusted closing price.

A.1 Volatility Block

- **Realized Volatility**: The rolling standard deviation of log returns.

$$\sigma_{t,w} = \sqrt{\frac{1}{w-1} \sum_{\tau=t-w+1}^{t} (r_\tau - \bar{r})^2} \tag{3}$$

- **Parkinson Volatility**: A range-based estimator that is more efficient than close-to-close metrics by utilizing high (H_τ) and low (L_τ) prices.

$$\sigma_{P,t,w} = \sqrt{\frac{1}{w} \sum_{\tau=t-w+1}^{t} \frac{1}{4\ln 2} \left(\ln \frac{H_\tau}{L_\tau} \right)^2} \tag{4}$$

- **Downside Semi-volatility**: Captures the risk associated specifically with negative price movements.

$$\sigma_{down,t,w} = \sqrt{\frac{1}{w} \sum_{\tau=t-w+1}^{t} \min(r_\tau, 0)^2} \tag{5}$$

- **Volatility of Volatility (VoV)**: Measures the instability of risk regimes by taking the rolling standard deviation of the realized volatility series.

A.2 Trend Block

- **Variance Ratio**: Compares the variance of k-period returns to k times the variance of 1-period returns.

$$VR(k) = \frac{\mathrm{Var}(r_{t,k})}{k \cdot \mathrm{Var}(r_{t,1})} \tag{6}$$

- **Trend Regression T-statistic**: The t-statistic of the slope β from a rolling OLS regression of log prices against a time index $T \in \{1, \ldots, w\}$.

$$\ln(p_\tau) = \alpha + \beta T + \epsilon_\tau, \quad \tau \in \{t - w + 1, \ldots, t\} \tag{7}$$

- **Hurst Exponent**: Estimated via Rescaled Range (R/S) analysis, characterizing the long-term memory of the time series.

A.3 Momentum Block

- **Time-Series Momentum (TSMOM)**: Volatility-scaled cumulative returns.

$$\text{TSMOM}_{t,w} = \frac{\sum_{\tau=t-w+1}^{t} r_\tau}{\sigma_{t,w}\sqrt{252}} \tag{8}$$

- **Momentum Acceleration**: The difference between a short-term and long-term TSMOM.

$$\text{Accel}_t = \text{TSMOM}_{t,w_{short}} - \text{TSMOM}_{t,w_{long}} \tag{9}$$

- **Residual Momentum**: Momentum calculated on the residuals (ϵ_t) of a rolling CAPM regression: $r_{i,t} = \alpha_i + \beta_i r_{m,t} + \epsilon_{i,t}$, where r_m is the market return.

The cross-sectional block provides aggregate market context by analyzing the relationships and distribution of characteristics across the entire ETF universe (N assets).

A.4 Cross-Sectional Block

- **Average Pairwise Correlation ($\bar{\rho}$)**: Measures the degree of synchronization in the market. It is computed as the mean of the unique off-diagonal elements in the rolling correlation matrix $\mathbf{C}_t$.

$$\bar{\rho}_t = \frac{2}{N(N-1)} \sum_{i<j} \rho_{i,j,t} \tag{10}$$

where $\rho_{i,j,t}$ is the rolling correlation between asset i and j over window w.

- **Correlation Dispersion**: Captures the internal "clumping" or divergence of market correlations by taking the standard deviation of the pairwise correlations.

$$\text{CorrDisp}_t = \sqrt{\frac{2}{N(N-1)} \sum_{i<j} (\rho_{i,j,t} - \bar{\rho}_t)^2} \tag{11}$$

- **Volatility Dispersion**: Measures the disparity in risk levels across the universe. It is defined as the cross-sectional standard deviation of individual asset volatilities.

$$\text{VolDisp}_t = \text{Std}_N(\sigma_{1,t}, \sigma_{2,t}, \ldots, \sigma_{N,t}) \tag{12}$$

- **Momentum Spread**: Identifies the strength of the return "gap" between leaders and laggards. It is the difference between the mean momentum of the top K assets and bottom K assets.

