Incorporating Market Regimes into Large-Scale Stock Portfolios: A Hidden Markov Model Approach

Ibanez, Francisco and Urga, Giovanni

July 2024

Incorporating Market Regimes into Large-Scale Stock Portfolios: *A Hidden Markov Model Approach*

Francisco A. Ibanez *   Giovanni Urga†

Abstract

We propose a portfolio construction method that accounts for the regime-dependent behavior of stocks, thereby impacting their expected returns. Using a hidden Markov model (HMM) and a regime-weighted least-squares approach, we estimate forward-looking regime-conditional factors. These factors help build large-scale stock portfolios for systematic investment management, considering financial market regimes. In historical simulations, our framework achieves superior risk-adjusted performance compared to passive portfolios in both relative and absolute management settings.

Keywords: Regime modeling, portfolio construction, hidden Markov model, least-squares, factor models

JEL Classification: C1, C2, C32, C4, C61, C63, E32, G11

*Bloomberg LP in New York, NY. Email: fibanez13@bloomberg.net
†Bayes Business School, London, UK. Email: g.urga@city.ac.uk

I Introduction

The contribution of this article is two-fold. First, we propose a regime-weighted least-squares method to estimate conditional regime-aware factor loadings based on the *Hidden Markov Model* (HMM) to obtain forward-looking factor loadings for a sizable universe of stocks. Projecting the risk model and expected returns from a factor model is not novel in the context of portfolio construction, however, incorporating market regime dynamics into the projection is absent in both the literature and practical applications. Second, we use these conditional factor loadings to construct large-scale stock portfolios that can be used to systematically manage investments in a regime-aware manner. We extend the existing literature on regime-aware portfolio construction by presenting a framework that addresses the computational and mathematical challenges associated with training and implementing such allocation frameworks in practice. We expect these results to be particularly interesting for market participants such as stock-focused hedge funds, who manage portfolios with hundreds or even thousands of stocks.

The economy is a complex system characterized by cyclical fluctuations, commonly known as the business cycle. These fluctuations significantly impact financial markets due to their inherent interconnectedness. For instance, during recessions, expansionary monetary policy and reduced investor risk appetite influence the behavior of financial assets like stocks and bonds. Market regimes, driven by the business cycle, have attracted significant research interest due to their potential implications for forecasting, investing, and risk management. In this article, we introduce a portfolio construction methodology that incorporates the regime-dependent behavior of stocks and its impact on their expected returns. The methodology proposed in this paper can be used to systematically manage large-scale stock portfolios while accounting for the impact of these financial market regimes. Our framework should be particularly interesting for quantitative asset managers and hedge funds alike, as these market participants tend to hold hundreds and potentially thousands of stocks within their portfolios. We incorporate market regime dynamics into portfolio choice by fitting a multivariate HMM to a parsimonious set of statistically relevant factors. Leveraging the state transition and emission probabilities, we estimate regime-dependent factor loadings for each stock in the investable universe, enabling us to project the regime-dependent joint distribution of factors back into stock-space. In a rich historical sample of several thousands of companies traded in the US, our framework is able to provide better risk-adjusted performance than passively held portfolios when implemented in both relative (against a benchmark) and absolute portfolio management settings.

The academic literature has extensively explored the different applications of regime-switching models in finance. These applications include, but are not limited to, interest rate and yield curve modeling (see, e.g., Ang and Bekaert 2002b, Bansal and Zhou 2002, Dai et al. 2007), regime changes in US monetary policy and other macro variables (see, e.g., Owyang and Ramey 2004, Hamilton 2005, Sims and Zha 2006, Hamilton 2010 Baele et al. 2015), option pricing (see, e.g., Chan 2014, Siu 2014), asset pricing, and portfolio choice. An in-depth review of their use cases and applications can be found in Ang and Timmermann (2012) and Mamon and Elliott (2014).

Within a regime-switching framework, Yang (2023) investigates the impact of business cycle-driven regime shifts on asset prices and return predictability. Using annual U.S. consumption data from 1929 to 2018, the author identifies two distinct regimes: a persistent low volatility, high growth regime, and a less persistent high volatility, low growth regime associated with deep recessions. The author develops a regime-switching asset pricing model with regime-dependent risk aversion and finds that regime-shift risk is a dominant factor in asset prices, especially during economic expansions. The model successfully explains key regime-dependent asset market phenomena, including the increased predictability of stock returns during recessions.

