



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

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

TESI DI LAUREA MAGISTRALE IN
MATHEMATICAL ENGINEERING - INGEGNERIA MATEMATICA

Author: **Alessandro Ciotti**

Student ID: 947754

Advisor: Prof. Daniele Marazzina

Co-advisors: Dott. Giorgio Bortolozzo – ARCA Fondi SGR

Dott. Marco Corradini – ARCA Fondi SGR

Academic Year: 2025-2026

Abstract

The identification of structural changes in financial market dynamics represents a central challenge in the development of adaptive and robust investment strategies. In this context, this thesis proposes a quantitative framework based on Hidden Markov Model (HMM), applied to the stocks belonging to the S&P 500 index over the period 2000–2026. The model allows for the identification of latent market regimes and the derivation of information conditioned on these states, both at the individual stock and sector levels, according to the GICS classification. This information is then combined through the temporal evolution of the regimes to construct dynamic ranking metrics capable of adapting to different market contexts. This procedure exhibits superior predictive capacity compared to the reference benchmark, producing higher cumulative and risk-adjusted returns. To improve the portfolio allocation process, a filtering mechanism based on probability thresholds associated with the identification of favorable regimes is introduced. This approach allows for a dynamic reduction in exposure during periods characterized by high uncertainty or adverse conditions, allowing more efficient risk management. The results of the analysis highlight how the proposed framework is able to significantly mitigate systemic risk, substantially improving key risk metrics. Furthermore, the integration of signals conditioned on latent regimes promotes the construction of more resilient portfolios compared to traditional benchmarks: the proposed strategy indeed shows greater efficiency in terms of transaction costs, with a significantly lower alpha decay compared to the classic momentum strategy. Overall, the study demonstrates how the joint use of probabilistic models and cross-sectional information can represent an effective approach for the construction of dynamic investment strategies capable of adapting to the evolution of market regimes.

Keywords: Momentum Factor, Hidden Markov Model, Volatility Clustering, Market Regimes, Risk Management, GICS sectors

Abstract in lingua italiana

L'identificazione dei cambiamenti strutturali nelle dinamiche dei mercati finanziari rappresenta una sfida centrale nello sviluppo di strategie di investimento adattive e robuste. In questo contesto, la presente tesi propone un framework quantitativo basato sul Modello di Markov a Stati Latenti (Hidden Markov Model, HMM), applicato ai titoli azionari appartenenti all'indice S&P 500 nel periodo 2000–2006. Il modello consente di identificare regimi di mercato latenti e di derivare informazioni condizionate a tali stati, sia a livello di singolo titolo sia di settore, secondo la classificazione GICS. Queste informazioni vengono quindi combinate tra loro attraverso l'evoluzione temporale dei regimi per costruire metriche di ranking dinamiche, in grado di adattarsi ai diversi contesti di mercato. Tale procedura evidenzia una capacità predittiva superiore rispetto al benchmark di riferimento, producendo rendimenti cumulati e risk-adjusted più elevati. Al fine di migliorare il processo di allocazione del portafoglio, viene introdotto un meccanismo di filtraggio basato su soglie di probabilità associate all'identificazione di regimi favorevoli. Questo approccio permette di ridurre dinamicamente l'esposizione nei periodi caratterizzati da elevata incertezza o condizioni avverse, favorendo una gestione più efficiente del rischio. I risultati dell'analisi evidenziano come il framework proposto sia in grado di mitigare in modo significativo il rischio sistemico, migliorando sensibilmente le principali metriche di rischio. Inoltre, l'integrazione di segnali condizionati ai regimi latenti favorisce la costruzione di portafogli più resilienti rispetto ai benchmark tradizionali: la strategia proposta mostra infatti una maggiore efficienza in termini di costi di transazione, con un decadimento dell'alpha sensibilmente inferiore rispetto alla classica strategia momentum. Nel complesso, lo studio dimostra come l'utilizzo congiunto di modelli probabilistici e informazioni cross-sectional possa rappresentare un approccio efficace per la costruzione di strategie di investimento dinamiche, capaci di adattarsi all'evoluzione dei regimi di mercato.

Parole chiave: Fattore Momentum, Modello di Markov a Stati Latenti, Clustering di Volatilità, Regimi di Mercato, Gestione del Rischio, Settori GICS

Foreword

Momentum is one of the most widely studied phenomena in finance and, at the same time, one of the most difficult to translate into a robust investment strategy. Its persistence coexists with sudden and often violent reversals, which tend to cluster precisely during post-crisis rebounds—when the impact of losses weighs most heavily on portfolio managers. This is a challenge well known to market practitioners, and it is the starting point of Alessandro Ciotti's work.

What sets the approach adopted in this thesis apart is the choice to eschew the more conventional path. While literature has already explored the use of Hidden Markov Models to identify market regimes – almost always utilizing them as a filter to activate or deactivate the momentum signal based on the estimated state – this work takes a different route. Instead, it integrates regime probabilities, propagated through the transition matrix, directly into the criterion used to rank and select securities within the S&P 500 universe. Alessandro distinguishes between two functions: a nowcasting component that estimates the current state to construct the stock ranking, and a forecasting component that projects regime dynamics forward to adjust overall exposure. This methodological distinction originally intervenes at the very core of portfolio construction.

The thesis combines a solid and coherent mathematical framework – from the treatment of market states to the definition of the final ranking – with a highly professional approach regarding practical applicability. Alessandro reconstructed the investable universe using point-in-time data, evaluated the impact of transaction costs across different scenarios, and maintained unallocated portfolio shares in cash rather than forcing redistribution. Concurrently, the limitations still present within the framework, from the Gaussian assumption to the stationarity of the transition matrix, are openly discussed and serve as a foundation for future analysis.

Together with our colleague Corradini, we had the opportunity to mentor Alessandro during the development of this work at Arca Fondi SGR and to engage with him throughout the various phases of the research.

The topic addressed in these pages touches upon one of the most compelling questions in quantitative finance: how to preserve the informational power of a signal across shifting market regimes. The work presented here offers a valuable contribution to this pursuit, demonstrating how market complexity can be successfully tackled through a rigorous and disciplined approach.

Giorgio Bortolozzo

Head of Equity Investments

Arca Fondi SGR

Contents

Abstract	i
Abstract in lingua italiana	iii
Foreword	v
Contents	vii
Introduction	1
1 Theoretical Foundations of Momentum Investing	3
1.1 Historical Development of Momentum Strategies	3
1.2 Implementation of Momentum Strategies	4
1.3 Market Regimes and Structural Limitations of Classical Momentum	5
1.3.1 Momentum Crashes and Regime Dependency	5
1.3.2 Structural Limitations	6
2 Hidden Markov Models: Theory and Financial Applications	7
2.1 Hidden Markov Models: Definitions and Theoretical Framework	7
2.2 Model Components	8
2.2.1 State Space Design	9
2.2.2 Transition Dynamics	9
2.2.3 Emission Distributions	10
2.3 Estimation and Inference Algorithms	11
2.3.1 Problem I: the Forward-Backward Algorithm	12
2.3.2 Problem II: the Viterbi Algorithm	15
2.3.3 Problem III: the Baum-Welch Algorithm	16
2.4 Applications of HMMs in Finance	17

3	Data and Methodology	19
3.1	Data Description	19
3.2	Proposed Strategy	24
3.2.1	Regime Identification Framework	24
3.2.2	Investment Universe and Benchmark Definition	25
3.2.3	Decision Rule	26
3.2.4	Rebalancing and Transaction Costs	27
3.3	Mathematical Formalism of the Proposed Strategy	28
3.4	Comparison with Existing Literature	30
4	Empirical Results	33
4.1	Implementation Details	33
4.2	Regime Characterization	34
4.2.1	2-States HMMs	34
4.2.2	3-States HMMs	35
4.3	Backtest Results	36
4.3.1	Performance Analysis	37
4.3.2	Risk Management	40
4.3.3	Sector Exposure Analysis	45
4.3.4	Trading Cost Analysis	48
4.4	Model Limitations	49
4.4.1	Return Distributions	49
4.4.2	Stationarity Hypothesis	49
4.4.3	Selection of the Number of States	50
4.4.4	Liquidity Issues	50
5	Conclusions and future developments	53
	Bibliography	55
A	Appendix A	59
A.1	2 States HMMs	59
A.2	3 States HMMs	65
A.3	Class Diagram	71

List of Figures	73
List of Tables	75
Acknowledgements	77

Introduction

The identification of structural changes in financial market dynamics has been a central objective for decades in the construction of investment portfolios resilient to various economic cycles. Although classical theory, based on the Efficient Market Hypothesis (EMH), suggests that prices instantaneously reflect all available information – theoretically rendering the systematic pursuit of excess returns impossible – a growing body of empirical evidence has identified persistent financial "anomalies". Among these, the Momentum effect – the tendency of assets that have performed well in the recent past to continue overperforming in the short term – constitutes one of the most significant challenges to the information efficiency paradigm.

However, implementing momentum-based strategies involves non-negligible intrinsic risks. The literature has extensively documented how these approaches are subject to periodic and violent reversals, known as momentum crashes, which typically occur during post-crisis recovery phases or in correspondence with abrupt shifts in the economic cycle. The vulnerability lies in the inherent non-stationarity of markets: indeed, financial returns are characterized by time-varying distributions and dependency structures that shift drastically between periods of stability and periods of financial panic. To address the problem of non-stationarity, econometric research has proposed the adoption of regime-switching models. Within this field lie Hidden Markov Models (HMMs) that allow for the modeling of the market as a system governed by an unobservable latent state, offering a superior representation of financial reality compared to traditional linear models due to their ability to capture complex phenomena such as volatility clustering and heavy tails in return distributions.

The application of regime-switching architectures in financial economics spans several decades, originating from the seminal work of Hamilton [15] on business cycles and later extended by [1] to capture nonlinear dynamics, volatility clustering, and distinct risk premium environments in equity markets. While HMMs are also used in areas like risk modeling [4, 16] and derivative pricing [6], their employment as systematic engines for cross-sectional stock selection remains relatively scarce. Early contributions in this direction include [19, 20], where macro-driven HMMs are utilized to generate scoring metrics

within broad equity universes. More recently, [26] pioneered the integration of HMMs into momentum investing by using hidden states as binary filters to time factor exposure at the stock, sector, and market levels. However, this existing approach relies on the model purely as a discrete timing device trained exclusively on aggregate sector indices, neglecting the idiosyncratic informational content embedded in individual stock price time series.

To bridge this gap, this thesis integrates the predictive capacity of multi-state HMMs within a stock selection framework applied to the constituents of the S&P 500 Index over the 2000–2026 period. The methodological innovation lies in moving beyond classic univariate approaches: rather than limiting the analysis to macro index dynamics, the model developed here simultaneously processes individual stock-specific price series and sector-level signals structured around the GICS classification, embedding them directly within the probabilistic evolution of the market regimes.

The empirical design transforms these latent state probabilities into a two-step framework. First, a nowcasting layer establishes a robust, real-time ranking metric for individual stocks. Second, a forecasting layer assesses regime stability and filter sector-level eligibility. Crucially, this model introduces an adaptive decision-threshold mechanism. This threshold dynamically triggers capital contraction into non-remunerated cash during high-volatility regimes, transforming a traditionally backward-looking anomaly into a state-aware, risk-managed strategy.

Finally, particular attention is paid to the operational feasibility of the investment strategy. In line with the studies of [23], a comprehensive sensitivity analysis is conducted to evaluate the impact of trading frictions and transaction costs. The empirical results highlight how this Markovian filtering architecture not only optimizes risk-adjusted performance metrics and mitigates tail-risk but also offers superior turnover management, significantly reducing net alpha decay compared to traditional momentum benchmarks.

The work is structured into five chapters. Chapter 1 analyzes the evolution of the literature regarding momentum strategies, identifying their structural weaknesses. Chapter 2 delves into the theory of HMM models, parameter estimation techniques, and the literature concerning their use in finance. Chapter 3 describes the dataset construction methodology and the architecture of the filtering model, integrating a review of application studies related to the research field of this thesis. Chapter 4 presents the results of the empirical tests, comparing the proposed strategy with classic benchmarks and discussing its robustness regarding transaction costs and application limits. Finally, Chapter 5 summarizes the concluding remarks.