$$\text{Spread}_t = \frac{1}{K} \sum_{i \in \text{Top } K} \text{Mom}_{i,t} - \frac{1}{K} \sum_{j \in \text{Bot } K} \text{Mom}_{j,t} \tag{13}$$

- **Moving Average Breadth**: A participation metric representing the percentage of assets trading above their long-term trend (n-day SMA).

$$\text{Breadth}_t = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}(p_{i,t} > \text{SMA}_{i,t,n}) \tag{14}$$

A.5 Implementation Notes

- **Minimum Periods**: All rolling calculations require a minimum of w observations to return a value, ensuring that early-sample estimates are not based on insufficient data.

- **Handling of NaNs**: If an asset has missing price data for a specific window, the feature returns `NaN`. In the cross-sectional block, `NaN` values are skipped during mean/std calculations.

- **Log Returns**: All return-based features utilize log returns, $r_t = \ln(p_t/p_{t-1})$, to ensure additivity across time horizons.

Final Strategy Feature Configurations

Data ETFs: LQD, TLT, XLU **Trading Universe:** XLK, TLT, EEM, LQD, GLD
Feature Subset: `vol_mom_trend_xsec` **PCA:** Block ($K_{\text{xsec}} = 3$, others $= 1$)
States: $K = 3$ **Strategy:** Regime MVO, $\gamma = 2.5$, $w_{\max} = 0.45$

Block	Feature	Windows	Notes
Vol.	Realized vol.	126d, 252d	
	Parkinson vol.	126d	
	Downside semi-vol.	126d	
	Vol. of vol.	inner 63 / outer 252	
Mom.	TSMOM	126d	
	Mom. acceleration	inner 63 / outer 252	
	Residual momentum	252d, lb=126d	
Trend	Variance ratio	63d, 252d (lag 5)	
	Directional persist.	63d, 252d (lag 1)	
X-Sec.	XS Vol. Disp.	63d, 126d	Over $\mathcal{U}_{\text{xsec}}$
	Avg Pairwise Corr.	63d, 126d	
	Corr. Dispersion	63d, 126d	
	MA Breadth	MA=200d	
	Mean TSMOM	63d, 126d	
	Mean Mom. Accel.	inner 21 / outer 126	
	XS Mom. Spread	63d, 126d	Top/Bot $k = 3$
	XS Mom. Disp.	63d, 126d	

B. Full Grid Search Configuration Space

This appendix enumerates the complete search space used in the primary and trading sweeps. The primary sweep fixes the trading set to `t21_broad_sso_def` and varies all other dimensions. The trading sweep crosses the primary sweep dimensions with all 64 trading sets.

B.1 Data ETF Configurations

Each configuration specifies the ETF universe used to compute per-asset features that drive the HMM. Cross-sectional features are always computed over the fixed `xsec_etfs` universe regardless of data ETF config.

ID	ETFs	Economic Rationale
d01	EEM, TLT, XLK	Growth + rates + EM
d03	LQD, TLT, XLU	Defensives + credit
d04	IWM, XLF, XLK	Growth + financial + small cap
d05	EEM, TLT, XLE	Commodity + EM + rates
d06	TLT, XLP, XLV	Defensive + healthcare + rates
d07	EEM, XLK, XLU	Growth + utilities + EM
d09	EEM, XLB, XLE	Materials + energy + EM
d10	TLT, XLP, XLY	Cyclical vs. defensive + rates
d11	EEM, TLT, XLF, XLK	Core growth + financial + rates + EM
d13	IWM, TLT, XLK, XLE	Growth + energy + rates + small cap
d16	EEM, TLT, XLB, XLI	Cyclical + rates + EM
d18	EEM, TLT, XLK, XLY	Consumer cycle + tech + EM + rates
d23	EEM, TLT, XLK, XLE, XLU	Growth + energy + utilities + EM
d26	LQD, TLT, XLP, XLU, XLV	Pure defensive + credit
d33	EEM, IWM, TLT, XLF, XLK, XLP	Growth + defensive + small cap
d34	EEM, LQD, TLT, XLP, XLU, XLV	All-defensive + credit + EM
d36	EEM, TLT, XLB, XLF, XLI, XLK	Materials + industrial + growth

Table 8: Data ETF configurations used in the trading grid search. Configurations containing poor results in the primary sweep were dropped in the trading sweep due to poor performance.