Incorporating market regime dynamics into asset allocation and portfolio construction has garnered significant interest from both academics and practitioners due to the potential for improved financial outcomes through a deeper understanding of regime dynamics in tradable assets. Ang and Bekaert (2002a) demonstrate the benefits of international diversification in the presence of regime changes, with a two-state Markov chain framework that dynamically allocates to US, UK, and German equity indices. Guidolin and Timmermann (2007) explore the implications of market regimes in the joint returns distribution of stocks and bonds. Employing a four-state framework with two transitory states (crash and recovery) and two persistent states (slow growth and bull), the authors find that optimal portfolio weights vary significantly across regimes, suggesting the potential benefits of regime-aware allocation. Bulla et al. (2011) fit a two-state HMM framework on 40 years of returns data from major markets across the United States, Japan and Germany. Their out-of-sample results conclude that an asset allocation strategy based on an HMM can outperform a buy-and-hold strategy. Kritzman et al. (2012) fit a HMM to four economic variables (equity turbulence, currency turbulence, inflation, and economic growth) and used the results to classify observations in their historical sample into two regimes. The study demonstrates that a dynamic strategy, which adjusts asset class allocations within a portfolio based on regime changes, is significantly more effective than static alternatives. Sheikh and Sun (2012) develop a regime-based approach to rebalance a portfolio of five asset classes based on a four-state switching model. Nystrup et al. (2015) proposes a regime-based asset allocation methodology built on a two-state (high- and low-volatility regimes) HMM that allocates between broad stock and bond indices, showing that their approach is able to deliver higher risk-adjusted performance than static strategies. Nystrup et al. (2019) study the use of model predictive control, in combination with forecasts from a multivariate HMM with two regimes and time-varying parameters, to dynamically manage a portfolio that is allocated across multiple asset classes. The author argues that their methodology can limit portfolio losses by adjusting the risk aversion parameter based on realized drawdowns. Kelliher et al. (2022) introduces a novel risk-parity approach based on a five-state HMM, fitted on four macroeconomic factors, which dynamically balances risk across a set of asset classes. The authors argue that their approach can deliver more consistent risk contributions and persistent outperformance, compared to traditional static allocations.

Despite the potential advantages of regime-aware portfolio construction, the estimation of unobservable regime-switching processes presents challenges. The biggest among these challenges is the so-called *curse of dimensionality* highlighted in Filardo (1998). This issue arises because the number of parameters to estimate exponentially increases with the number of hidden states and variables used in the system. For this reason, the previously mentioned studies on regime-aware investing tend to focus on a rather limited set of investable assets, comprising broad indices and asset class-related ETFs. However, these applications have limited practical relevance for market participants, such as stock-focused hedge funds, who manage portfolios with hundreds or even thousands of stocks, as the computational burden of training these models on such large datasets is often prohibitive. However, the low-rank nature of the equity data could allow us instead to select a parsimonious set of k factors and train a regime switching model on this reduced set, rather than over the entire investable universe. Consequently, the return and covariance estimates coming from the HMM trained on the set of k factors can then be projected from factor space back into stock space through a vector of regime-aware factor loadings. These factor loadings contain the sensitivity of each of the n stocks in the investable universe to the k factors, in each of the m hidden states of the process. Projecting risk models and expected returns from a factor model with the intention of addressing the *curse of dimensionality* in portfolio construction is not novel (see, e.g., Perold 1984 and Jacobs et al. 2005). However, incorporating market regime dynamics into the projection is absent in both the literature and practical applications.

We present a systematic sector rotation strategy, commonly implemented by portfolio managers in the financial industry, that systematically allocates to large universe of thousands of stocks, based on the current and expected market regimes identified by our model. In a 20-year long historical simulation, our framework is able to systematically manage a large-scale investable universe that encompasses nearly the entirety of the U.S. stock market. The empirical results indicate that our methodology is capable of providing higher risk-adjusted returns compared to passively held portfolios, at different levels of tracking error budget commonly used by active portfolio managers in the investment industry. Both our framework and its empirical results expand on the current state of the literature by providing a way of implementing regime modeling in the context of a large-scale portfolio, working around the *curse of dimensionality*, and allowing its implementation in a real-world setting.

The remainder of the paper is organized as follows. In Section II, we describe our dataset and analyze the relationship between the business cycle and the behavior of the stock performance. In Section III, we discuss the concepts behind the Hidden Markov models, which will be the backbone of our methodology, and we outline the regime-weighted least-squares methodology we propose to arrive at the stock-level return and variance estimates. Section IV reports a useful application of our regime-aware framework to build large-scale stock portfolios. Section V concludes.

II Stock Performance Across the Business Cycle

II.A. Sourcing and Data Preparation

We work with a US-centric dataset, due to the size of this market, the number of companies listed, its high liquidity, and data availability. The historical performance of our investable universe is made up of US-based common stocks and is sourced from the return files of the Center for Research in Security Prices (CRSP) database. We select all securities identified as common stock (share codes 10 or 11) listed on the NYSE, AMEX, and NASDAQ (exchange codes 1, 2, 3, 31, 32, and 33). For each of these stocks, we collect daily returns

Figure 1: **Reprentativeness of the training and validation samples**

Figure 1 plots the number of stocks contained in each of the samples used in this study, and how much coverage, in terms of total market capitalization, they represent. In order to analyze the regime-dependent behavior of the US equity market, we collect a comprehensive dataset of stocks by combining the Center for Research in Security Prices (CRSP) database and COMPUSTAT data bases. The complete dataset comprises every stock available in these databases, while the investable universe is defined as a subset of stocks that meet data availability criteria. Despite the slightly fewer stocks contained in the investable universe, it is still of the stock market, covering more than 90% of the bread US stock market, in terms of market capitalization, at any point in time. The historical sample is divded into a training sample (from July 31, 1963 to December 31, 2002), which is only used for estimation purposes, and a validation sample (from January 1, 2003 to December 31, 2022), which is used in the empirical out-of-sample exercise presented in this study.