1 | Theoretical Foundations of Momentum Investing

1.1. Historical Development of Momentum Strategies

Momentum strategies represent one of the most robust and well-documented empirical anomalies in modern finance. Their development lies at the core of the market efficiency debate, challenging the hypothesis that prices instantaneously incorporate all available information. The literature has extensively demonstrated that past returns possess significant predictive power, giving rise to two distinct methodological approaches: relative momentum and absolute momentum. Relative momentum (cross-sectional momentum, CSM) is based on the relative strength among assets: the strategy involves comparing performance within a defined universe, buying “winner” stocks and selling “loser” stocks. In contrast, absolute momentum (time-series momentum, TSM), often identified with trend-following, evaluates each asset solely based on its own historical performance. The operational divergence between the two approaches becomes particularly evident during bear markets: while absolute momentum may reduce exposure to zero or take short positions if trends are negative, relative momentum remains consistently invested in assets exhibiting the least negative performance (the “relative winners”). Although this exposes CSM to deeper drawdowns, TSM risks missing market rebounds due to its directional nature.

The seminal contribution of Jegadeesh and Titman [18] established the benchmark for cross-sectional momentum based on total returns. Subsequent research has developed new momentum signals. Rachev et al. [25] propose ranking assets based on reward-to-risk metrics, such as the Sharpe ratio and the Sortino ratio. One of the most relevant variants was introduced by Blitz et al. [5]. The authors construct the momentum signal using the residuals from a regression based on the Fama-French three-factor model. This approach removes the asset’s exposure to systematic risk factors (Market, Size, Value),

isolating the idiosyncratic component of momentum. A paradigm shift is represented by Novy-Marx [22], who shows that momentum can be extracted not only from prices but also from firm fundamentals. Moving the time horizon to higher frequencies, studies such as Gao et al. [10] document the existence of intraday momentum, suggesting that investor flows generate predictable short-term trends even within a single trading day.

1.2. Implementation of Momentum Strategies

The classical cross-sectional momentum strategy proposed by Jegadeesh and Titman [18] relies on the cumulative total return of an asset over a predefined formation period. Specifically, the implementation requires selecting both a formation window J (in months), over which past returns are computed, and a holding period K (in months), during which the portfolio positions are maintained. For example, a strategy with $J = 12$ and $K = 6$ implies that cumulative returns are calculated over the previous twelve months and that the resulting portfolio is held for six months. In computing the total return, the most recent month is typically excluded in order to avoid short-term reversal effects. Once the momentum signal has been computed, assets are sorted into quantiles based on their past performance. The long leg of the strategy consists of the assets in the top quantile (the past winners), while the short leg comprises the assets in the bottom quantile (the past losers).

A distinctive feature of the Jegadeesh and Titman [18] implementation is the use of overlapping portfolios. Rather than forming a single portfolio that is fully liquidated and reconstituted every K months, a new portfolio is formed each month based on the preceding J -month performance, and each such portfolio is held for K months. At any given time, the strategy's aggregate position consists of K portfolios simultaneously active, with one portfolio closing and a new one opening each month. This overlapping structure smooths turnover, reduces the impact of entry and exit timing, and provides greater temporal diversification. This mechanism is central to the original methodology and distinguishes it from simpler rolling strategies that rebalance the entire portfolio at fixed intervals.

Empirical evidence suggests that formation windows from $J = 6$ to $J = 12$ months provide an optimal trade-off between signal strength and stability. Regarding to the holding period K , shorter horizons—typically between 3 and 6 months—are generally preferred, as longer holding periods tend to generate long-term mean reversion phenomena. The choice of the parameters J and K has been, and continues to be, the subject of extensive research concerning the temporal structure of momentum. Novy-Marx [21] shows that

a momentum signal constructed using returns from 12 to 7 months prior ($r_{t-12,t-7}$) outperforms signals based on more recent performance, suggesting that intermediate-term returns contain stronger predictive information. Moreover, Goyal et al. [11] examine significantly longer evaluation windows – up to five years – and observe that momentum may re-emerge following reversal periods. Their findings indicate that the optimal choice of the formation horizon should account for the asset’s life cycle and broader return dynamics, rather than focusing exclusively on the most recent year.

1.3. Market Regimes and Structural Limitations of Classical Momentum

While classical momentum has demonstrated robustness across markets and time periods, the strategy exhibits significant structural vulnerabilities that become particularly pronounced during certain market regimes.

1.3.1. Momentum Crashes and Regime Dependency

Momentum strategies are subject to rare but severe drawdowns known as momentum crashes. These crashes are characterized by sharp reversals in which past losers rebound violently while past winners stagnate or decline, generating losses for the long-short portfolio. This phenomenon is driven by the time-varying beta exposure of the two legs of the strategy. As studied in [13], the market beta of momentum strategies is a function of past market returns. Following market downturns, the ranking mechanism selects low-beta stocks as winners and high-beta stocks as losers. This composition results in a negative net beta for the overall long-short portfolio. During a sharp market rebound, this negative beta exposure becomes highly harmful: high-beta losers recover violently, generating substantial losses on the short positions, while low-beta winners participate only modestly in the rally. The most documented example occurred in the second quarter of 2009, following the Global Financial Crisis. As markets rapidly reversed from extreme pessimism, shorted losers surged by double digits within weeks, while long positions in past winners underperformed. The resulting winner-minus-loser portfolio losses ranged from -35% to -50% depending on implementation details. This episode illustrates a fundamental vulnerability: the performance of momentum varies dramatically across market regimes.

Wang and Xu [27] find that momentum crashes cluster in months immediately following large market drawdowns exceeding -30% , when fear-driven selling gives way to coordinated rebounds in oversold assets. Barroso and Santa-Clara [2] demonstrate that scaling

momentum positions inversely with ex-ante volatility substantially improves risk-adjusted returns and mitigates crash risk, suggesting that time-varying leverage adjustments can partially address regime-dependent vulnerabilities.

1.3.2. Structural Limitations

Beyond regime-specific vulnerabilities, momentum strategies face several persistent structural challenges. First, momentum returns exhibit substantial time-series instability. Rolling Sharpe ratios vary widely across decades. This instability complicates ex-ante risk assessment and challenges the assumption of stable factor premia.

Second, as documented in the previous section, momentum generates high portfolio turnover, depending on the J and K parameters. This turnover creates substantial transaction costs from bid-ask spreads, market impact, and tax inefficiency. Novy-Marx and Velikov [23] estimate that implementation costs reduce momentum's net returns by 1–3% annually, with the erosion particularly severe in small-cap and emerging markets where spreads and impact are higher.

Third, momentum portfolios often exhibit unintended exposures to sectors and risk factors. Winners tend to cluster in cyclical sectors during economic expansions, creating concentration risk. Sharp sector rotations, such as the tech crash of 2000-2002 or the energy collapse of 2014-2015, can devastate momentum portfolios that inadvertently overweight these areas.

Fourth, momentum strategies are vulnerable to crowding and liquidity spirals. As momentum has become widely adopted among quantitative asset managers since the 2000s, simultaneous de-risking during stress periods amplifies crashes.

2 | Hidden Markov Models: Theory and Financial Applications

Hidden Markov Models (HMMs) are a class of probabilistic models that describe systems evolving over time through a sequence of unobservable (hidden) states. Each hidden state governs the probability distribution of observable variables, allowing the model to capture complex temporal dependencies and regime-switching behavior. Originally developed in the 1960s for speech recognition, HMMs have since become a fundamental tool in signal processing, bioinformatics, and increasingly in financial econometrics.

2.1. Hidden Markov Models: Definitions and Theoretical Framework

A Hidden Markov Model is a stochastic model in which an observable process $\mathbf{Y}$ depends on an underlying latent Markov process $\mathbf{X}$. The key feature of the model is that the state process is not directly observable and must be inferred from the observed data.

Definition 2.1.1. *Let $\{X_t\}_{t \geq 0}$ and $\{Y_t\}_{t \geq 0}$ be two discrete-time stochastic processes. The pair (X_t, Y_t) is defined a discrete-time Hidden Markov Model if:*

1. $\{\mathbf{X}_n\}$ is a Markov process, i.e.

$$\mathbb{P}(X_{n+1} \in B \mid X_n, \dots, X_1) = \mathbb{P}(X_{n+1} \in B \mid X_n) \quad (2.1)$$

for every $n \geq 1$ and every Borel set B .

2. The observation $\mathbf{Y}_n$ depends only on the current hidden state $\mathbf{X}_n$, i.e.

$$\mathbb{P}(Y_n \in A \mid X_n = x_n, \dots, X_1 = x_1) = \mathbb{P}(Y_n \in A \mid X_n = x_n) \quad (2.2)$$

for every $n \geq 1$, all $x_1, \dots, x_n$ and every Borel set A .

Definition 2.1.2. Let $\{X_t\}_{t \geq 0}$ and $\{Y_t\}_{t \geq 0}$ be two continuous-time stochastic processes.

The pair (X_t, Y_t) is defined a continuous-time Hidden Markov Model if:

1. $\{X_n\}$ is a Markov process which is not directly observable. It fulfills (2.1).
2. For every $T > 0$ and every Borel set A ,

$$\mathbb{P}(Y_T \in A \mid \sigma(X_s : s \leq T)) = \mathbb{P}(Y_n \in A \mid X_T) \quad (2.3)$$

where $\sigma(X_s : s \leq T)$ denotes the σ -algebra generated by the process X_t up to time T , representing all information available up to that time.

2.2. Model Components

A Hidden Markov Model is formally defined by a five-tuple $\lambda = (X, Y, A, B, \pi)$, where:

1. State Space $X = \{x_1, x_2, \dots, x_N\}$: A finite set of N hidden states that the system can occupy at any discrete time t .
2. Observation Space $Y = \{y_1, y_2, \dots, y_T\}$: The set of possible observable outcomes.
3. Transition Probability Martix $A = \{a_{ij}\}$: An $N \times N$ matrix where $a_{ij} = \mathbb{P}(X_{t+1} = x_j \mid X_t = x_i)$ represents the probability of transitioning from state i to state j . The matrix satisfies $\sum_{j=1}^N a_{ij} = 1$.
4. Emission Probability Distribution B :
 - Discrete Case: $B = \{b_j(\bar{y})\}$ where $b_j(\bar{y}) = \mathbb{P}(y_t = \bar{y} \mid x_t = s_j)$ is the probability of observing symbol $\bar{y} \in \mathcal{Y}$ when the system is in state s_j . The density satisfies

$$\sum_{\bar{y} \in \mathcal{Y}} b_j(\bar{y}) = 1 \quad \forall j$$

- Continuous Case: $B = \{b_j(\cdot)\}$ where $b_j(z) = f(z \mid x_t = s_j; \theta_j)$ is a probability density function parameterized by θ_j . It holds that

$$\int_{-\infty}^{+\infty} b_j(z) dz = 1 \quad \forall j.$$

5. Initial State Distribution $\pi = \{\pi_i\}$: A vector of probabilities where $\pi_i = \mathbb{P}(q_1 = s_i)$ represents the probability that the process starts in state s_i at time $t = 1$. The distribution satisfies the normalization condition $\sum_{i=1}^N \pi_i = 1$.

As stated previously, HMMs rely on two fundamental Markov assumptions:

1. **Markov Property:** The probability of transitioning to the next state depends only on the current state, not on the history of past states.
2. **Output Independence:** The observation at time t depends only on the current hidden state, not on past observations or states.

2.2.1. State Space Design

The choice of the state space is crucial and depends on the specific research question under investigation. In financial applications, both the number of hidden states N and their economic interpretation are determined by the phenomenon being modeled.

The most common specification involves two hidden states, which are typically interpreted as risk-on and risk-off market regimes. Two-state models offer greater interpretability and parameter parsimony; however, they may lead to an oversimplification of complex market dynamics. For example, Ang et. al [1] employ a two-state regime-switching model to distinguish between bull and bear regimes in international equity markets. In contrast, Guidolin et al. [14] adopt Hidden Markov Model with three and four states in an asset allocation framework, showing that the inclusion of intermediate regimes improves out-of-sample performance and economic interpretability. Higher-order models (i.e., $N \geq 4$) are also present in the literature. However, increasing the number of states raises concerns about overfitting and may reduce the economic interpretability of the model.