B.2 Feature Subsets

Subset Name	**Feature Blocks and Windows**
`vol_short`	Volatility: realized (21d, 63d), Parkinson (21d, 63d), downside semi (21d, 63d), VoV (inner21/outer63)
`vol_full`	Volatility: realized (21d, 63d, 126d), Parkinson (21d, 63d, 126d), downside semi (21d, 63d, 126d), VoV (inner63/outer252)
`mom_short`	Momentum: TSMOM (63d), accel (inner21/outer126), residual (lb63)
`mom_long`	Momentum: TSMOM (126d), accel (inner63/outer252), residual (lb126)
`mom_full`	Momentum: TSMOM (63d, 126d), accel (both windows), residual (lb63, lb126)
`trend_short`	Trend: variance ratio (63d), directional persistence (63d), trend regression t-stat (126d)
`trend_long`	Trend: variance ratio (252d), directional persistence (252d), trend regression t-stat (252d), Hurst (126d)
`trend_full`	Trend: all above windows combined
`xsec_full`	Cross-sectional: all 8 features, windows 63d and 126d
`trend_mom`	Trend (full) + Momentum: TSMOM (63d, 126d), residual (lb63)
`mom_trend_xsec`	Momentum (core) + Trend (full) + XS dispersion features
`vol_mom_trend_xsec`	Volatility (core+VoV) + Momentum (core) + Trend (full) + XS dispersion
`kitchen_sink`	All four blocks, all windows (minus VoV and Hurst for memory constraints)

Table 9: Feature subsets used in the grid search. Many sets were dropped after the primary sweep for having no Sharpe lift and no statistical validation.

B.3 Trading ETF Sets

Key	N	ETFs
t01_lean_growth	4	XLK, TLT, EEM, IWM
t02_lev_growth	4	XLK, TLT, SSO, EEM
t03_defensive	4	XLU, TLT, LQD, EEM
t04_safe_haven	4	XLK, XLU, TLT, GLD
t05_fin_small	4	XLF, TLT, EEM, IWM
t06_pure_def	4	XLP, TLT, LQD, GLD
t07_energy_em	4	XLK, XLE, TLT, EEM
t08_health_credit	4	XLV, TLT, LQD, IWM
t09_core_lev	5	XLK, XLU, TLT, EEM, SSO
t10_core_nolev	5	XLK, XLU, TLT, EEM, IWM
t11_credit_gold	5	XLK, TLT, EEM, LQD, GLD
t12_fin_lev	5	XLF, XLK, TLT, SSO, EEM
t13_def_credit	5	XLU, XLP, TLT, LQD, EEM
t14_energy_gold	5	XLK, XLE, TLT, EEM, GLD
t15_growth_def	5	XLK, XLU, XLP, TLT, IWM
t16_health_def	5	XLV, XLU, TLT, LQD, EEM
t17_lev_gold	5	XLK, TLT, SSO, GLD, EEM
t18_fin_util	5	XLF, XLU, TLT, EEM, IWM
t19_consumer_small	5	XLK, XLY, TLT, EEM, IWM
t20_energy_credit	5	XLE, XLK, TLT, EEM, LQD
t21_broad_sso_def	6	XLK, XLU, XLP, TLT, EEM, SSO
t22_broad_energy	6	XLK, XLU, TLT, EEM, SSO, XLE
t23_fin_nolev	6	XLK, XLF, XLU, TLT, EEM, IWM
t24_def_credit_nolev	6	XLK, XLU, XLP, TLT, LQD, EEM
t25_health_lev	6	XLK, XLV, XLU, TLT, EEM, SSO
t26_energy_gold6	6	XLK, XLE, XLU, TLT, EEM, GLD
t27_fin_lev_gold	6	XLF, XLK, TLT, SSO, EEM, GLD
t28_def_gold	6	XLK, XLU, XLP, TLT, EEM, GLD
t29_health_small	6	XLV, XLP, XLU, TLT, LQD, IWM
t30_consumer_fin	6	XLK, XLY, XLF, TLT, EEM, IWM