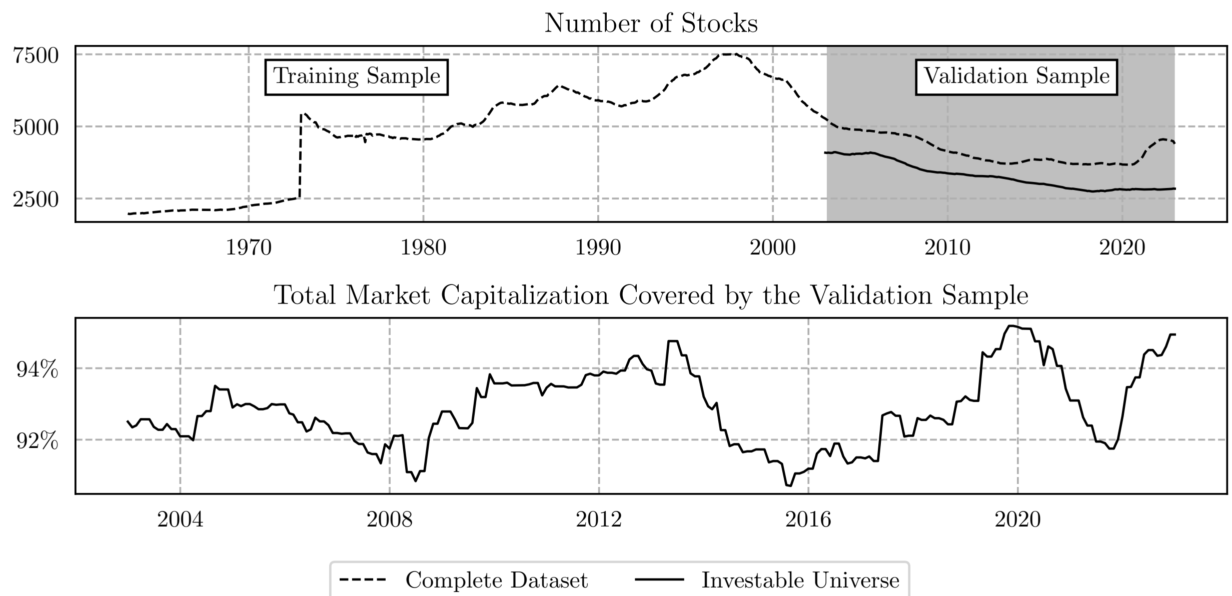

starting in July 1963, and ending in December 2022. We adjust these returns in the case of stock delisting in the same way as described in Shumway (1997), also explained in Bali et al. (2016). From these returns, we subtract the risk-free rate from Kenneth French's website to compute the excess returns. In addition, for each of the stocks in the universe, we also collect time series of daily closing prices (field PRC), number of shares outstanding (field SHROUT), to calculate each stock's market capitalization. Finally, we obtain the Global Industry Classification Standard (GICS) codes from COMPUSTAT, and link them to our return dataset.

We separate the long historical sample into two subsamples; a *training sample* that starts on July 31, 1963 and ends on December 31, 2002, and a *validation sample* that starts on January 1, 2003 and ends on December 31, 2022. The number of companies, and the total market capitalization covered by both samples is visualized in Figure 1. The *training sample* will be used to train and determine the parameterization of the HMM, while the validation sample will be used for the historical portfolio simulation. In terms of size, the complete dataset ranges from 1,965 to 7,513 stocks, averaging 4,180 during the training sample and 4,877 during the validation sample. For the purpose of the empirical exercise to be described below, **we subset** complete dataset based on data availability in order to arrive at our investable universe. We define the investable universe on a given date as every stock in the dataset with valid returns, market capitalization, and industry classification entries on that particular day. On top of this, we filter out stocks with 5 or fewer years of historical data, so we can work with enough degrees of freedom to fit our framework. This leaves us with an investable universe with an average of 3,175 stocks during the validation sample window, ranging from 2,684 to 3,913 companies, covering more than 90% of the total capitalization of the US stock market at each point in time, and averaging nearly 93% of market capitalization coverage during the validation sample. Finally, we downsample the return frequency of our stock dataset from daily to weekly. For the purpose of our exercise, we think that weekly frequency is optimal, as this avoids the intricacies of daily frequency data, such as a large number of outliers and high negative autocorrelation, while still leaving enough us observations to fit the large number of parameters to be estimated. Additionally, we sample the returns on Wednesdays to avoid the so-called day-of-the-week effect that may occur on days such as Fridays and Mondays.

II.B. Distinguishing Defensive from Cyclical Stocks

Fundamentally, the behavior of the business cycle significantly influences the stock market, as corporate earnings are closely tied to consumer spending within the economy.

During periods of economic growth, consumers are likely to spend more freely, using their stable and increasing disposable income on items like technology, luxury goods, and recreational activities. Conversely, when consumers perceive their future earnings as uncertain, they may cut back on non-essential purchases, preserving their disposable income for vital needs such as food and healthcare. Consequently, certain sectors of the stock market are anticipated to outperform during tough economic times (known as defensive sectors), whereas others prosper when economic conditions are positive (referred to as cyclical sectors).

Let us analyze our dataset to identify indications of regime-dependent behavior within the US stock market. Should such patterns exist, we anticipate observing the previously mentioned tug-of-war dynamic between the cyclical and defensive stock sectors. Prior to initiating this analysis, it is imperative to categorize the companies in our sample as either cyclical or defensive. Notably, there is no consensus regarding the number of constituents or the composition of these categories across various financial services firms and market practitioners. Nonetheless, a majority of these actors agree that Consumer Staples, Healthcare, and Utilities can be deemed *defensive sectors*, while Consumer discretionary, Financials, and Materials can be considered *cyclical sectors*. We classify the stocks that are part of our investable universe into these two groups of sectors using the GICS codes from COMPUSTAT.

