POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Regime-Switching Momentum: Robust Portfolio Selection among S&P500 Constituents

Laurea Magistrale in Mathematical Engineering - Ingegneria Matematica

Author: Alessandro Ciotti

Advisor: Prof. Daniele Marazzina

Co-advisors: Dott. Giorgio Bortolozzo, Dott. Marco Corradini

Academic year: 2025-2026

1. Introduction

The identification of structural changes in financial market dynamics represents a central challenge in developing investment strategies that are both adaptive and resilient across economic cycles. Classical asset pricing theory, rooted in the Efficient Market Hypothesis (EMH), claims that prices instantaneously incorporate all available information, rendering systematic excess returns impossible. Yet a substantial body of empirical literature documents persistent market anomalies. Among these, the Momentum effect – the tendency of recently outperforming assets to continue outperforming in the short run – constitutes one of the most robust challenges to the informational efficiency paradigm.

Despite its documented profitability, classical momentum is vulnerable to periodic and severe reversals known as momentum crashes, which arise during market recovery phases and regime transitions [3]. These drawdowns come from the strategy's purely backward-looking nature and its susceptibility to the inherent non-stationarity of financial markets: return distributions and dependency structures shift dramatically between calm and crisis regimes.

This thesis addresses these limitations by proposing a quantitative framework that integrates Hidden Markov Models (HMMs) into a stock selection procedure applied to S&P 500 constituents over the 2000–2026 period. The core methodological innovation lies in moving beyond univariate regime detection: the model combines stock-specific probabilistic signals and sector-level information – classified according to the GICS classification – through the temporal evolution of latent regime distributions, yielding dynamic ranking metrics and risk management filters.

2. Theoretical Background

2.1. Momentum Strategies

The seminal contribution of Jegadeesh and Titman [7] established the cross-sectional momentum benchmark by ranking assets based on cumulative past returns over a formation window J, maintaining long positions in top-decile "winners" for a holding period K. The strategy employs overlapping portfolios to reduce turnover impact and improve temporal diversification.

A well-documented structural weakness is the time-varying beta of the strategy's legs. As analyzed by Grundy and Martin [4], following market downturns the winner portfolio loads on low-beta stocks while the loser portfolio concentrates in high-beta assets, creating a negatively-exposed long-short position that suffers severely during sharp market rebounds. Furthermore, Novy-Marx and Velikov [9] estimate that implementation costs erode momentum's net returns by 1–3% annually, with the erosion amplified by sector concentration risk and crowding effects.

2.2. Hidden Markov Models in Finance

A Hidden Markov Model is a doubly stochastic process

$$\lambda = (\mathcal{X}, \mathcal{Y}, A, B, \pi) \tag{1}$$

in which an observable sequence $\mathcal{Y} = \{Y_t\}_{t\geq 0}$ is governed by an unobservable latent Markov chain $\mathcal{X} = \{X_t\}_{t\geq 0}$ with state space of cardinality N. The model is defined by the transition matrix $A = \{a_{ij}\}$, emission distributions $B = \{b_j(\cdot)\}$, and initial distribution π. In the Gaussian case: $Y_t \mid X_t = j \sim \mathcal{N}(\mu_j, \sigma_j^2)$.
Parameter estimation is performed via the Baum-Welch algorithm [2], an Expectation-Maximization procedure that iteratively maximizes the likelihood $P(\mathcal{Y} \mid \lambda)$. State inference relies on the Forward-Backward algorithm, which computes both posterior probabilities $\gamma_t(i) = P(X_t = s_i \mid \mathcal{Y}, \lambda)$ for historical analysis, and filtered probabilities $P(X_t = s_i \mid y_1, \dots, y_t, \lambda)$ for real-time, look-ahead-free applications. HMMs have been applied to equity regime detection [1], volatility modeling [5], and more recently to stock selection [8].

3. Data and Methodology

3.1. Data

The study employs the full point-in-time universe of S&P 500 constituents from January 2000 to December 2025, reconstructed quarterly to prevent look-ahead bias. Weekly dividend-adjusted prices and market capitalizations are collected, and each firm is assigned its current GICS first-level sector classification. Over the sample period, Information Technology grew from a moderate share to approximately 35% of total market capitalization by end-2025, with five sectors (IT, Financials, Health Care, Communication Services, Consumer Discretionary) jointly representing nearly 80% of index value.