Key	N	ETFs
t31_safe_havens	6	XLK, XLU, TLT, EEM, LQD, GLD
t32_mat_energy_lev	6	XLB, XLE, XLK, TLT, EEM, SSO
t33_fin_lev_credit	6	XLK, XLF, XLU, TLT, SSO, LQD
t34_health_credit6	6	XLK, XLV, TLT, EEM, IWM, LQD
t35_broad7_lev	7	XLK, XLF, XLU, XLP, TLT, EEM, SSO
t36_broad7_nolev	7	XLK, XLU, XLP, TLT, EEM, IWM, LQD
t37_health_lev_gold	7	XLK, XLV, XLU, TLT, EEM, SSO, GLD
t38_energy_lev_small	7	XLK, XLE, XLU, TLT, EEM, SSO, IWM
t39_fin_credit_small	7	XLF, XLK, XLU, TLT, LQD, EEM, IWM
t40_consumer_lev_gld	7	XLK, XLY, XLP, TLT, EEM, SSO, GLD
t41_health_def_credit	7	XLK, XLV, XLP, TLT, EEM, LQD, IWM
t42_energy_gold_cred	7	XLK, XLU, XLE, TLT, EEM, GLD, LQD
t43_fin_lev_credit7	7	XLF, XLK, XLP, TLT, SSO, EEM, LQD
t44_mat_energy_small	7	XLK, XLB, XLE, TLT, EEM, SSO, IWM
t45_health_gold_small	7	XLV, XLK, XLU, TLT, EEM, IWM, GLD
t46_consumer_health	7	XLK, XLY, XLV, TLT, EEM, SSO, LQD
t47_broad8_gold	8	XLK, XLF, XLU, XLP, TLT, EEM, SSO, GLD
t48_broad8_nolev	8	XLK, XLF, XLU, XLP, XLV, TLT, EEM, IWM
t49_def_lev_gold_cred	8	XLK, XLU, XLP, TLT, EEM, SSO, LQD, GLD
t50_energy_fin_lev	8	XLK, XLE, XLF, XLU, TLT, EEM, SSO, IWM
t51_health_def_cred8	8	XLK, XLV, XLU, XLP, TLT, LQD, EEM, IWM
t52_consumer_fin_gld	8	XLK, XLY, XLF, XLU, TLT, EEM, SSO, GLD
t53_mat_fin_cred	8	XLK, XLB, XLF, XLU, TLT, EEM, LQD, IWM
t54_energy_fin_gold	8	XLF, XLE, XLK, XLP, TLT, EEM, SSO, GLD
t55_health_energy_crd	8	XLK, XLV, XLE, XLU, TLT, EEM, LQD, IWM
t56_consumer_health8	8	XLK, XLY, XLP, XLV, TLT, EEM, SSO, LQD
t57_broad10_credit	10	XLK, XLF, XLU, XLP, XLV, TLT, EEM, SSO, LQD, IWM
t58_broad10_energy	10	XLK, XLF, XLE, XLU, XLP, TLT, EEM, SSO, GLD, IWM