To differentiate periods of strong economic growth marked by high consumer spending, from periods of economic downturn characterized by reduced consumer expenditure, we utilize the US business cycle dates published by the Business Cycle Dating Committee of the National Bureau of Economic Research (NBER). The NBER compiles a timeline of US business cycles, pinpointing the dates of peaks and troughs that together distinguish economic expansions from recessions. A *peak* is identified as the month when various economic indicators reach their maximum, which is then followed by a noticeable drop in economic activity. Conversely, a *trough* is noted as the month when economic activity hits its lowest, right before it begins to steadily recover. A *recession* is defined as the interval from a peak to its following trough, the lowest point. The phase between a trough and the next peak is considered an *economic expansion*. Typically, the economy is in a state of expansion, with recessions usually short-lived, lasting between 2 (COVID-19 pandemic) and 18 (global financial crisis) months, and averaging 11 months, considering the periods from the beginning of our training sample, until the end of our validation sample.

To assess whether there is differentiated performance of the cylical and defensive sectors in recession and expansion periods, we look at Figure 2. The two panels in the figure show the empirical cumulative density function of the monthly performance of cyclical and defensive stocks in excess of the overall market return, during the expansionary (left) and recessionary (right) months, as defined by the NBER. In economic expansion, the performance of both groups of sectors seems to be visually indistinguishable from each other. However, this changes when the economy is in recession; defensive stocks outperform their counterparts. This is further confirmed by a one-sided Kolmogorov–Smirnov test, which finds statistical evidence of the outperformance of defensive stocks over cyclical stocks, during economic recessions with a 95% confidence interval. The results are in line with economic rationale, as companies within a defensive sector, such as the Consumer Staples sector, for example, provide goods and services that are likely to be resilient to economic contractions because they are deemed essential, such as food and beverages, household products, and clothing. On the other hand, companies in a cyclical sector such as Materials derive their profits from activities such as manufacturing and construction, which are mainly driven by the aggregate level of consumption and interest rates. Consequently, companies in cyclical sectors tend to suffer the most when the economy enters a contractionary state. The cyclicality in the performance of stocks and the uncertainty surrounding this process is commonly referred to as *market regimes*.

At times, the business cycle and the market regimes do not necessarily concur, as stock prices tend to incorporate expectations of the future state of the economy rather than

Figure 2: **Sector performance in economic expansions and recessions (1963-2022)**

Figure 2 plots the cumulative empirical densities corresponding to the performance of the defensive and cyclical sectors during economic expansions and recessions. Defensive sectors are expected to be resilient to economic downturns and include Consumer Staples, Healthcare, and Utilities. On the other hand, cyclical sectors are expected to thrive in expansionary economic regimes and suffer in turbulent market environments, and include sectors such as Consumer discretionary, Financials, and Materials. Using the monthly recession indicator published by the NBER, we separated our long historical sample into periods of economic expansion (left-hand side panel) and economic contraction (right-hand side panel). It becomes apparent that performance of the defensive and cyclical sectors has differed during economic recessions, favoring the former over the latter. This difference is also statistically significant, from a one-sided Kolmogorov–Smirnov test point of view, with a 95% confidence interval, evidencing that some stocks, can provide partial downside protection during economic downturns.