2.2.2. Transition Dynamics

The transition matrix A governs regime persistence and the switching dynamics between states. The diagonal elements a_{ii} represent the probability of remaining in regime i from one period to the next, capturing regime persistence. From these diagonal probabilities, it is also possible to compute the expected duration of regime i before a switch occurs. The expected duration is given by:

$$\mathbb{E}[D_i] = \frac{1}{1 - a_{ii}}$$

Off-diagonal elements a_{ij} , with $i \neq j$, measure the probability of transitioning from regime i to regime j , thus quantifying regime-switching dynamics. The transition matrix is not necessarily symmetric. In financial markets, transitions from a bull to a bear regime often occur more abruptly than recoveries from bear to bull markets. This asymmetry can be analyzed by comparing the relevant transition probabilities—for instance, if the

probability of remaining in a bull state is higher than that of remaining in a bear state, the implied expected duration of bull markets will be longer.

Although rare in financial applications, it is theoretically possible for a state to be defined as absorbing, i.e., $a_{ii} = 1$. In such a case, once the process enters that state, it cannot exit.

2.2.3. Emission Distributions

The emission distribution B specifies how observations are generated conditional on the hidden regime. As previously discussed, these distributions may be either discrete or continuous, depending on the nature of the observed data. In the discrete setting, typical choices include the Bernoulli, Poisson, and multinomial distributions, the latter being widely employed in text and speech recognition applications where observations take categorical values.

In financial applications, however, observed variables are generally continuous. Consequently, emission distributions must be modeled through continuous probability laws. The most common specification assumes that observations are conditionally Gaussian. Formally,

$$Y_t \mid X_t = j \sim \mathcal{N}(\mu_j, \sigma_j^2),$$

where the parameter vector for regime j is given by $\theta_j = \{\mu_j, \sigma_j^2\}$, with μ_j and σ_j^2 denoting the regime-dependent mean and variance. This framework is extensively treated in the regime-switching literature on asset returns, starting from the seminal contribution of Hamilton [15], which introduced Markov-switching models for macroeconomic time series.

The Gaussian specification offers important advantages in terms of computational efficiency and interpretability. Parameter estimation is performed via the Baum-Welch algorithm [3], an Expectation-Maximization procedure that provides a solution to the maximum likelihood problem. Moreover, regime-specific means and variances allow for an economically intuitive characterization of the states, such as distinguishing between high-volatility and low-volatility environments or identifying transitions toward favorable market conditions.

Despite these advantages, the normal distribution is often inadequate for financial data. Empirical evidence consistently shows that asset returns exhibit heavy tails and, in some cases, skewness which are not captured by the Gaussian assumption. As a consequence, tail risk measures derived under normality may be underestimated, leading to unreliable assessments of extreme risk.

To address these limitations, emission distributions can instead be modeled using a Student- t specification,

$$Y_t \mid X_t = j \sim t_{\nu_j}(\mu_j, \sigma_t^2),$$

where the degrees of freedom ν_j also be regime-dependent. The Student- t distribution provides greater flexibility in capturing heavy tails. However, unlike the Gaussian case, the likelihood function does not admit a simple closed-form maximization, and parameter estimation requires numerical optimization methods.

An additional layer of flexibility can be introduced by modeling emissions as mixture distributions. In this case, the conditional density within regime J is expressed as

$$b_j(z) = \sum_{k=1}^{K_j} w_{jk} * f_k(z, \theta_{jk})$$

where w_{jk} are mixture weights satisfying $\sum_k w_{jk} = 1$, and f_k are component densities (e.g. Gaussians). This approach allows for greater heterogeneity within a single regime. Nevertheless, the increase in the number of parameters substantially raises computational complexity and may introduce overfitting concerns, particularly in finite samples.

2.3. Estimation and Inference Algorithms

Estimating the parameters of a Hidden Markov Model and inferring the hidden states from observed data involves addressing three fundamental problems, originally formalized by Rabiner [24]. These problems constitute the theoretical backbone of HMM methodology. The associated algorithms extend naturally from the discrete to the continuous case by replacing probability mass functions with probability density functions, without altering the general structure of the procedures.

The first problem is the evaluation problem. Given a model λ and an observation sequence Y , the objective is to compute the likelihood $\mathbb{P}(Y \mid \lambda)$. This quantity measures how well the specified model explains the observed data and plays a central role in model comparison and likelihood-based inference.

The second problem is the decoding problem. Given a model λ and an observation sequence Y , the goal is to determine the most likely sequence of hidden states, formally defined as

$$Q^* = \arg \max_Q \mathbb{P}(Q \mid Y, \lambda).$$

This step allows to recover the underlying state path that most plausibly generated the

data. In financial applications, decoding corresponds to identifying historical market regimes.

The third problem is the parameter estimation problem. Given an observation sequence Y , the aim is to determine the model parameters

$$\lambda^* = \arg \max_{\lambda} \mathbb{P}(Y | \lambda),$$

that maximize the likelihood function. This is typically the most complex of the three problems, as it requires iterative optimization procedures. In practice, the Expectation–Maximization framework is employed to obtain maximum likelihood estimates of the parameters.

Together, these three problems define the core computational and inferential tasks in Hidden Markov Model and provide the foundation for their application in economics and finance.

2.3.1. Problem I: the Forward-Backward Algorithm

The forward-backward algorithm provides a procedure to compute the likelihood $\mathbb{P}(Y | \lambda)$ and the posterior probabilities of the hidden states. The algorithm applies identically to both discrete and continuous HMMs; the only difference lies in replacing probability mass functions with probability density functions in the case of continuous observations.

Forward Pass

The forward procedure is based on the definition of the forward variable

$$\alpha_t(i) = \mathbb{P}(y_1, y_2, \dots, y_t, X_t = s_i | \lambda),$$

which represents the joint probability of observing the sequence up to time t and being in state s_i at time t , given the model λ .

The recursion begins with the initialization step at $t = 1$,

$$\alpha_1(i) = \pi_i b_i(y_1),$$

where π_i denotes the initial state probability and $b_i(y_1)$ is either the probability mass (discrete case) or the probability density (continuous case) evaluated at the observation y_1 .

For $t = 1, \dots, T - 1$, the forward variables evolve according to the recursion

$$\alpha_{t+1}(j) = \left(\sum_{i=1}^N \alpha_t(i) a_{ij} \right) b_j(y_{t+1}),$$

where a_{ij} denotes the transition probability from state i to state j .

The likelihood of the entire observation sequence is then obtained through the termination step

$$\mathbb{P}(Y \mid \lambda) = \sum_{i=1}^N \alpha_T(i).$$

This algorithm has computational complexity $O(N^2T)$, which represents a substantial improvement over the naive approach requiring summation over all N^T possible state sequences. For long observation sequences, numerical underflow may occur due to repeated multiplication of probabilities or densities smaller than one. To ensure numerical stability, a scaling procedure is introduced. At each time t , a scaling coefficient is defined as

$$c_t = \sum_{i=1}^N \alpha_t(i),$$

and the forward variables are rescaled as

$$\hat{\alpha}_t(i) = \frac{\alpha_t(i)}{c_t}.$$

Under this convention, the likelihood can be reconstructed as the product of scaling coefficients,

$$\mathbb{P}(Y \mid \lambda) = \prod_{t=1}^T c_t,$$

so that the log-likelihood is given by

$$\log \mathbb{P}(Y \mid \lambda) = \sum_{t=1}^T \log c_t.$$

This procedure prevents numerical underflow while preserving exact likelihood computation.

Backward Pass

The backward procedure complements the forward pass. It is based on the backward variable

$$\beta_t(i) = \mathbb{P}(y_{t+1}, \dots, y_T \mid X_t = s_i, \lambda),$$

which represents the probability of observing the remaining part of the sequence, conditional on being in state s_i at time t . The recursion is initialized at the terminal time T by setting

$$\beta_T(i) = 1 \quad \forall i.$$

For $t = T - 1, \dots, 1$, the backward variables satisfy

$$\beta_t(i) = \sum_{j=1}^N a_{ij} b_j(y_{t+1}) \beta_{t+1}(j).$$

When scaling is employed in the forward pass, the backward variables must be scaled consistently using the same coefficients c_t to maintain numerical coherence.

Posterior State Probabilities (Smoothing)

By combining forward and backward variables, one obtains the posterior probability of being in state s_i at time t , conditional on the entire observation sequence:

$$\gamma_t(i) = \mathbb{P}(X_t = s_i \mid Y, \lambda) = \frac{\alpha_t(i) \beta_t(i)}{\sum_{j=1}^N \alpha_t(j) \beta_t(j)}.$$

Under the scaling convention, this simplifies to

$$\gamma_t(i) = \hat{\alpha}_t(i) \hat{\beta}_t(i).$$

These posterior probabilities incorporate information from both past and future observations and therefore provide the best ex-post estimate of the hidden regime at each time point. In financial applications, they are typically used for historical regime identification.

Filtered Probabilities (Real-Time Inference)

In contrast, real-time inference relies on filtered probabilities, which condition only on past and current observations. The filtered probability of state s_i at time t is given by

$$\mathbb{P}(X_t = s_i \mid y_1, \dots, y_t, \lambda) = \frac{\alpha_t(i)}{\sum_{j=1}^N \alpha_t(j)}.$$

Filtered probabilities avoid look-ahead bias and are therefore suitable for live trading or forecasting applications, although they are generally noisier than posterior probabilities because they do not incorporate future information.

2.3.2. Problem II: the Viterbi Algorithm

The Viterbi algorithm [9] finds the single most likely state sequence $X^* = (X_1^*, \dots, X_T^*)$ using dynamic programming. The procedure is identical for both discrete and continuous HMMs. It identifies the path that maximizes the joint probability of observations and states.

Define the variable

$$\delta_t(i) = \max_{X_1, \dots, X_{t-1}} \mathbb{P}(X_1, \dots, X_{t-1}, X_t = s_i, Y_1, \dots, Y_t \mid \lambda),$$

which represents the maximum probability of reaching state s_i at time t along any path. The recursion begins with the initialization at $t = 1$:

$$\delta_1(i) = \pi_i \cdot b_i(Y_1),$$

$$\psi_1(i) = 0,$$

where $\psi_t(j)$ stores the predecessor state that maximizes the path probability.

For $t = 2, \dots, T$, the recursion step is

$$\delta_t(j) = \max_{i \in \{1, \dots, N\}} [\delta_{t-1}(i) \cdot a_{ij}] \cdot b_j(Y_t),$$

$$\psi_t(j) = \arg \max_{i \in \{1, \dots, N\}} [\delta_{t-1}(i) \cdot a_{ij}].$$

At termination, the probability of the most likely path is

$$\mathbb{P}^* = \max_{i \in \{1, \dots, N\}} \delta_T(i),$$

and the final state is

$$X_T^* = \arg \max_{i \in \{1, \dots, N\}} \delta_T(i).$$

Backtracking allows recovery of the optimal state sequence for $t = T - 1, \dots, 1$:

$$X_t^* = \psi_{t+1}(X_{t+1}^*).$$

The Viterbi path provides a single deterministic regime classification for each time point, which is particularly useful for historical analysis and visualization. However, it discards uncertainty. For instance, if $\gamma_t(1) = 0.51$ and $\gamma_t(2) = 0.49$, the Viterbi algorithm assigns state 1 with full certainty, whereas posterior probabilities preserve the near-ambiguity.

2.3.3. Problem III: the Baum-Welch Algorithm

The Baum-Welch algorithm [3] is an Expectation-Maximization (EM) procedure for learning HMM parameters $\lambda = (A, B, \pi)$ from data. It applies to both discrete and continuous HMMs, with the M-step formulas depending on the chosen emission distribution.

E-Step (Expectation)

Given current parameter estimates $\lambda^{(k)}$, the posterior state probabilities $\gamma_t(i)$ are calculated using the forward-backward algorithm, as well as the posterior transition probabilities

$$\xi_t(i, j) = \mathbb{P}(X_t = s_i, X_{t+1} = s_j \mid Y, \lambda) = \frac{\alpha_t(i) \cdot a_{ij} \cdot b_j(Y_{t+1}) \cdot \beta_{t+1}(j)}{\sum_{i'=1}^N \sum_{j'=1}^N \alpha_t(i') \cdot a_{i'j'} \cdot b_{j'}(Y_{t+1}) \cdot \beta_{t+1}(j')}$$

where, with scaling, the denominator simplifies to c_{t+1} .