Analysis of sector-level risk profiles reveals significant heterogeneity: defensive sectors such as Health Care (CAGR 11.79%, Sharpe 0.80, volatility 15.48%) and Consumer Staples (CAGR 9.40%, Sharpe 0.75, volatility 13.16%) substantially outperform cyclical sectors such as Energy (vol. 30.30%) and IT (vol. 28.10%) on a risk-adjusted basis over the full horizon. Intra-sector pairwise correlations increase markedly during crisis episodes (2008–09, 2020), reducing diversification benefits precisely when they are most needed.

3.2. Proposed Strategy

The strategy conditions momentum signals on latent market regimes, exploiting the empirically documented persistence of expansionary phases. Gaussian HMMs with $N \in \{2, 3\}$ states are estimated at the sector level using a five-year rolling window of constituent return series. To ensure robust estimation, only stocks with at least five years of continuous price history are included in the training universe.

Nowcasting – Stock Ranking. For each stock i and each regime h, regime-conditional expected returns $\mu_i^{(h)}$ and volatilities $\sigma_i^{(h)}$ are computed as probability-weighted averages over the period. An Exponential Moving Average (EMA) with smoothing factor $\alpha = 2/(J+1)$, with $J = 12$, is applied to the posterior probabilities $\pi_t^{(h)}$ to obtain a regime-inertia distribution $\widehat{EMA}_T^{(h)}$, formalizing the hypothesis that persistent regimes carry predictive weight. The composite expected Sharpe ratio is:

$$\widehat{SR}_{i,T} = \frac{\hat{\mu}_{i,T}}{\hat{\sigma}_{i,T}} \tag{2}$$

where $\hat{\mu}_{i,T}$ and $\hat{\sigma}_{i,T}$ represent respectively the expected mean and volatility of stock i computed with the regime-inertia distribution $\widehat{EMA}_T$. Stocks in the top decile of $\widehat{\mathrm{SR}}_{i,T}$ are selected for the long-only portfolio.

Forecasting – Regime Filter. A sector-level risk filter is computed by projecting the current regime distribution π_T forward over the holding period $K = 6$ months using the transition matrix M (monthly, derived from the weekly matrix as $M = W^4$). The forecast probability is the harmonic mean across the horizon:

$$\pi_{T+K}^{(h)} = \frac{1}{\sum_{k=1}^{K} k^{-1}} \sum_{k=1}^{K} \frac{1}{k} \left[\pi_T \cdot M^k\right]_h . \tag{3}$$

Sector s is included in the portfolio only if $\pi_{T+K}^{\text{Growth}} > \gamma$, where three threshold specifications are examined:

- Unfiltered: $\gamma = 0$
- Naive: $\gamma = 1/N$
- Stationary: $\gamma = \tilde{\pi}^{\text{Growth}}$, the long-run stationary probability

When the condition is not met, the corresponding allocation is held in cash rather than redistributed, providing a systematic defensive mechanism.

Portfolio rebalancing follows the overlapping structure of [7] with $K = 6$ months. Transaction costs are modeled as

$$TC_t = c \cdot \sum_i S_{i,t} |Q_{i,t} - Q_{i,t-1}| \tag{4}$$

with $c \in [5, 20]$ bps. The short leg is intentionally excluded, consistent with evidence that momentum crashes concentrate in the short portfolio [3] and that the long-leg effect dominates among large-cap liquid stocks [6].

4. Empirical Results

4.1. Regime Characterization

Two-state models consistently identify a stress regime with strongly negative mean returns and high volatility (e.g., IT: $\mu_0 = -19.03\%$, $\sigma_0 = 38.36\%$) and an expansionary regime with positive returns and subdued volatility ($\mu_1 = 26.01\%$, $\sigma_1 = 15.29\%$). Regime persistence is high across all sectors ($P_{11} \approx 94\%$), supporting the hypothesis of regime inertia. Three-state models further differentiate a bear, an intermediate/mean-reverting, and a bull state, enabling finer segmentation at the cost of increased sensitivity to noise.

4.2. Performance Analysis

Table 1 reports annualized out-of-sample performance metrics net of 15 bps transaction costs. All HMM-based strategies outperform the classical momentum benchmark (MOM: CAGR 5.76%, Sharpe 0.57) both in terms of absolute returns and risk-adjusted efficiency.

Table 1: Annual Performance: HMM Strategies vs. Benchmark.