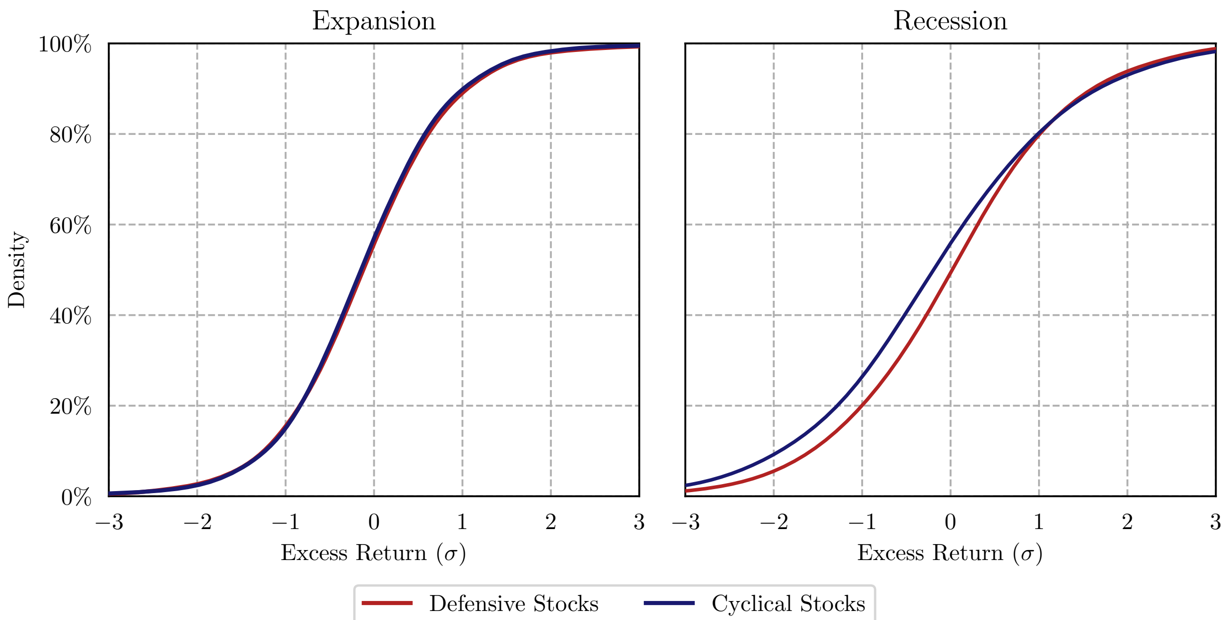

the current one. For this reason, we can consider other variables more closely linked to the market regime, such as the Chicago Board Options Exchange CBOE Martket Volatility Index (VIX), which is usually deemed as a gauge for investor *fear*. Figure 3 shows the empirical joint distribution of the spread between the monthly return of the defensive and cyclical sectors, and the level VIX. In concordance with the case of economic recessions and expansions, the return difference between the defensive and cyclical sectors is centered around zero when the short-term implied volatility, measured by the VIX, is within its interquartile range (13.8 - 23.5). However, in times when the level of market uncertainty has

Figure 3: **Defensive-minus-cyclical return spread and VIX levels (1990-2022)**

Figure 3 displays the joint empirical distribution of defensive-minus-cyclical monthly return spread and VIX levels, using data from 1990 (close to the beginning of the VIX time series) to 2022 (end of our validation sample). Taking a closer look of the joint empirical distribution of defensive-minus-cyclical historical monthly return spread and VIX levels over the last thirty years. In periods where the stock market overall volatility is close to typical levels, the difference between the performance of defensive and cyclical is both centered around 0% and symmetrical, However, conditional to turbulent periods of high market volatility, the return spread distribution shifts towards the positive side, as a result of the protection provided by defensive stocks.

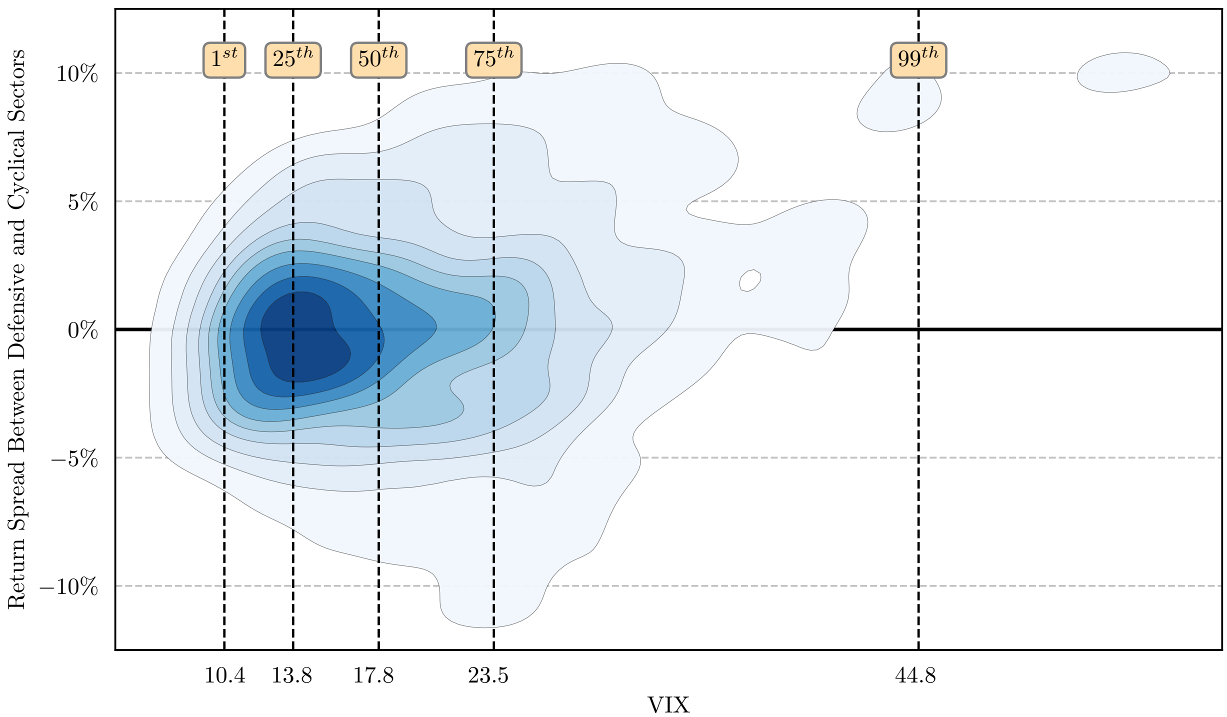

spiked beyond the 75th percentile, the mean returns differential between the defensive and cyclical sectors progressively moves toward positive values with the level of VIX, exhibiting the protection provided by defensive sectors in periods of extreme market volatility.

The main conclusion from the inspection of our dataset is two-fold. First,there seems to be a cyclicality in the performance of stocks that could be linked to the business cycle, which is commonly referred to as *market regimes*. Second, identifying and modeling market regimes, together with understanding the conditional performance of stocks in each of these regimes, could be beneficial for investment managers, as defensive stocks tend to outperform cyclical stocks in dire market environments.

In the next section, we propose a portfolio construction framework that takes into account the regime-dependent behavior of stock prices by modeling and describing the dynamics of the market regimes driving the joint distribution of stock returns.