M-Step (Maximization)

The parameters are updated to maximize the expected log-likelihood. The initial distribution is updated as

$$\pi_i^{(k+1)} = \gamma_1(i),$$

while the transition probabilities are updated according to

$$a_{ij}^{(k+1)} = \frac{\sum_{t=1}^{T-1} \xi_t(i, j)}{\sum_{t=1}^{T-1} \gamma_t(i)},$$

which represents the expected number of transitions from state i to state j divided by the expected number of times in state i .

For the emission parameters, the update depends on the chosen distribution. For Gaussian emissions $\mathcal{N}(\mu_j, \sigma_j^2)$, the updates are

$$\mu_j^{(k+1)} = \frac{\sum_{t=1}^T \gamma_t(j) \cdot Y_t}{\sum_{t=1}^T \gamma_t(j)}, \quad (\sigma_j^2)^{(k+1)} = \frac{\sum_{t=1}^T \gamma_t(j) \cdot (Y_t - \mu_j^{(k+1)})^2}{\sum_{t=1}^T \gamma_t(j)}.$$

For Student's t emissions $t_{\nu_j}(\mu_j, \sigma_j^2)$, the degrees of freedom ν_j require numerical optimization:

$$\nu_j^{(k+1)} = \arg \max_{\nu} \sum_{t=1}^T \gamma_t(j) \log b_j(Y_t; \mu_j, \sigma_j, \nu).$$

The algorithm iterates E-step and M-step until convergence, typically defined as

$$|\log \mathbb{P}(Y \mid \lambda^{(k+1)}) - \log \mathbb{P}(Y \mid \lambda^{(k)})| < \epsilon,$$

where ϵ is a small threshold. During the procedure, the log-likelihood increases (or remains constant) at each iteration, and the algorithm converges to a local maximum, which is not necessarily global. In practice, multiple initializations are recommended. Baum-Welch can be run from different starting points and the solution with the highest likelihood is selected.

2.4. Applications of HMMs in Finance

Hidden Markov Models have found numerous applications in financial economics, ranging from macroeconometrics to asset pricing and portfolio management. The purpose of this section is to briefly review some of the main contributions in order to highlight the flexibility and applicability of this modeling framework.

Hamilton [15] was the first to introduce a Markov regime-switching model to identify different phases of the U.S. business cycle. Building on this framework, [1] applied Markov-switching models with Gaussian emissions to equity returns, distinguishing between bull and bear market regimes. Their results showed that asset returns and risk premia exhibit markedly different statistical properties across regimes. In the fixed-income literature, [12] introduced a regime-switching term structure model in which interest rate dynamics depend on latent states. Similarly, [8] applied regime-switching techniques to exchange rate dynamics, identifying different degrees of mean reversion across states and providing evidence of nonlinear adjustment mechanisms in foreign exchange markets.

Regime-switching models have also played an important role in risk management. [16] proposed a Markov-Switching ARCH model, extending the traditional ARCH framework by allowing volatility dynamics to differ across regimes. This approach captures changes in market volatility more effectively than constant-parameter GARCH models. In a multi-variate context, [4] developed a regime-switching Dynamic Conditional Correlation (DCC) model, allowing the correlation matrix to vary across latent states, thus providing a time-varying dependence structures.

HMMs have further been applied to derivative pricing. For instance, [6] studied option valuation under regime-switching dynamics, showing that allowing parameters such as volatility to change across states.

More recently, regime-switching models have been employed as engines for stock selection and tactical asset allocation. For example, [19] and [20] explored the use of HMM-based regime identification for enhancing equity selection strategies. Closely related to the focus of this thesis is the application of HMMs to momentum investing, where regime classification is used to condition momentum signals on the prevailing market state (e.g., [26]).

Overall, the literature demonstrates that HMMs provide a powerful and versatile framework for modeling nonlinear dynamics, structural breaks, and state-dependent behavior in financial time series.

3 | Data and Methodology

This section aims to describe the data employed in the empirical analysis and the methodology adopted to construct the portfolio, which relies on a scoring procedure derived from Hidden Markov Model.

3.1. Data Description

The study is conducted on the full universe of stocks included in the S&P 500 index from the beginning of 2000 to the end of 2025. To avoid look-ahead bias, the incorrect use of current index constituents to evaluate past performance, the historical point-in-time composition of the index is reconstructed at the beginning of each quarter. This choice is consistent with the index rebalancing schedule, which takes place on the third Friday of March, June, September, and December. Using the quarterly point-in-time composition ensures that only the information available at each rebalancing date is incorporated into the analysis.




Figure 3.1: Performance Comparison: S&P500 EW vs. S&P500

For each stock in the sample, weekly dividend-adjusted prices and market capitalization were collected. Additionally, each company was classified according to the first-level Global Industry Classification Standard (GICS).

It is crucial to emphasize that all firms, both currently listed and delisted, are assigned the GICS sector classification currently in effect. Consequently, no adjustments were made to reflect potential historical changes in sector classification at the individual firm level. Adopting a static classification based on today's standards introduces a potential look-ahead bias: the sector-level data analyzed in the following chapters may be affected by this temporal discrepancy, as past sector compositions are projected according to contemporary industrial logic.

A prime example of this dynamic is Alphabet (Google) and Meta (Facebook), which were reclassified in 2018 from the Information Technology sector to the then-newly created Communication Services sector. Similarly, companies like Visa and Mastercard are currently listed under the Financials sector, despite being classified as Information Technology prior to the 2023 GICS restructuring.

Despite the introduction of this potential look-ahead bias, its impact on the final strategy results is considered negligible for two primary reasons. First, the equal-weighted scheme adopted for portfolio construction ensures that no single reclassified large-cap stock can disproportionately distort the aggregate sector-level metrics. Second, since the strategy is designed to capture regime inertia based on latent statistical patterns rather than sector-specific fundamental mandates, the persistence of the identified trends remains the primary driver of performance. The model responds to the underlying price dynamics of the assets, which remain consistent regardless of the formal sectoral label assigned to them at any given time.




Figure 3.2: Point-in-Time GICS sectors market capitalization weight.

Figure 3.2 reports the percentage weight of each GICS sector in terms of total market capitalization. As of the end of 2025, the Information Technology sector represents approximately 35% of the total market capitalization, making it by far the largest sector within the investment universe. A similar expansion can be observed for Communication Services. More generally, by the end of 2025, roughly 80% of the total market capitalization is concentrated in only five sectors: Information Technology, Financials, Health Care, Communication Services, and Consumer Discretionary. On the opposite side, the Energy sector, after reaching a peak in 2008 when oil prices hit historical highs, has progressively declined and now accounts for only a marginal share of the overall index.




Figure 3.3: Intra-Sector Volatility

As widely documented in the literature, GICS sectors can be broadly classified into two macro-categories: cyclical and defensive sectors. Cyclical sectors are those that tend to exhibit strong performance during periods of economic expansion and to underperform during recessions. These typically include Information Technology, Energy, Financials, Consumer Discretionary, and, to a slightly lesser extent, Industrials and Materials. The remaining sectors are generally classified as defensive, as they tend to display greater resilience during economic downturns.

Figure 3.3 reports the annual volatility quintiles for each sector. Volatility is computed by considering each sector as an equal-weight portfolio composed of the stocks belonging to that sector and included in the universe on the first trading day of each year. The analysis reveals substantial heterogeneity in risk dynamics across sectors, consistent with the cyclical versus defensive distinction. The burst of the dot-com bubble in the early 2000s had a particularly severe impact on the Information Technology and Communication Services sectors. In contrast, the 2008 subprime mortgage crisis affected all sectors simultaneously, representing an unprecedented systemic event. A similar pattern emerged in 2020 during the COVID-19 pandemic, when financial markets experienced extreme volatility over very short time intervals. More recently, the Energy sector has exhibited exceptional volatility levels, driven by the sharp fluctuations in energy commodities. These prices transitioned from the historic lows observed during the pandemic – when they briefly entered negative territory – to new highs reached in 2022, following the supply shocks and geopolitical tensions triggered by the war in Ukraine.




Figure 3.4: Intra-Sector Average Pairwise Correlation

The same methodology adopted for the construction of sector indices has been applied to analyze the average pairwise correlation within each sector. This measure is particularly relevant for achieving effective portfolio diversification. As shown in Figure 3.4, the aver-

age correlation among stocks belonging to the same sector is generally positive. However, a relevant feature emerges during periods of market stress: correlations increase significantly, reducing diversification benefits and making the portfolio more exposed to systemic market movements. Over the sample period, the Energy, Real Estate, Utilities, and Financials sectors have exhibited relatively higher levels of intra-sector correlation compared to other sectors. Conversely, defensive sectors such as Health Care and Consumer Staples not only display lower average volatility but also show lower levels of internal correlation.

	CAGR	Vol	Skew	Exc. Kurt.	VaR95%	CVaR95%	MDD	Sharpe
Utilities	10.39%	17.91%	-1.04	4.21	-7.09%	-12.21%	-48.03%	0.65
Comm. Serv.	6.40%	21.31%	0.09	2.84	-10.05%	-13.76%	-64.85%	0.40
Health Care	11.79%	15.48%	-0.76	2.45	-6.79%	-10.39%	-40.17%	0.80
Cons. Staples	9.40%	13.16%	-0.40	2.04	-5.17%	-8.09%	-34.05%	0.75
Industrials	11.71%	21.81%	-0.03	4.80	-8.10%	-13.97%	-59.44%	0.62
Materials	9.79%	22.94%	0.29	5.36	-9.72%	-14.72%	-59.95%	0.52
Financials	9.13%	26.66%	1.33	16.91	-11.56%	-16.80%	-80.54%	0.46
Cons. Discr.	9.53%	25.41%	0.55	8.55	-9.16%	-15.73%	-66.84%	0.48
IT	8.98%	28.10%	0.34	3.68	-12.35%	-18.00%	-79.56%	0.45
Energy	8.62%	30.30%	0.38	4.77	-12.06%	-18.67%	-74.78%	0.42
Real Estate	8.51%	25.10%	-0.27	7.62	-9.67%	-16.78%	-75.79%	0.46

Table 3.1: Annualized sector performances and risk statistics, 2000-2026

The empirical results presented in Table 3.1 provide a comprehensive overview of the risk-return profiles across different U.S. equity sectors over the 2000–2026 period, revealing a distinct dichotomy between cyclical and defensive sectors. The analysis confirms the resilient nature of defensive sectors, most notably Health Care, which achieved a CAGR of 11.79%, and Consumer Staples, with a CAGR of 9.40%. These clusters not only attained superior risk-adjusted returns—evidenced by Sharpe Ratios of 0.80 and 0.75, respectively—but also exhibited the lowest annualized volatilities in the sample at 15.48% and 13.16%. In contrast, cyclical sectors such as Energy and Information Technology display significantly higher volatility, reaching 30.30% and 28.10% respectively. Interestingly, despite the rapid technological expansion observed in recent years, the long-term CAGR of the IT sector remains lower than that of Health Care over this extended horizon, primarily due to the severe drawdowns experienced during systemic crises which reduced the long-term compounding of returns. This highlights a critical investment principle where capital preservation during market contractions, typical of defensive sectors, often proves more deterministic of terminal wealth than the capture of cyclical peaks.

The excess kurtosis values provide additional insight into tail risk dynamics, with Finan-

cial and Consumer Discretionary exhibiting high values of 16.91 and 8.55, respectively. Such leptokurtic distributions indicate a higher propensity for black swan events, violent shocks that far exceed the expectations of a normal distribution. Conversely, Health Care and Consumer Staples show significantly lower excess kurtosis, suggesting a return distribution closer to Gaussian and less prone to extreme outliers. The economic rationale for this divergence lies in the business structure; cyclical industries are linked to credit availability and consumer confidence, leading to aggressive deleveraging and panic selling during downturns, whereas defensive sectors provide essential services that maintain stable demand, effectively smoothing the tails of their return distributions.

3.2. Proposed Strategy

The analysis of volatility and intra-sector correlations plays a central role in the construction of momentum-based strategies. Selecting stocks solely on the basis of their past realized performance may represent a weak criterion. In contrast, conditioning the selection on the broader dynamics of the sector, can lead to more robust investment decisions. This highlights the importance of decomposing aggregate market volatility into its sector-specific components, in order to better identify the underlying drivers of risk. Such considerations motivate the adoption of models capable of detecting sector-specific clustering patterns and regime dynamics, thereby supporting the construction of portfolios that remain resilient across different phases of the economic cycle.