Key	N	ETFs
t59_broad11	11	XLK, XLF, XLU, XLP, XLV, XLE, TLT, EEM, SSO, LQD, IWM
t60_full12	12	XLK, XLF, XLU, XLP, XLV, XLE, XLB, TLT, EEM, SSO, GLD, IWM
t61_consumer10	10	XLK, XLY, XLV, XLF, XLU, TLT, EEM, SSO, LQD, IWM
t62_cyclical10	10	XLK, XLF, XLE, XLB, XLI, TLT, EEM, SSO, GLD, LQD
t63_defensive10	10	XLV, XLP, XLU, XLF, XLK, TLT, LQD, EEM, IWM, GLD
t64_realestate12	12	XLK, XLY, XLF, XLE, XLU, TLT, EEM, SSO, GLD, LQD, IWM, VNQ

Table 10: All 64 trading ETF sets. N denotes the number of assets in the portfolio universe. SSO is the 2× leveraged S&P 500 ETF; GLD is gold; VNQ is real estate.

B.4 Sweep Constants

Parameter	Value
N_{states}	$\{2, 3, 4, 5\}$
PCA modes	{flat, block}
Test / step months	$\{(12, 12)\}$
Training years (minimum)	5
HMM iterations (`n_iter`)	500
HMM random fits (`n_fits`)	3
Covariance type	Full
Bootstrap samples (`n_boot`)	100
Block size (bootstrap)	5 days
Flat PCA components	10
Block PCA xsec components	3
Block PCA other components	1
Start date	2006-06-21

Table 11: Fixed hyperparameters used across all grid search configurations.

C. Walk-Forward Pipeline and Window Details

C.1 End-to-End Pipeline Description

The walk-forward pipeline is designed with strict no-lookahead guarantees at every stage. The full sequence of operations is as follows.

Stage 1: Raw Data Acquisition. Daily OHLC prices and log returns are pulled from Polygon.io for the full ETF universe over the period 2006-06-21 to the most recent completed trading day. The end date is capped to the prior business day at pull time to exclude partial intraday bars. All data is stored as a single flat DataFrame indexed by date with columns `{ticker}_open`, `{ticker}_high`, `{ticker}_low`, `{ticker}_close`, `{ticker}_return`.

Stage 2: Raw Feature Computation. All features described in Section 4 are computed over the full date range using only past data at each point in time (each rolling computation uses only data up to and including day t). This produces a single unnormalized feature matrix $\mathbf{F} \in \mathbb{R}^{T\times p}$ which is saved to disk as `raw_features.csv`. Critically, *no normalization is applied at this stage.* The raw features are computed once and reused across all walk-forward windows; normalization is deferred to within each training fold.

Stage 3: Window Generation. The walk-forward schedule partitions the full date range into a sequence of non-overlapping train/test splits. With training parameter `train_years`$= 5$, `test_months`$= 12$, and `step_months`$= 12$, the n-th window is defined as:

$$
\begin{aligned}
\text{train\_start}^{(n)} &= t_0 \quad \text{(fixed, expanding)} \\
\text{train\_end}^{(n)} &= t_0 + 5 \text{ yr} + (n-1) \times 12 \text{ mo} \\
\text{test\_start}^{(n)} &= \text{train\_end}^{(n)} + 1 \text{ day} \\
\text{test\_end}^{(n)} &= \text{test\_start}^{(n)} + 12 \text{ mo} - 1 \text{ day}
\end{aligned}
$$

Test windows are adjacent and non-overlapping by construction, so that concatenating all test-window outputs produces a single contiguous out-of-sample series with no gaps. For example, with $t_0 = $ 2007-09-19:

Window	Train Start	Train End	Test Start	Test End	N Train
0	2007-09-19	2012-09-19	2012-09-20	2013-09-19	1262
1	2007-09-19	2013-09-19	2013-09-20	2014-09-19	1512
2	2007-09-19	2014-09-19	2014-09-22	2015-09-19	1764
3	2007-09-19	2015-09-19	2015-09-21	2016-09-19	2015
⋮	⋮	⋮	⋮	⋮	⋮
13	2007-09-19	2025-09-19	2025-09-22	2026-04-14	4530

Table 12: Walk-forward window schedule for the 12-month test/step configuration. The training window expands by one year at each step (expanding window design). Test windows are exactly adjacent: the test start of window $n+1$ is the calendar day after the test end of window n.