Model	γ	**CAGR**	**Sharpe**	α
2-State	0%	7.28%	0.68	2.62%
	50%	6.71%	0.67	2.56%
	$\tilde{\pi}$	6.13%	0.80	3.87%
3-State	0%	7.16%	0.67	2.44%
	33.3%	6.54%	0.68	2.62%
	$\tilde{\pi}$	5.78%	0.81	3.82%
MOM	–	5.76%	0.57	–

As the threshold γ increases, CAGR and volatility decrease while the Sharpe ratio improves significantly, reaching 0.80–0.81 under the Stationary specification. Market beta decreases from ≈ 0.96 (Unfiltered) to ≈ 0.51–0.57 (Stationary), indicating that tighter filters yield more market-independent return streams – an attractive property for multi-asset portfolio integration.

4.3. Risk Management

The regime-filtering mechanism produces substantial improvements across all tail-risk metrics (Table 2). Maximum drawdown is reduced from approximately −57% (benchmark) to −19.9% under the 3-state Stationary strategy. CVaR at 95% confidence improves from −12.55% to −6.12%. Return distributions are less negatively skewed and exhibit lower excess kurtosis than the benchmark, confirming a systematic reduction in downside tail risk.

Table 2: Risk Metrics vs. Benchmark (15 bps costs).

Model	γ	**MDD**	$\mathbf{CVaR_{95\%}}$
2-State	0%	−55.97%	−11.42%
	$\tilde{\pi}$	−23.34%	−6.56%
3-State	0%	−56.35%	−11.67%
	$\tilde{\pi}$	−19.92%	−6.12%
MOM	–	−57.14%	−12.55%

4.4. Sector Dynamics and Transaction Costs

Sector exposure analysis reveals that the HMM framework implicitly encodes a flight-to-quality rotation: cyclical sectors (e.g., Consumer Discretionary) decrease in weight during stress phases while defensive sectors (e.g., Health Care) gain allocation – without any explicit sector constraints. This emergent behavior confirms that regime-conditioned selection captures macroeconomic cycle dynamics.

Transaction cost analysis demonstrates that HMM strategies are significantly more efficient than the classical benchmark. Under a 20 bps cost scenario, the benchmark suffers an annualized alpha decay of -63 bps, compared to -30 to -34 bps for the proposed strategies. This advantage reflects the greater stability of the regime-conditioned ranking signal, which produces lower portfolio turnover.

5. Conclusions

This thesis demonstrates that integrating Hidden Markov Models into a momentum-based stock selection framework represents a robust and practically viable enhancement over classical strategies. By conditioning portfolio construction on latent regime probabilities – rather than relying solely on backward-looking returns – the proposed approach achieves three simultaneous improvements: higher risk-adjusted performance (Sharpe ratio up to 0.81 vs. 0.57 for the benchmark), substantially reduced tail risk (MDD from -57% to -20%), and lower sensitivity to transaction costs (alpha decay reduced by approximately 50%).

The regime-filtering mechanism directly addresses the momentum crash problem by dynamically contracting market exposure during adverse conditions. The sector rotation patterns induced by the model reflect genuine macroeconomic regime dynamics, lending economic credibility to the statistical framework. Operationally, the strategy's lower turnover offers meaningful cost advantages relative to traditional implementations.

Future developments include: (i) adoption of Student-t emission distributions to better capture heavy tails; (ii) extension to Hidden Semi-Markov Models to relax the geometric sojourn-time assumption; (iii) dynamic selection of the number of states guided by economic interpretability criteria; and (iv) integration with macroeconomic variables to improve detection of gradual regime transitions such as the 2022 inflationary bear market.

References

[1] Andrew Ang and Geert Bekaert. International asset allocation with regime shifts. *Review of Financial Studies*, 15(4):1137–1187, 2002.

[2] Leonard E. Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Annals of Mathematical Statistics*, 41(1):164–171, 1970.

[3] Kent Daniel and Tobias J. Moskowitz. Momentum crashes. *Journal of Financial Economics*, 122(2):221–247, 2016.

[4] Bruce D. Grundy and J. Spencer Martin. Understanding the nature of the risks and the source of the rewards to momentum investing. *Review of Financial Studies*, 14(1):29–78, 2001.

[5] James D. Hamilton and Raul Susmel. Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64(1-2):307–333, 1994.

[6] Harrison Hong, Terence Lim, and Jeremy C. Stein. Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. *The Journal of Finance*, 55(1):265–295, 2000.

[7] Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1):65–91, 1993.

[8] Nguyet Nguyen and Dung Nguyen. Hidden markov model for stock selection. *Risks*, 3(4):455–473, 2015.

[9] Robert Novy-Marx and Mihail Velikov. A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1):104–147, 2016.