Moreover, the core hypothesis of this thesis is that market regimes exhibit significant serial correlation, a phenomenon often referred to as regime persistence. Unlike standard momentum strategies that rely solely on past returns, this framework assumes that the underlying latent state of the market possesses an inherent stability. By identifying these states through a Hidden Markov Model, the strategy aims to exploit the tendency of expansionary phases to persist over time, thereby aligning portfolio exposure with the long-term market inertia.

3.2.1. Regime Identification Framework

For this reason, this thesis develops a momentum strategy conditioned on the volatility state of each sector, with the objective of identifying sectors and stocks that are in an expansionary phase. Regime identification is performed through a Hidden Markov Model, which is widely documented in the literature as a robust framework for modeling regime shifts, state transitions, and persistent trends in financial time series.

In the following analysis, Gaussian Hidden Markov Models with both two and three latent states are estimated. The two-state specification allows for a clear distinction between low-volatility (expansionary) regime and high-volatility (stress) regime. This binary classification provides a relatively robust segmentation of market conditions; however, it may fail to adequately capture borderline or transitional dynamics. By contrast, the three-state specification introduces an additional intermediate regime, typically associated with mean-reverting or transitional market behavior. While this structure enables a more granular characterization of volatility dynamics, it also increases the model's sensitivity to noise and may reduce the stability of state classification.

3.2.2. Investment Universe and Benchmark Definition

To ensure the robustness of the estimation process, the models are trained on a dynamic universe of stocks that satisfy two key criteria:

- Point-in-Time Membership: Stocks must be constituents of the index at the time of analysis to prevent look-ahead bias;
- Historical Depth: Stocks must have a complete price history for the preceding five years.

This selection reflects a necessary trade-off: since the HMM requires consistent time-series length for accurate fitting, a five-year lookback period was chosen to provide sufficient data for reliable state identification without excessively restricting the sample size or excluding too many firms from the investable universe.

As a consequence, the actual training universe consistently includes fewer than 500 stocks, as observable in Fig. 3.5, since all companies with less than five years of data are systematically excluded. Starting from this universe, a long-only momentum strategy is also estimated as in [18], using the Sharpe ratio as ranking metric. In the following sections, this strategy will be referred to as the main benchmark to evaluate and compare the performance results obtained through the proposed HMM frameworks.




Figure 3.5: Evolution of the investment universe size used for portfolio construction

3.2.3. Decision Rule

Each sector is analyzed independently by estimating a sector-level Hidden Markov Model using the return series of all constituent stocks. The model identifies a common latent regime process at the sector level, which is assumed to affect all stocks within the sector. However, the application of this regime dynamic is divided into two distinct phases: nowcasting and forecasting.

Nowcasting for Stock Ranking

To identify the most promising assets within each sector, the model generates a stock-specific ranking metric based on a nowcasting approach. Instead of relying on posterior probabilities, which can be sensitive to short-term market noise, an Exponential Moving Average (EMA) is applied to the state probabilities.

This choice is rooted in the hypothesis of regime momentum, where a persistent expansionary regime is more likely to continue than a transient spike. This inertia is captured by smoothing the probabilities, creating a robust state distribution that serves as the basis for the cross-sectional ranking of stocks.

Forecasting for Regime Filtering

The forecasting component is utilized as a risk-management filter to determine the overall market exposure. The sectoral regime dynamics are projected to a $t + 6M$ horizon using the transition matrix. To ensure the stability of this long-term view, the harmonic mean of the transition probabilities is computed across the forecast horizon.

The decision to invest in stocks belonging to a specific sector is then governed by comparing this forecasted probability against three distinct threshold levels:

- Unfiltered: No regime-based restriction is applied.
- Naive: A static threshold represents the neutral probability of being in a specific state, set at 50% for 2-states HMM and 33.3% for 3-states one.
- Stationary: A dynamic threshold based on the stationary distribution of the Markov chain, which accounts for the structural features of each sector.

By separating these two dimensions, the strategy ensures that stock selection is driven by persistent regime inertia, while capital preservation is managed through a forward-looking assessment of regime stability.

3.2.4. Rebalancing and Transaction Costs

The model is re-estimated on a yearly basis, using a five-year rolling window. However, portfolio formation occurs on a monthly basis: each month a new portfolio is initiated while an existing one exits, with its holding period determined by K , as in the original overlapping framework. The choice of monthly rebalancing helps mitigate noise in state estimation while containing portfolio turnover. The total transaction cost at time t is modeled as:

$$TC_t = c \cdot \sum_{i \in \mathcal{U}} S_{i,t} \cdot |Q_{i,t} - Q_{i,t-1}|$$

where the parameter c represents the proportional transaction cost coefficient, expressed in basis points (bps) of the total notional value traded. The second term indicates the total notational traded between $t - 1$ and t . This cost is then deducted from the gross performance in order to account for the impact of trading activity.

Moreover, just the long-only momentum strategy is implemented. The short leg is intentionally excluded from the strategy for both empirical and theoretical reasons. Academically, it has been demonstrated that the momentum effect is significantly stronger and more persistent in the long leg than in the short leg. [17] observed a pronounced asymmetry in information diffusion: negative news are incorporated into prices more slowly than positive news, primarily due to short-selling constraints and management's tendency to withhold adverse information. Furthermore, the authors highlight that the profitability of the short leg is heavily concentrated among small-cap stocks with low analyst coverage. Since the universe analyzed in this thesis consists of the 500 largest U.S. companies by market capitalization, characterized by extensive analyst coverage and efficient information dissemination, the decision to eliminate the short leg is aligned with these findings.

Furthermore, past performance alone is frequently an insufficient predictor for short positions due to momentum crashes. As shown by [7], short portfolios are prone to sudden and violent reversals during market rebounds, creating a convex risk profile that can lead to catastrophic losses, potentially offsetting the gains generated by the long positions.

3.3. Mathematical Formalism of the Proposed Strategy

This section formalizes the mathematical framework underlying the sector-based momentum strategy described previously. Using the trained HMM, posterior probabilities are extracted for each hidden regime $h = 0, 1, \dots, H$:

$$\pi_t = \begin{bmatrix} \pi_t^{(0)} \\ \pi_t^{(1)} \\ \dots \\ \pi_t^{(H)} \end{bmatrix} = \begin{bmatrix} \mathbb{P}(X_t = 0 \mid \mathcal{F}_T) \\ \mathbb{P}(X_t = 1 \mid \mathcal{F}_T) \\ \dots \\ \mathbb{P}(X_t = H \mid \mathcal{F}_T) \end{bmatrix}, \forall t \in [t_0, T]$$

where $\mathcal{F}_T$ denotes the filtration at time T , representing the totality of information available up to the final time, which coincides with the endpoint of the rolling training window.

These probabilities are then used to define, for each stock i , the conditional expected return, volatility, and risk-adjusted return:

$$\begin{aligned} \mu_i^{(h)} &= \mathbb{E}[r_{i,t} \mid X_t = h] = \sum_{t=T-\Delta}^T \frac{\pi_t^{(h)}}{\sum_{t=T-\Delta}^T \pi_t^{(h)}} r_{i,t} \\ \sigma_i^{(h)} &= \sqrt{\text{Var}[r_{i,t} \mid X_t = h]} = \sqrt{\sum_{t=T-\Delta}^T \frac{\pi_t^{(h)}}{\sum_{t=T-\Delta}^T \pi_t^{(h)}} (r_{i,t} - \mu_i^{(h)})^2} \\ SR_i^{(h)} &= \frac{\mu_i^{(h)}}{\sigma_i^{(h)}} \end{aligned}$$

where $r_{i,t}$ denotes the observed return of asset i at time t and X_t is the hidden state at t .

The conditional metrics are aggregated using the regime distribution that represents market inertia. As described in Section 3.2.3, this distribution is obtained by applying an Exponential Moving Average (EMA) to the posterior probabilities. To align with the formation period J established by [18], the probabilities from the last 12 monthly obser-

vations are utilized. The distribution is calculated as follows:

$$\begin{aligned} EMA_t^{(h)} &= \alpha \cdot \pi_t^{(h)} + (1 - \alpha) \cdot EMA_{t-1}^{(h)} \quad \forall t, \forall h \\ EMA_0 &= \pi_0 \end{aligned}$$

where $\alpha = \frac{2}{J+1}$ and $J = 12$. As previously noted, the decision to use an exponential smoothing mechanism integrates the hypothesis of regime momentum, or regime inertia, into the stock selection process. This approach mitigates the impact of high-frequency noise inherent in the point-in-time observation π_t , providing a more stable signal for asset allocation. Using this probability distribution, the expected metrics for each stock i are computed:

$$\begin{aligned} \hat{\mu}_{i,T} &= \sum_{h=0}^H EMA_T^{(h)} \cdot \mu_i^{(h)} \\ \hat{\sigma}_{i,T}^2 &= \sum_{h=0}^H EMA_T^{(h)} \cdot (\sigma_i^{(h)})^2 + \sum_{h=0}^H EMA_T^{(h)} \cdot (\mu_i^{(h)} - \hat{\mu}_{i,T})^2 \\ \hat{\sigma}_{i,T} &= \sqrt{\hat{\sigma}_{i,T}^2} \\ \hat{SR}_{i,T} &= \frac{\hat{\mu}_{i,T}}{\hat{\sigma}_{i,T}} \end{aligned}$$

As described in 3.2.3, stocks are ranked according to their expected Sharpe ratio ($\hat{SR}_{i,t}$), and assets belonging to the top decile of the distribution are selected for the portfolio.

In addition to individual stock selection, a sector-level filtering process is performed as explained in 3.2.3. Starting from the instantaneous distribution π_T at the time of portfolio rebalancing, a 6-month forecast is calculated as the harmonic mean of the future distributions from one to six months. This approach aligns the forecast with the holding period of each sub-portfolio ($K = 6$) while accounting for the increased uncertainty associated with longer-term predictions. Given that the predicted distribution of market states at $T + \tau$ months is defined as:

$$\pi_{T+\tau} = \pi_T \cdot \mathbf{M}^\tau$$

where $\mathbf{M}$ represents the monthly transition matrix for each specific sector – derived from the weekly matrix $\mathbf{W}$ as $\mathbf{M} = \mathbf{W}^4$ – the probability for each regime $h \in [0, H]$ used to filter the sectors is calculated as:

$$\bar{\pi}_{T+K}^{(h)} = \frac{1}{\sum_{k=1}^K \frac{1}{k}} \cdot \sum_{k=1}^K \frac{1}{k} \pi_T \cdot (\mathbf{M}^k)_{,h}$$

From this distribution, the probability associated with the growth regime ($\bar{\pi}_{T+K}^{Growth}$) is

extracted and compared against a threshold value γ to determine whether the stocks belonging to that sector should be included in the portfolio. Three scenarios for the threshold γ are considered:

- Unfiltered case: No restrictions are applied.
- Naive case: Uses a static neutral threshold.
- Stationary case: Recognizes the unique structural nature of each sector.

In the Stationary case, γ is defined as the long-term probability of the growth state. This stationary distribution, $\tilde{\pi}$, is obtained by solving:

$$\tilde{\pi}\mathbf{M} = \tilde{\pi}, \quad \sum_{h=0}^H \tilde{\pi}_i = 1$$

It represents the probability of the sector being in a specific state over the long run, regardless of its current condition. Consequently, if the following sector-level constraint is satisfied:

$$\bar{\pi}_{T+K}^{Growth} > \tilde{\pi}^{Growth}$$

the stocks within that sector are permitted in the portfolio. If this inequality is not met, the sector's allocation is not redistributed among other assets but is instead held in cash, providing a strategic defensive advantage during turbulent market periods.

3.4. Comparison with Existing Literature

Nevertheless a large part of literature investigates the use of HMMs for modeling market regimes, relatively few studies explore their application in stock selection and even fewer in the context of momentum investing.