Stage 4: Per-Window Processing. For each window n, the following operations are performed in strict sequence:

1. **Feature selection**: Columns of $\mathbf{F}$ corresponding to the configured feature subset and data ETFs are selected, yielding $\mathbf{F}^{(n)} \in \mathbb{R}^{T \times p_n}$.

2. **Train/test split**: The feature matrix is split into $\mathbf{F}^{(n)}_{\text{train}}$ (rows in $[\text{train\_start}^{(n)}, \text{train\_end}^{(n)}]$) and $\mathbf{F}^{(n)}_{\text{test}}$ (rows in $[\text{test\_start}^{(n)}, \text{test\_end}^{(n)}]$).

3. **Feature normalization**: Compute column-wise mean $\hat{\boldsymbol{\mu}}_{\text{tr}}$ and standard deviation $\hat{\boldsymbol{\sigma}}_{\text{tr}}$ from $\mathbf{F}^{(n)}_{\text{train}}$ only. Apply to both splits:

$$\tilde{\mathbf{F}}^{(n)}_{\text{train}} = \frac{\mathbf{F}^{(n)}_{\text{train}} - \hat{\boldsymbol{\mu}}_{\text{tr}}}{\hat{\boldsymbol{\sigma}}_{\text{tr}}}, \quad \tilde{\mathbf{F}}^{(n)}_{\text{test}} = \frac{\mathbf{F}^{(n)}_{\text{test}} - \hat{\boldsymbol{\mu}}_{\text{tr}}}{\hat{\boldsymbol{\sigma}}_{\text{tr}}}$$

4. **PCA fit and transform**: Fit PCA (flat or block, per configuration) on $\tilde{\mathbf{F}}^{(n)}_{\text{train}}$ only. Apply the fitted transform to both splits:

$$\mathbf{Z}^{(n)}_{\text{train}} = \text{PCA}\left(\tilde{\mathbf{F}}^{(n)}_{\text{train}}\right), \quad \mathbf{Z}^{(n)}_{\text{test}} = \text{PCA}_{\text{fitted}}\left(\tilde{\mathbf{F}}^{(n)}_{\text{test}}\right)$$

5. **PC normalization**: Normalize the principal components using statistics from the training PCs only:

$$\hat{\mathbf{Z}}^{(n)}_{\text{train}} = \frac{\mathbf{Z}^{(n)}_{\text{train}} - \hat{\boldsymbol{\mu}}_{\text{PC,tr}}}{\hat{\boldsymbol{\sigma}}_{\text{PC,tr}}}, \quad \hat{\mathbf{Z}}^{(n)}_{\text{test}} = \frac{\mathbf{Z}^{(n)}_{\text{test}} - \hat{\boldsymbol{\mu}}_{\text{PC,tr}}}{\hat{\boldsymbol{\sigma}}_{\text{PC,tr}}}$$

6. **HMM training**: Fit the Gaussian HMM with K states on $\hat{\mathbf{Z}}^{(n)}_{\text{train}}$ using Baum-Welch EM (uses hmmlearn.GaussianHMM).

7. **Walk-forward decoding**: Run the fitted HMM on $\hat{\mathbf{Z}}^{(n)}_{\text{test}}$ to produce posterior probabilities $\gamma_t(k)$ and Viterbi labels z_t^* for all test-window dates.

The outputs of step 7 across all windows are concatenated to produce the full OOS regime series used for portfolio construction and evaluation.

Remark 1. *At no point do any statistics computed on the test window (or any future data) influence the normalization, PCA, or HMM parameters. The only information flowing from train to test is the fitted normalization parameters, PCA components, and HMM parameters—all estimated exclusively from* $\mathbf{F}^{(n)}_{train}$.