III The Regime-Conditional Distribution of Stock Returns

III.A. Hidden Markov Models

The distribution of some stochastic processes, such as the performance of a particular financial asset, might not be properly described by a single probability distribution but rather by a combination of densities. Consider the case of stocks, whose returns distribution is known to be empirically left-skewed and characterized by excess kurtosis. Furthermore, and just for the sake of illustration, let us assume that the equity market alternates between market regimes over time. Historically, the large losses that have caused the fat left tail in the empirical distribution of stock returns have occurred primarily in consecutive periods of market turmoil, commonly known as *bear markets*, which are also related to the well-documented volatility clustering phenomenon. In contrast, markets also go through stable growth regimes usually characterized by low stock volatility and positive returns, colloquially dubbed *bull markets*. It would be fair to assume, in this illustration, that the distribution of stock returns is not driven by a single density, but rather by two: a Gaussian distribution for *bull market* regimes with positive mean return and low variability, and another one for *bear market* regimes, located to the left of the previous distribution with significantly higher volatility.

Representing and modeling an overall population distribution as a combination of m densities is known as a *mixture model*. Bouguila et al. (2022) argues that *Hidden Markov Models* (HMM) can be considered an extension of mixture models along the temporal axis, which are capable of modeling and incorporating space-time features. Formally, HMMs are double-stochastic discrete processes within the family of generative machine learning algorithms. They can be trained in both supervised and unsupervised manners to model the dynamics of a *hidden* stochastic process $\boldsymbol{s}$. Although $\boldsymbol{s}$ is not directly observable, it can be inferred through a set of visible stochastic variables $\boldsymbol{\mathcal{Z}}$ emitted by the *hidden* process. In the stock market, regimes such as turbulent *bear markets* and *bull markets* markets are not directly observable or measurable by investors. However, most portfolio managers can make an educated guess of the regime the market is currently going through, by observing the behavior of traditional asset such as equities and bonds, and place bets and construct portfolios based on their observations.

At time t, the *hidden* underlying stochastic process can assume any of the market regimes in $\boldsymbol{s} \in (1, 2, \ldots, m)$. When the process is in state s, a set of observations is emitted by the hidden process (in our case, the stock returns), drawn from the corresponding state-dependent distribution. The hidden process evolves over time in an equally spaced sequence, forming a first-order discrete-time Markov chain (DTMC). A process is said to be a DTMC if it satisfies the so-called *Markovian property*:

$$P\left(s_{t+1} \mid s_t, s_{t-1}, \ldots, s_0\right) = P\left(s_{t+1} \mid s_t\right)$$

where the future state of the underlying stochastic process (i.e., the market regimes in our case) depends only on its current state. The collection of the probabilities of transitioning from and to every state in $\boldsymbol{s}$ is organized and stored in the so-called one-step *transition* *matrix* $\mathbf{\Pi} \in \mathbb{R}^{m \times m}$,

$$\mathbf{\Pi} = \begin{pmatrix} \pi_{11} & \pi_{12} \dots \pi_{1m} \\ \pi_{21} & \pi_{22} \dots \pi_{2m} \\ \vdots & \ddots \quad \vdots \\ \pi_{m1} & \pi_{m2} \dots \pi_{mm} \end{pmatrix}_{m \times m}$$

where $P(s_{t+1} = j \mid s_t = i) = \pi_{ij}$. HMMs can be trained on various component densities, both discrete and continuous. Given our focus on modeling stock returns, we assume that our model combines a finite collection of m Gaussian densities. In the case of continuous return distributions, the component densities overlap, meaning that a single observed realized return z_t could have been emitted by more than one, and potentially each of the m market regimes. The likelihood od a particular observation z_t being emitted by each of the m states is stored in $\boldsymbol{\gamma}_t \in \mathbb{R}^{1 \times m}$, where $\mathbf{1}^T \boldsymbol{\gamma}_t = 1$. In a multivariate framework of k observable variables, the collection of probabilities that a particular realization (row) in $\boldsymbol{\mathcal{Z}} \in \mathbb{R}^{T \times k}$ was emitted when the system was in a given state of $\boldsymbol{s}$ are stored in the corresponding row of the *emission probabilities matrix* $\mathbf{\Gamma} \in \mathbb{R}^{T \times m}$

$$\mathbf{\Gamma} = \begin{pmatrix} \gamma_{\{1\},t-T} & \gamma_{\{2\},t-T} & \cdots & \gamma_{\{m\},t-T} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{\{1\},t-1} & \gamma_{\{2\},t-1} & \cdots & \gamma_{\{m\},t-1} \\ \gamma_{\{1\},t} & \gamma_{\{2\},t} & \cdots & \gamma_{\{m\},t} \end{pmatrix}_{T \times m} \tag{1}$$

where $\gamma_{\{s\}t}$ is the probability that the observations $(z_{1,t}, z_{2,t}, \dots z_{k,t})$ were generated by state $s \in (1, 2, \dots, m)$ at time t. Figure 4 offers a graphical representation of an HMM process. The hidden layer of the process, the market regimes in our case, evolves over time following $\mathbf{\Pi}$. At each time t, the dominant market regime emits the set of observations z_t (e.g., security returns), with probabilities described in γ_t.