An early contribution in this direction is provided by [19], where a Markov-switching framework trained on macroeconomic variables is employed to perform stock selection within the S&P 500 universe. For each macroeconomic input, the regime for the following month is predicted, and historical periods characterized by similar regime configurations are identified. Based on these analogues, a score is computed for each stock in the index and only the top 50 ranked securities are included in the portfolio. A similar methodology is adopted in [20], where the selection procedure is extended to the an All Country World Index.

With specific reference to the integration of HMMs into momentum strategies, one of the main contributions is [26], which also serves as an initial inspiration for the present work,

although methodological modifications are introduced here. The investment universe considered in that study consists of approximately 900 Korean stocks from January 2000 to December 2018. The Hidden Markov Model is employed in three distinct ways. First, using returns computed over formation periods ranging from one to twelve months, stocks classified as being in a rising regime are selected to construct an equal-weighted portfolio. Second, sector indices are incorporated, and only stocks in a rising state within rising sectors are selected. Finally, the approach is further conditioned on the aggregate market regime, proxied by the KOSPI index. Overall, the HMM is primarily used as a timing and filtering device, determining whether momentum signals should be activated under specific regime configurations.

In contrast, as extensively discussed in the previous sections, the present study embeds the estimated parameters and probabilistic outputs of each HMM directly into the construction of the ranking metrics. Rather than using the model solely as a binary regime filter, regime dynamics are incorporated into the estimation of forward-looking expected returns and volatilities through the evolution of predictive regime probabilities. This allows stock ranking to be informed not only by backward-looking information but also by the forward propagation of regime probabilities via the transition matrix. Consequently, regime information affects both asset selection and sector allocation in a fully integrated and internally coherent framework.

4 | Empirical Results

4.1. Implementation Details

Each model is trained using a five-year rolling window of returns and subsequently used to perform inference on the immediately following year. It should be noted that the training phase is conducted on a single sequence constructed by concatenating the return time series of the individual stocks belonging to the same sector. By explicitly specifying the length of each sub-series, it is possible to preserve the desired informational structure: observations are pointwise independent, yet they share a common latent process represented by the market regime. The implementation of the portfolio strategy involves a holding period K of 6 months, accordingly with [18]. The empirical framework considers model specifications with two and three latent regimes. For each configuration, various levels of the probability threshold γ are analyzed to evaluate their impact on stock selection and overall performance.

From an implementation standpoint, the model architecture is organized into two distinct blocks of classes. The first block includes the classes *SectorHMM.py* and *MultiSectorHMM.py*. The *SectorHMM.py* class is responsible for managing individual sectoral models, including estimation procedures, inference, and the calculation of relevant metrics. Instead, the *MultiSectorHMM.py* class is designed for the centralized management of sectoral models and is used to construct monthly portfolios, coordinating information from the individual models.

The second block consists of the classes *SinglePortfolio.py* and *PortfolioManager.py*. The former allows for the explicit modeling of the overlapping portfolios logic described in [18]. The latter manages a deque (double-ended queue) data structure containing *SinglePortfolio* objects and implements the necessary methods for calculating transaction costs as well as for storing and updating the strategy's performance. Access to this information is provided through appropriate *getter* methods. The overall class structure, including attributes and methods, is illustrated in Figure A.23.

4.2. Regime Characterization

To provide an economic interpretation of the hidden regimes, the models analyzed in this section were estimated using the entire return series of the sector index, rather than individual stocks. This choice is motivated by several factors, including the use of a rolling window in the training process, which makes a direct economic interpretation at the individual asset level less straightforward.

It is important to emphasize that these models will not be employed in subsequent analysis. They are used exclusively for descriptive purposes, with the aim of providing an economic perspective on the dynamics associated with the identified regimes.

4.2.1. 2-States HMMs

Using two states allows for a clear identification of the market's bimodal nature. As shown in Table 4.1, State 0 captures periods of stress, characterized by high volatility and negative average returns. Conversely, State 1 represents an expansionary phase common to all sectors, featuring positive average returns and a contained level of uncertainty.

The persistence probabilities, around 90% for all sectors, further highlight the typical phenomenon of market cycle clustering, which is extensively documented in the literature. A detailed sector-by-sector analysis is provided in Section A.1, displaying the model's full transition matrix, the estimated distribution of weekly returns, and the evolution of probabilities and regimes over time.

Sector	μ_0 (%)	σ_0 (%)	$\mathbb{P}_{00}$ (%)	μ_1 (%)	σ_1 (%)	$\mathbb{P}_{11}$ (%)
Comm. Serv.	-11.75	31.42	90.64	15.06	14.13	94.45
Cons. Discr.	-7.38	31.88	91.55	21.45	12.82	94.00
Cons. Staples	-7.72	31.87	91.55	21.07	12.81	94.00
Energy	-7.82	40.11	90.06	18.24	17.20	94.25
Financials	-19.22	48.07	89.21	17.75	14.81	94.69
Health Care	-7.73	26.91	88.82	15.52	11.65	94.03
Industrials	-7.41	34.47	89.88	15.61	13.35	94.69
IT	-19.03	38.36	91.06	26.01	15.29	94.43
Materials	-14.93	34.58	89.36	20.54	14.60	93.55
Real Estate	-16.43	44.14	90.84	18.11	15.23	95.41
Utilities	-21.70	33.81	89.71	15.85	13.15	95.48

Table 4.1: 2-States HMM Estimated Parameters: High Vol (0) vs. Low Vol (1).

4.2.2. 3-States HMMs

The introduction of a third state, in addition to the previous two, allows for a more complex representation of market behavior. While the two-state case primarily distinguished between high-volatility versus low-volatility regimes, the addition of a third state enables a more informative classification into bear, bull, and intermediate states.

State 0 is identified as a bear market, characterized by significantly negative average returns and high volatility. State 1 represents an intermediate phase of uncertainty, with average returns near zero but still elevated volatility levels; this phase is typically associated with mean-reverting behavior in the indices. Finally, State 2 represents the most favorable regime for investors, featuring positive average returns and low uncertainty. Similarly to the previous case, a detailed sector-by-sector analysis based on the 3-state model is provided in Section A.2.

Sector	μ_0 (%)	σ_0 (%)	$\mathbb{P}_{00}$ (%)	μ_1 (%)	σ_1 (%)	$\mathbb{P}_{11}$ (%)	μ_2 (%)	σ_2 (%)	$\mathbb{P}_{22}$ (%)
Comm. Serv.	-15.39	36.91	90.20	-2.77	22.76	90.88	16.30	13.68	94.80
Cons. Discr.	-24.30	43.73	89.79	3.05	24.51	91.92	21.43	12.58	94.64
Cons. Staples	-26.05	45.23	89.70	1.97	24.72	92.06	21.11	12.58	94.66
Energy	-19.26	54.82	89.37	5.60	28.87	91.44	17.49	16.49	94.52
Financials	-18.04	81.30	89.76	-5.66	28.51	92.06	18.92	13.57	94.86
Health Care	-18.27	36.97	88.64	0.79	19.12	90.77	16.70	11.32	94.40
Industrials	-35.08	48.80	89.59	4.83	24.16	91.72	16.71	12.85	95.08
IT	-29.28	45.11	90.97	4.77	26.35	90.83	26.17	14.53	94.42
Materials	-39.15	46.89	89.11	1.71	25.68	91.24	20.65	14.30	94.21
Real Estate	-45.66	73.55	89.95	-1.16	29.44	92.02	17.86	14.83	95.82
Utilities	-70.26	52.14	89.27	-1.29	23.06	91.12	16.64	12.71	95.65

Table 4.2: HMM Estimated Parameters (3-State): Bear (0), Mid (1), and Bull (2).

4.3. Backtest Results

The models are utilized to perform stock selection with the objective of constructing a portfolio consistent with the momentum factor. As described earlier, at each rebalancing date, all available stocks are ranked based on their expected Sharpe ratio $\hat{S}R_{i,T}$; subsequently, stocks belonging to the top decile are selected for portfolio construction. As an additional filtering layer, a constraint is imposed requiring that the sector to which each selected stock belongs is in a growth market phase.

In an operational context, it is reasonable to expect a portfolio manager to choose thresholds higher than the Naive case to increase the robustness of the selection. However, in the following analysis, different values of γ are considered. Specifically, using a 0%-threshold allows to isolate the pure contribution of the model as a statistical selection tool.

All reported results are calculated in an out-of-sample context. The strategies obtained through these models are finally compared to the long-only momentum strategy described in [18]. To ensure a consistent comparative framework, as specified in 3.2.2, the momentum benchmark is constructed using the exact same investment universe employed for the training and estimation of the HMM models. The evaluation and ranking of these strategies are conducted using the Sharpe ratio, in accordance with the methodology proposed in [25].

4.3.1. Performance Analysis

Table 4.3 presents the annualized performance metrics for the strategies, net of transaction costs equal to 15 basis points on the value traded at each monthly rebalancing. For both models considered, it is observed that as the probability threshold γ increases, both the CAGR (Compound Annual Growth Rate) and the volatility of the strategies progressively decrease. This behavior is consistent with expectations: the probability constraint tends to exclude the riskiest stocks, leaving a larger portion of the portfolio in non-remunerated cash.

The Sharpe ratio remains substantially stable in the Unfiltered and Naive implementations, while showing a significant increase at $\gamma = \tilde{\pi}_{Growth}$. This result suggests that a more stringent selection based on the probability of being in a bull state allows for an improvement in the strategy's risk-return profile, which represents one of the primary objectives in portfolio management. In this context, a portfolio manager can adopt two distinct approaches:

- perform numerical optimization to identify the probability value that satisfies specific predetermined constraints;
- select the threshold discretionarily, based on their own market view: in phases of high uncertainty, higher values will be prioritized, while in expansionary contexts, more moderate thresholds may be adopted.

The parameters α and β play a central role in terms of outperformance and portfolio diversification, respectively. Within the context of the CAPM (Capital Asset Pricing Model), these quantities are estimated via ordinary least squares regression:

$$r_t = \alpha + \beta r_t^{bmk} + \epsilon_t$$

where r_t represents the strategy's return at time t , r_t^{bmk} the benchmark return, and ϵ_t the error term. In this framework, α measures the excess return generated by the strategy, while β quantifies its market sensitivity.

Generating outperformance while simultaneously maintaining a low correlation with the market—and thus the ability to produce returns under different market conditions—represents one of the main challenges in asset management. In both the 2-state and 3-state models, all strategies exhibit β values below 1, with a decreasing trend as the threshold γ increases. The generated outperformance stands at approximately 2.5% annually, reaching values near 4% for $\gamma = \tilde{\pi}_{Growth}$. These results indicate that as the probability threshold increases,

the strategies tend to become increasingly attractive both in terms of diversification and as a source of additional return within a multi-asset portfolio.

Model	γ	CAGR (%)	Vol (%)	Sharpe	α (%)	β
2-States	0%	7.28	17.32	0.68	2.62	0.96
	50%	6.71	15.91	0.67	2.56	0.86
	$\tilde{\pi}_{Growth}$	6.13	11.59	0.80	3.87	0.57
3-States	0%	7.16	17.20	0.67	2.44	0.95
	33.3%	6.54	15.15	0.68	2.62	0.81
	$\tilde{\pi}_{Growth}$	5.78	10.76	0.81	3.82	0.51
Benchmark	MOM	5.76	16.64	0.57	-	-

Table 4.3: Annual Performance Comparison: 2-States vs. 3-States HMM.




Figure 4.1: 2-State HMM Strategies - Cumulative Return




Figure 4.2: 3-State HMM Strategies - Cumulative Return

In Figure 4.1 and Figure 4.2, the performances generated by the individual strategies can be compared with the benchmark. As previously noted, lower threshold values favor higher returns, at the cost of greater volatility. Conversely, higher thresholds, as in the Stationary case, tend to improve the overall risk-return profile.

Figure 4.3 and Figure 4.4 display the returns of the different strategies side-by-side, aggregated on an annual basis. Beyond the numerical values, it clearly emerges that there is not one universally optimal threshold level γ . Indeed, a dynamic adaptation of this parameter to different economic cycles could lead to better results in terms of overall returns. As an illustrative example, it can be observed that during the major crisis phases of recent years, the use of high thresholds would have made it possible to avoid the most significant drawdowns recorded by global indices. However, during expansionary periods, such choices would have resulted in a performance drag compared to strategies characterized by more moderate thresholds.

As previously discussed, the selection of the γ level can therefore be guided either by integration with other models or by the manager's discretion and sensitivity toward the market context.