C.2 Full Window Table

Win.	Train Start	Train End	Test Start	Test End	N Train	N Test
0	2007-09-19	2012-09-19	2012-09-20	2013-09-19	1262	250
1	2007-09-19	2013-09-19	2013-09-20	2014-09-19	1512	252
2	2007-09-19	2014-09-19	2014-09-22	2015-09-19	1764	251
3	2007-09-19	2015-09-19	2015-09-21	2016-09-19	2015	252
4	2007-09-19	2016-09-19	2016-09-20	2017-09-19	2267	252
5	2007-09-19	2017-09-19	2017-09-20	2018-09-19	2519	252
6	2007-09-19	2018-09-19	2018-09-20	2019-09-19	2771	251
7	2007-09-19	2019-09-19	2019-09-20	2020-09-19	3022	252
8	2007-09-19	2020-09-19	2020-09-21	2021-09-19	3274	251
9	2007-09-19	2021-09-19	2021-09-20	2022-09-19	3525	252
10	2007-09-19	2022-09-19	2022-09-20	2023-09-19	3777	251
11	2007-09-19	2023-09-19	2023-09-20	2024-09-19	4028	252
12	2007-09-19	2024-09-19	2024-09-20	2025-09-19	4280	250
13	2007-09-19	2025-09-19	2025-09-22	2026-04-14	4530	141

Table 13: Complete walk-forward window schedule for the 12-month test/step configuration (best configuration data ETF: EEM, TLT, XLK). Window 13 is a partial year as it extends to the current date at time of writing. The training window grows by approximately 252 trading days per step (expanding design).

D. Software Stack and Reproducibility

D.1 Compute Infrastructure

All grid search computations were performed on the PACE ICE high-performance computing cluster at Georgia Institute of Technology. Jobs were submitted via SLURM using a rolling pool of up to 490 concurrent single-core jobs. Each job processed one configuration with 16GB of allocated memory and a 2-hour wall time limit.

Component	**Specification**
Cluster	PACE ICE, Georgia Tech
Scheduler	SLURM
Concurrent jobs	490 (rolling pool)
Memory per job	16 GB
Wall time per job	2 hours
Total configurations	84,864 (primary + trading sweep)

Table 14: Compute infrastructure specifications.

D.2 Python Environment

Package	**Version**
Python	3.10.x+
`hmmlearn`	0.3.3
`scikit-learn`	1.7.2
`pandas`	2.2.x
`numpy`	1.26.x+
`polygon-api-client`	1.16.3
`matplotlib`	3.10.6
`seaborn`	0.13.2
`scipy`	1.16.2
`cvxpy`	1.7.x

Table 15: Python package versions. The `polygon-api-client` version is pinned to 1.16.3 across all environments to ensure data reproducibility; a version discrepancy between 1.15.4 and 1.16.3 was found to produce systematically different intraday low prices, affecting feature computation.

D.3 Data Reproducibility

- All raw price data is cached to `asset_data.csv` and `market_data.csv` at pull time. Subsequent runs load from cache rather than re-pulling, ensuring identical data across all configurations in a given sweep.
- The end date for all data pulls is capped to the prior business day (`pd.Timestamp.today() - BDay(1)`) to exclude partial intraday bars that vary depending on the time of day the pull is executed.
- Random seed: `random_state=42` is passed to the HMM initializer. However, multiple random initializations (`n_fits=3`) are used and the best is selected by log-likelihood, so the final model is deterministic conditional on the data but not on a single seed.

References

Hamilton, J.D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.

Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *Review of Financial Studies*, 15(4), 1137–1187.

Guidolin, M., & Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11), 3503–3544.

Kritzman, M., Page, S., & Turkington, D. (2012). Regime shifts: Implications for dynamic strategies. *Financial Analysts Journal*, 68(3), 22–39.

Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411.

Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, 34(4), 768–783.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.

Baum, L.E., Petrie, T., Soules, G., & Weiss, N. (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Annals of Mathematical Statistics*, 41(1), 164–171.

Moskowitz, T.J., Ooi, Y.H., & Pedersen, L.H. (2012). Time series momentum. *Journal of Financial Economics*, 104(2), 228–250.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business*, 53(1), 61–65.