Figure 4: **Graphical Representation for Hidden Markov Models**

Figure 4 provides a graphical abstraction of a *Hidden Markov Model* (HMM) that models hidden market regimes and observable asset returns. In a HMM, an unobservable stochastic process s_t evolves through time following a Markov chain, assuming any of the market regimes in $\boldsymbol{s} \in (1, 2, \ldots, m)$. At time t, the hidden variable s_t (i.e., the market regime) emits a realization of an observable process z_t (e.g., stock returns) with emission probability γ. HMMs are particularly useful to describe the dynamics between unobservable market regimes, and the episodic performance of financial assets.

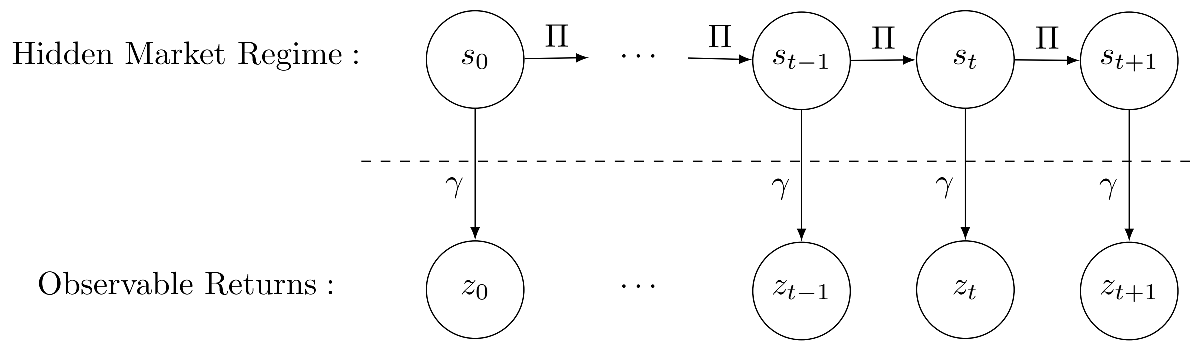

III.B. Regime-Weighted Least Squares

Real-world portfolios often contain a large number of securities. Unfortunately, training an HMM on a large set of stocks as observable variables presents major complications, as the number of parameters to estimate quickly explodes as the dimensionality of $\boldsymbol{s}$ and $\boldsymbol{\mathcal{Z}}$ increases. The low-rank structure of stock return data allows us to select a parsimonious set of k factors and fit the regime-switching model on this reduced set, mitigating the curse of dimensionality associated with modeling the entire investable universe. After fitting the model to the k factors, these regime-dependent estimates can be projected back into stock-space through regime-aware factor loadings. These loadings represent the sensitivity of each of the n stocks in the investable universe to the k factors, in each of the m hidden states of the process.

Consider an investable universe of n stocks (where n is large) whose returns can be explained by a set of k factors (where k is considerably smaller than n). We also assume that the dynamics of these k factors can be characterized by a finite set of m market regimes that evolve in time following a DTMC. Given $\mathcal{F}_t$, let $\boldsymbol{x} \in \mathbb{R}^{T\times 1}$ be the T most recent historical returns of a given stock in the investable universe of size n; $\boldsymbol{\mathcal{Z}} \in \mathbb{R}^{T\times k}$ be the sequences of T historical returns for each of the k factors:

$$\boldsymbol{\mathcal{Z}} = \begin{pmatrix} z_{1,t-T} & z_{2,t-T} & \cdots & z_{k,t-T} \\ \vdots & \vdots & \ddots & \vdots \\ z_{1,t-1} & z_{2,t-1} & \cdots & z_{k,t-1} \\ z_{1,t} & z_{2,t} & \cdots & z_{k,t} \end{pmatrix}_{T\times k}$$

where $z_{i,t}$ is the realized return of factor i at time t. The approriate parameters used in the projection of $\boldsymbol{x}$ onto $\boldsymbol{\mathcal{Z}}$ will be subject to the dominant regime $\boldsymbol{s} \in (1, 2, \ldots, m)$ in which the system is in:

$$\boldsymbol{x} = \begin{cases} \boldsymbol{\mathcal{Z}\theta}_{\{s=1\}} + \boldsymbol{\varepsilon}_{\{s=1\}}, & \text{if } s = 1 \\ \boldsymbol{\mathcal{Z}\theta}_{\{s=2\}} + \boldsymbol{\varepsilon}_{\{s=2\}}, & \text{if } s = 2 \\ \quad\vdots & \\ \boldsymbol{\mathcal{Z}\theta}_{\{s=m\}} + \boldsymbol{\varepsilon}_{\{s=m\}}, & \text{if } s = m \end{cases} \tag{2}$$

where $\boldsymbol{\theta}_{\{s\}} \in \mathbb{R}^{k\times 1}$ is the regime-dependent set of factor loadings within the market regime s, and $\varepsilon_{\{s\}} \in \mathbb{R}^{T\times 1}$ is the vector of residuals of the linear projection of $\boldsymbol{x}$ on $\boldsymbol{\mathcal{Z}}$ within the same regime. By leveraging the emission probabilities in (1) and reorganizing the elements in (2), we can estimate the regime-dependent factor loadings.