Figure 4.3: 2-State HMM Strategies - Annual Returns



Figure 4.4: 3-State HMM Strategies - Annual Returns

4.3.2. Risk Management

In this section, the primary risk management metrics are analyzed and discussed. As previously highlighted, an overall improvement across all considered metrics is observed as the probability threshold increases. In particular, the maximum drawdown – the largest percentage drop from a previous all-time high – shows a significant reduction:

values decrease from approximately -55% in the unconstrained case to about -20% when the most stringent threshold is applied.

Model	γ	maxDD (%)	VaR95% (%)	CVaR95% (%)	Skewness	Exc. Kurt.
2-States	0%	-55.97	-7.25	-11.42	-0.54	2.03
	50%	-52.04	-6.66	-10.31	-0.70	1.82
	$\tilde{\pi}_{Growth}$	-23.34	-4.73	-6.56	-0.20	1.23
3-States	0%	-56.35	-7.21	-11.67	-0.54	2.15
	33.3%	-45.25	-6.33	-9.81	-0.52	1.77
	$\tilde{\pi}_{Growth}$	-19.92	-4.38	-6.12	-0.22	1.66
Benchmark	MOM	-57.14	-7.11	-12.55	-0.70	3.42

Table 4.4: Risk Metrics: 2-States vs. 3-States HMM across Probability Thresholds.

Furthermore, as shown in Figure 4.5 and Figure 4.6, the left tail of the distribution, usually associated with downside risk, is less heavy than that of the benchmark.

Analyzing specific values, the 95% Value-at-Risk (VaR), which measures the maximum expected loss with a 95% confidence level, is slightly worse than the benchmark only in the unfiltered case, while it progressively improves as the threshold γ increases. The behavior of the Conditional Value-at-Risk (CVaR), which measures the average loss in the most extreme scenarios, is different: regardless of the γ value, the CVaR is consistently lower than that of the benchmark, further decreasing as the threshold rises.

Additionally, the skewness and kurtosis of the return distribution are measured. All distributions are less skewed toward negative returns compared to the benchmark, except for the Naive threshold of the 2-states HMM, as indicated by less negative skewness values. Regarding kurtosis, the observations concerning the tails are confirmed: although the monthly distributions are leptokurtic relative to a Gaussian, they are nevertheless more platykurtic than the momentum strategy used as a benchmark.

Overall, these results highlight that the use of HMM models not only improves returns but also contributes to modifying their distribution, mitigating the tail risk components typically associated with a pure momentum strategy.




Figure 4.5: 2-State HMM Strategies - Return Distributions




Figure 4.6: 3-State HMM Strategies - Return Distributions

In Figure 4.7 and Figure 4.8, the so-called underwater plots of the strategies are presented, showing the evolution of drawdowns over time compared to the benchmark. At the extreme values of the threshold γ , the two models exhibit very similar profiles. Differences emerge, however, in the naive cases, particularly during the 2008 financial crisis. In this context, the 3-state HMM model manages to contain the drawdown more effectively, limiting the loss by up to approximately 15% more than the 2-state model. This result is attributable to a superior regime classification capability: the introduction of a third state allows for a more accurate segmentation of market dynamics, enabling a more precise distinction between sectors in growth, decline, or retracement phases.

It is also interesting to analyze the drawdown profile in relation to the 2022 bear market. Observing the charts, it is notable that the strategies exhibit an earlier drawdown compared to the benchmark. Furthermore, the model shows some difficulty in promptly identifying the regime change, thus failing to effectively mitigate the downward phase. While the first phenomenon is primarily attributable to stock selection, the second can be explained by several macroeconomic dynamics. Indeed, the 2022 bear market was linked to the inflationary shock that emerged in late 2021, which forced central banks to

rapidly increase interest rates. This was compounded by rising energy prices, partly due to the conflict in Ukraine. However, unlike other crises, these factors did not reflect on the markets through a sudden and violent shock, but rather through a more gradual process. To confirm this, one can observe the trend of the VIX Index, which measures the 30-day implied volatility derived from S&P 500 options: during 2022, no significant spikes were recorded, in contrast to what was observed in previous crisis phases. For example, during the 2008–2009 financial crisis and the 2020 COVID-19 pandemic, the index reached levels exceeding 80 points. This makes the 2022 period particularly complex for models of this type to detect, as they rely heavily on variations in volatility levels. The absence of a clear structural break in return dynamics reduced the model’s ability to identify the regime change in a timely manner.




Figure 4.7: 2-State HMM Strategies - Underwater plots




Figure 4.8: 3-State HMM Strategies - Underwater plots

4.3.3. Sector Exposure Analysis

This section analyzes the portfolio's exposure to GICS sectors, measured as dollar weights relative to the overall NAV. Considering the results obtained for both models in the $\gamma = 0\%$ case, it emerges that the dynamics of sector allocations consistently reflect the distinction between cyclical and defensive sectors, suggesting that the model effectively captures variations in macroeconomic regimes.

In particular, cyclical sectors—typically more sensitive to economic conditions—show a weight reduction during periods characterized by high volatility and deteriorating market conditions, while increasing their incidence during expansionary phases. This behavior is consistent with an implicit risk-oriented selection, where the model tends to reduce exposure to assets more exposed to the economic cycle when stress signals are present.

A significant example, shown in Figure 4.9 and Figure 4.10, is represented by the Con-

sumer Discretionary sector: its weight decreases near the 2008 crisis, only to progressively increase during the subsequent expansion phase until 2018, highlighting a reallocation consistent with the improving macroeconomic context. Conversely, defensive sectors exhibit the opposite dynamic. Exposure to the Health Care sector, for example, systematically increases during market stress phases and decreases during recovery periods. This behavior suggests that the model, while not explicitly imposing sector constraints, implicitly incorporates a flight-to-quality logic, reallocating capital toward sectors less sensitive to the cycle during times of greater uncertainty.

Overall, such evidence indicates that the HMM-based approach is not limited to the cross-sectional selection of stocks but also induces a sector rotation consistent with market regimes, contributing to both risk management and the improvement of the portfolio's risk-return profile.




Figure 4.9: 2-State HMM Strategies - Sector Exposures




Figure 4.10: 3-State HMM Strategies - Sector Exposures

The use of a higher probability threshold leads to a greater incidence of the cash component within the portfolio. This translates into a significant reduction in drawdowns, as highlighted in Table 4.4, though at the expense of overall profitability, reported in Table 4.3. As previously discussed, the introduction of a third state makes the classification of market regimes more sensitive, while simultaneously making the condition for belonging to the bull market regime more stringent. Consequently, strategies based on 3-state HMM models systematically exhibit a higher cash allocation compared to those based on 2-state models. This factor helps explain why drawdown profiles are generally superior in the former case. In general terms, it is observed that for both models and across all levels of the threshold γ , the portfolios exhibit similar compositions.

However, some differences emerge over time: in the early years of the analyzed period, there is greater sectoral diversification, whereas from 2016, an increasing concentration in cyclical sectors is observed, particularly in the technology sector. This latter piece

of evidence provides further support for the model’s validity, as it reflects the ability to capture structural market dynamics, such as the robust growth of the Information Technology sector that began in those years and continues to this day.

4.3.4. Trading Cost Analysis

This section aims to analyze the impact of trading costs on the considered strategies. As previously described, the cost incurred at each portfolio rebalancing date is calculated as the actual turnover value multiplied by a unit cost per dollar traded.

To this end, various cost levels were considered, ranging from 5 basis points – representing an optimistic lower bound – to 20 basis points, which was selected as a conservative hypothesis to test the robustness of the strategies.

Model	γ	5 bps (%)	10 bps (%)	15 bps (%)	20 bps (%)
2-States	0%	-0.08	-0.16	-0.25	-0.33
	50%	-0.08	-0.16	-0.25	-0.33
	$\tilde{\pi}_{Growth}$	-0.07	-0.15	-0.23	-0.30
3-States	0%	-0.08	-0.17	-0.25	-0.34
	33.3%	-0.08	-0.16	-0.24	-0.32
	$\tilde{\pi}_{Growth}$	-0.07	-0.15	-0.22	-0.30
Benchmark	MOM	-0.15	-0.30	-0.46	-0.63

Table 4.5: Transaction Cost Impact on Annualized Excess Returns.

The values reported in Table 4.5 represent the annualized alpha of the strategies net of transaction costs, with respect to the corresponding gross version for the same threshold level γ . This approach was adopted to effectively isolate the impact of trading costs on strategy returns. As expected, these values are negative, reflecting the performance erosion due to transaction costs.

As highlighted in [23], traditional momentum strategies are generally classified as high-turnover anomalies, as they require high trading frequency to capture stocks with the best recent performance. This generates significant execution costs, which negatively impact the strategy’s performance. This evidence is also confirmed by the analysis of the benchmark used in this study: a consistent performance drag is observed, ranging from approximately 15 bps annually in the 5 bps cost scenario to 63 bps in the 20 bps cost scenario. Furthermore, this underperformance is greater than that observed in strategies based on 2 and 3-state models, for a given level of costs. This result suggests that the momentum signal generated by the model is more stable over time, thus producing a more

reliable alpha.

Analyzing the individual models in detail, it is observed that the impact of costs decreases significantly as the threshold γ increases, due to two main reasons: a reduction in trading activity, also visible in Figure 4.9 and Figure 4.10, and greater accuracy and stability in stock selection.

Overall, the operational efficiency analysis highlights how these types of models are capable of significantly optimizing the impact of transaction costs on strategy performance.

4.4. Model Limitations

While the preceding analysis demonstrates that the use of these models can provide an effective solution to various issues typically associated with the development of momentum strategies, this section aims to highlight and analyze the weaknesses of the proposed framework, as well as to identify potential improvement.

4.4.1. Return Distributions

The implementation of HMM models assumes that, within each state, returns follow a Gaussian distribution. While this hypothesis may serve as a theoretical approximation, it is often unrealistic in practice. Specifically, high-volatility regimes are characterized by numerous tail events that a normal distribution cannot effectively capture. Similarly, more stable regimes exhibit a high concentration of observations around the mean, often featuring kurtosis levels higher than those of a Gaussian distribution.

For this reason, employing alternative distributions capable of modeling skewness and heavy tails – such as the Student’s t – could improve the model’s descriptive power. However, this approach entails an increase in the number of parameters to be estimated, making the optimization procedure more complex and potentially reducing the stability of the model’s convergence.

4.4.2. Stationarity Hypothesis

The use of HMM models assumes the stationarity of the transition matrix between regimes, implying that the transition probabilities between states remain constant over time. Although the models employed in this study are estimated using five-year rolling windows, the structure of financial markets is subject to shifts over time as underlying macroeconomic conditions change. Consequently, the model parameters may only be

partially representative of the specific context in which they are intended to operate.

For this reason, Hidden Semi-Markov Models (HSMM) have been developed with the aim of overcoming the assumption that state durations follow a geometric distribution. HSMMs explicitly model the time spent in each state (*sojourn time*) using an arbitrary distribution. This allows for a more flexible and realistic representation of regime persistence, enhancing the model's ability to accurately describe economic cycles.

4.4.3. Selection of the Number of States

Starting from the idea that a model could potentially be described by an infinite number of states, the primary constraint in this context remains the interpretability of the results. Selecting the number of latent states based on information criteria such as AIC or BIC does not always guarantee a clear economic interpretation of the model and can also be sensitive to the noise present in the data. Conversely, choosing to use a limited number of states – as in the case of the two or three states adopted in this study – allows for greater readability of the results, potentially oversimplifying the underlying dynamics. For this reason, before defining the model structure, it is necessary to carefully analyze the application context and identify the objective of the analysis to properly balance interpretability and descriptive power.

In this analysis, a high number of states would have led to the definition of regimes that are difficult to interpret economically, potentially associated with extreme return values and thus rarely observable. By limiting the number of states, these observations are instead absorbed as extreme events within the state distributions, reducing their impact on the overall model structure and maintaining a more stable and interpretable representation of market regimes.

4.4.4. Liquidity Issues

In the conducted analysis, fixed transaction costs ranging from 5 to 20 basis points per dollar traded were considered. However, in practice, trading costs are not constant over time but tend to fluctuate based on market conditions.

Specifically, during periods of crisis – when the HMM model tends to reduce exposure and thus liquidate assets – market liquidity decreases. This occurs even for stocks belonging to the S&P 500 index, which generally feature high capitalization and robust liquidity. Under such conditions, a widening of the bid-ask spread is typically observed, leading to an increase in implicit execution costs.

Consequently, actual trading costs tend to rise precisely at the moments of greatest need for risk reduction, generating an additional performance decay beyond what was observed in Table 4.5.

5 | Conclusions and future developments

Hidden Markov Models (HMMs) were developed to model sequential data in which a system evolves according to latent states that are not directly observable. Although their early applications were in speech recognition and bioinformatics, their use subsequently expanded into the financial field, thanks to their ability to identify structural changes in market dynamics.

While these models are often employed for ex-post analysis, this thesis has opted to utilize them in an operational stock-picking context, with the aim of developing an alternative to the classic momentum factor introduced by Jegadeesh and Titman [18].

By leveraging the temporal evolution of regime distributions and their inherent persistence, it was possible to construct regime-conditioned metrics for each stock that integrate shared information within the same GICS sector. Aggregating this information through a distribution that reflects the concepts of Regime Momentum, or Regime Inertia, resulted in a more informed and robust ranking mechanism compared to traditional momentum strategies. Consequently, the analyzed models functioned as effective selection tools, capable of generating outperformance without a significant increase in portfolio volatility.

Furthermore, the introduction of probability thresholds for identifying favorable market regimes represented an additional layer of control, making it possible to address one of the primary limitations of momentum strategies known as momentum crash. This constraint not only reduces the correlation with the reference benchmark but also contributes to an overall improvement in risk indicators, which are already superior in the version without such restrictions.

Finally, the model-based strategy proved resilient to another limitation of the momentum factor: high turnover and the resulting high management costs. As previously analyzed, even in the version without operational constraints, the model significantly reduces the performance drag associated with the benchmark's trading costs. The selection process

was tested across various cost levels to evaluate its robustness.

However, these models do not represent a final stage in strategy development. Indeed, the assumptions of Gaussian distributions and regime stationarity constitute a first approximation of financial market reality. Similarly, the modeling of trading costs represents an initial level of abstraction that could be further refined by introducing elements of market microstructure. Finally, the requirement for a five-year price history introduces a significant selection constraint, as it systematically excludes recently listed firms and IPOs from the investable universe. This limitation may prevent the strategy from capturing the growth potential often associated with younger, high-momentum companies, potentially reducing the analyzed opportunity set.

Overall, the results obtained suggest that studying markets through statistical models capable of identifying latent regimes, estimating transition matrices and posterior probabilities, represents a valuable resource for constructing investment strategies that are robust across different market cycles.

Bibliography

- [1] A. Ang and G. Bekaert. International asset allocation with regime shifts. *Review of Financial Studies*, 15(4):1137–1187, 2002.
- [2] P. Barroso and P. Santa-Clara. Momentum has its moments. *Journal of Financial Economics*, 116(1):111–120, 2015.
- [3] L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Annals of Mathematical Statistics*, 41(1):164–171, 1970.
- [4] M. Billio and M. Caporin. Multivariate markov switching dynamic conditional correlation garch representations for contagion analysis. *Journal of Financial Econometrics*, 3(4):542–572, 2005.
- [5] D. Blitz, J. Huij, and M. Martens. Residual momentum. *Journal of Empirical Finance*, 18(6):1121–1137, 2011.
- [6] N. P. B. Bollen. Valuing options in regime-switching models. *Journal of Derivatives*, 2(4):38–49, 1995.
- [7] K. Daniel and T. J. Moskowitz. Momentum crashes. *Journal of Financial Economics*, 122(2):221–247, 2016.
- [8] C. Engel and J. D. Hamilton. Long swings in the dollar: Are they in the data and do markets know it? *American Economic Review*, 80(4):689–713, 1990.
- [9] G. D. Forney. The viterbi algorithm. *Proceedings of the IEEE*, 61(3):268–278, 1973.
- [10] Y. Gao, Y. Han, M. Li, and W. Zhou. Intraday momentum in stock returns. *Journal of Financial Markets*, 40:1–21, 2018.
- [11] A. Goyal, N. Jegadeesh, and A. Subrahmanyam. Cross-sectional and time-series tests of return predictability: What is the difference? *Review of Financial Studies*, 28(7):1993–2028, 2015.

- [12] S. F. Gray. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1):27–62, 1996.
- [13] B. D. Grundy and J. S. 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.
- [14] M. Guidolin and A. Timmermann. Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11):3503–3544, 2007.
- [15] J. D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989.
- [16] J. D. Hamilton and R. Susmel. Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64(1-2):307–333, 1994.
- [17] H. Hong, T. Lim, and J. C. Stein. Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. *The Journal of Finance*, 55(1):265–295, 2000.
- [18] N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1):65–91, 1993.
- [19] N. Nguyen and D. Nguyen. Hidden markov model for stock selection. *Risks*, 3(4):455–473, 2015.
- [20] N. Nguyen and D. Nguyen. Global stock selection with hidden markov model. *Risks*, 9(1):9, 2020.
- [21] R. Novy-Marx. Is momentum really momentum? *Journal of Financial Economics*, 103(3):429–453, 2012.
- [22] R. Novy-Marx. Fundamental momentum: How to profit from value and momentum everywhere. *Journal of Financial Economics*, 116(2):329–350, 2015.
- [23] R. Novy-Marx and M. Velikov. A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1):104–147, 2016. doi: 10.1093/rfs/hhv033.
- [24] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [25] S. T. Rachev, S. V. Stoyanov, A. Biglova, and F. J. Fabozzi. Momentum strategies based on reward-risk stock selection criteria. *Journal of Banking & Finance*, 32(7):1146–1158, 2008.

- [26] H. Ryou, H. H. Bae, H. S. Lee, and K. J. Oh. Momentum investment strategy using a hidden markov model. *Sustainability*, 12(17):7031, 2020.
- [27] H. Wang and J. Xu. Idiosyncratic volatility, momentum, and reversals. *Journal of Financial Economics*, 118(3):507–525, 2015. doi: 10.1016/j.jfineco.2015.08.002.

A | Appendix A

A.1. 2 States HMMs




Figure A.1: Communication Services




Figure A.2: Consumer Discretionary




Figure A.3: Consumer Staples




Figure A.4: Energy




Figure A.5: Financials




Figure A.6: Health Care




Figure A.7: Industrials




Figure A.8: Information Technology




Figure A.9: Materials




Figure A.10: Real Estate




Figure A.11: Utilities

A.2. 3 States HMMs




Figure A.12: Communication Services




Figure A.13: Consumer Discretionary




Figure A.14: Consumer Staples




Figure A.15: Energy




Figure A.16: Financials




Figure A.17: Health Care




Figure A.18: Industrials




Figure A.19: Information Technology




Figure A.20: Materials




Figure A.21: Real Estate




Figure A.22: Utilities

A.3. Class Diagram




Figure A.23: UML Class Diagram

List of Figures

3.1	Performance Comparison: S&P500 EW vs. S&P500	19
3.2	Point-in-Time GICS sectors market capitalization weight.	21
3.3	Intra-Sector Volatility	21
3.4	Intra-Sector Average Pairwise Correlation	22
3.5	Evolution of the investment universe size used for portfolio construction . .	26
4.1	2-State HMM Strategies - Cumulative Return	38
4.2	3-State HMM Strategies - Cumulative Return	39
4.3	2-State HMM Strategies - Annual Returns	40
4.4	3-State HMM Strategies - Annual Returns	40
4.5	2-State HMM Strategies - Return Distributions	42
4.6	3-State HMM Strategies - Return Distributions	43
4.7	2-State HMM Strategies - Underwater plots	44
4.8	3-State HMM Strategies - Underwater plots	45
4.9	2-State HMM Strategies - Sector Exposures	46
4.10	3-State HMM Strategies - Sector Exposures	47
A.1	Communication Services (2 States)	59
A.2	Consumer Discretionary (2 States)	60
A.3	Consumer Staples (2 States)	60
A.4	Energy (2 States)	61
A.5	Financials (2 States)	61
A.6	Health Care (2 States)	62
A.7	Industrials (2 States)	62
A.8	Information Technology (2 States)	63
A.9	Materials (2 States)	63
A.10	Real Estate (2 States)	64
A.11	Utilities (2 States)	64
A.12	Communication Services (3 States)	65
A.13	Consumer Discretionary (3 States)	65

A.14 Consumer Staples (3 States)	66
A.15 Energy (3 States)	66
A.16 Financials (3 States)	67
A.17 Health Care (3 States)	67
A.18 Industrials (3 States)	68
A.19 Information Technology (3 States)	68
A.20 Materials (3 States)	69
A.21 Real Estate (3 States)	69
A.22 Utilities (3 States)	70
A.23 UML Class Diagram	71

List of Tables

- 3.1 Annualized sector performances and risk statistics, 2000-2026 23
- 4.1 2-States HMM Estimated Parameters: High Vol (0) vs. Low Vol (1). . . . 35
- 4.2 HMM Estimated Parameters (3-State): Bear (0), Mid (1), and Bull (2). . . 36
- 4.3 Annual Performance Comparison: 2-States vs. 3-States HMM. 38
- 4.4 Risk Metrics: 2-States vs. 3-States HMM across Probability Thresholds. . 41
- 4.5 Transaction Cost Impact on Annualized Excess Returns. 48

Acknowledgements

Giunto al termine di questo percorso, desidero dedicare un pensiero a tutte le persone che, in modi diversi, hanno contribuito a renderlo possibile. Dietro ogni traguardo raggiunto ci sono volti, parole e gesti che hanno lasciato un segno. A tutti loro va la mia più sincera gratitudine.

Desidero innanzitutto ringraziare il Professor Daniele Marazzina per il tempo dedicato alla lettura di questo elaborato, per i preziosi spunti e suggerimenti che hanno arricchito il valore del lavoro svolto.

Un sincero ringraziamento va ai miei correlatori, Giorgio e Marco. Grazie per il costante supporto e per l'opportunità di sviluppare questa tesi all'interno della società, ma soprattutto per l'entusiasmo che mi avete trasmesso fin dal primo momento. La vostra fiducia e il vostro coinvolgimento hanno reso questa esperienza ancora più stimolante. Ringrazio inoltre tutti i colleghi che, sin dal primo giorno, mi hanno accolto facendomi sentire parte di un gruppo che va ben oltre il semplice contesto professionale.

Un pensiero speciale va a tutti i miei amici. Non è un'esagerazione dire che in questa tesi c'è un pezzetto di ognuno di voi: ogni chiaccherata, ogni risata e ogni momento condiviso hanno contribuito, direttamente o indirettamente, a portarmi fin qui.

Desidero poi ringraziare di cuore tutta la mia famiglia. Un grazie agli zii, ai cugini e a tutte le persone che, dopo ogni esame, non hanno mai fatto mancare un messaggio di affetto o di congratulazioni. Un ringraziamento speciale va a mia nonna: le sue instancabili preghiere, a giudicare dal risultato, sembrano aver funzionato davvero.

Grazie a mio fratello, per quel suo immancabile "Grande Brother" accompagnato da una pacca sulla spalla dopo ogni esame, a prescindere dal voto.

Grazie Bea, o forse dovrei dire Bubu, per essermi stata accanto in ogni istante, con la parola giusta sia nei giorni felici sia in quelli più complessi. Se chiudo gli occhi ti vedo seduta a pochi metri da me durante il mio esame di maturità e mi viene da sorridere, perché col senno di poi non avrei potuto chiedere una compagna di viaggio migliore per affrontare tutto ciò che ne è seguito.

Infine, un grazie immenso ai miei genitori, per la fiducia e per il sostegno incondizionato dimostratomi in questi anni ma soprattutto per avermi insegnato che un traguardo non rappresenta mai un punto d'arrivo, bensì la linea di partenza per un passo successivo ancor più importante.

Vorrei concludere con una citazione che ho letto di recente e che spero possa essere d'ispirazione per chi la leggerà, così come lo è stata per me durante la stesura di questo lavoro:

«Vinca dunque la perseveranza, perché, se la fatica è tanta, il premio non sarà mediocre. Tutte le cose preziose son poste nel difficile.»

Giordano Bruno, *La cena de le ceneri*