Let:

$$\boldsymbol{\mathcal{D}} = \operatorname{diag}\left(\operatorname{vec}\left(\boldsymbol{\Gamma}\right)\right)_{mT \times mT}$$

$$\boldsymbol{M} = \begin{pmatrix} \boldsymbol{\mathcal{Z}} & 0 & \cdots & 0 \\ 0 & \boldsymbol{\mathcal{Z}} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \boldsymbol{\mathcal{Z}} \end{pmatrix}_{mT \times mk}$$

$$\boldsymbol{\delta} = \begin{pmatrix} \boldsymbol{x} & 0 & \cdots & 0 \\ 0 & \boldsymbol{x} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \boldsymbol{x} \end{pmatrix}_{mT \times m}$$

$$\boldsymbol{\Theta} = \begin{pmatrix} \boldsymbol{\theta}_{\{s=1\}} & 0 & \cdots & 0 \\ 0 & \boldsymbol{\theta}_{\{s=2\}} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \boldsymbol{\theta}_{\{s=m\}} \end{pmatrix}_{mk \times m}$$

where block diagonal matrix $\boldsymbol{\mathcal{D}}$ rearrages $\boldsymbol{\Gamma}$ so that each block contains the information stored in the columns of the latter as diagonal entries, both $\boldsymbol{M}$ and $\boldsymbol{\delta}$ are block matrices that store the historical performance of risk factors and stock returns, respectively, $\boldsymbol{\Theta}$ contains the regime-dependent set of factor loadings $\boldsymbol{\theta}_{\{s\}} \in \mathbb{R}^{k \times 1}$ corresponding to regime s, previously introduced in (2), in a block-diagonal way.

We can now obtain the m sets of k regime-aware factor loadings by solving

$$\underset{\boldsymbol{\Theta}}{\text{argmin}} \quad (\boldsymbol{\delta} - \boldsymbol{M\Theta})^T \boldsymbol{\mathcal{D}} (\boldsymbol{\delta} - \boldsymbol{M\Theta}) \tag{3}$$

In the unconstrained case, the first-order condition of (3) yields

$$\frac{\partial}{\partial \boldsymbol{\Theta}} \; \boldsymbol{\delta}^T \boldsymbol{\mathcal{D}} \boldsymbol{\delta} - 2\boldsymbol{\delta}^T \boldsymbol{\mathcal{D}} \boldsymbol{M\Theta} + \boldsymbol{\Theta}^T \boldsymbol{M}^T \boldsymbol{\mathcal{D}} \boldsymbol{M\Theta} = 0$$

$$\boldsymbol{\Theta} = \left(\boldsymbol{M}^T \boldsymbol{\mathcal{D}} \boldsymbol{M}\right)^{-1} \boldsymbol{M}^T \boldsymbol{\mathcal{D}} \boldsymbol{\delta}$$

In the above results, $\boldsymbol{\Theta}$ is a $mk \times m$ matrix that contains the sets of regime-aware factor loadings in a block-diagonal way. To simplify the subsequent analysis, we employ a more compact representation of Θ instead:

$$\hat{\boldsymbol{\Theta}} = \boldsymbol{\mathcal{I}}^T \left(\boldsymbol{M}^T \boldsymbol{\mathcal{D}} \boldsymbol{M}\right)^{-1} \boldsymbol{M}^T \boldsymbol{\mathcal{D}} \boldsymbol{\delta} \tag{4}$$

where

$$\boldsymbol{\mathcal{I}} = \begin{pmatrix} \boldsymbol{I}_k \\ \boldsymbol{I}_k \\ \vdots \\ \boldsymbol{I}_k \end{pmatrix}_{mk \times mk}$$

is used to aggregate the results, so that $\hat{\boldsymbol{\Theta}}$ is a $k \times m$ matrix that contains the estimated factor loading of the returns of the stock of interest to the k factors (along the rows), in each of the m regimes (along the columns):

$$\hat{\boldsymbol{\Theta}} = \begin{pmatrix} \hat{\boldsymbol{\theta}}_{\{s=1\}} & \hat{\boldsymbol{\theta}}_{\{s=2\}} & \dots & \hat{\boldsymbol{\theta}}_{\{s=m\}} \end{pmatrix}_{k \times m}$$

The results of (4) represent a market regime-weighted least squares (RWLS) estimator, which serves as the foundation of our methodology for obtaining regime-aware means and covariances of large-scale stock portfolios.

III.C. From factors to stocks

In a multivariate HMM with Gaussian emissions, the m independent densities are characterized by different distribution parameters. At time t, the probability distributions of the k observable factors will be given by

$$p\left(\boldsymbol{z}_t \mid \mathcal{F}_t\right) = \sum_{s=1}^{m} \gamma_{\{s\},t}\, p\,\left(\boldsymbol{z} \mid \boldsymbol{\mu}_{\{s\}}, \boldsymbol{\Sigma}_{\{s\}}\right) \tag{5}$$

where $\gamma_{\{s\}t}$ is the probability at time t that the observed factor returns $\boldsymbol{z}_t$ were emitted by regime $s \in (1, 2, \ldots, m)$, $\boldsymbol{\mu}_{\{s\}} \in \mathbb{R}^{k\times 1}$ is the vector of expected returns for each of the k factors in regime s, and $\boldsymbol{\Sigma}_{\{s\}} \in \mathbb{R}^{k\times k}$ their variance-covariance matrix within that regime. Under (5), the conditional mean $\boldsymbol{\mu}_t$ and covariance $\boldsymbol{\Sigma}_t$ of the k observable factors in the HMM are given by

$$\begin{aligned} \boldsymbol{\mu}_t &= \sum_{s=1}^{m} \gamma_{\{s\}t} \boldsymbol{\mu}_{\{s\}} \\ \boldsymbol{\Sigma}_t &= \sum_{s=1}^{m} \gamma_{\{s\}t} \left(\boldsymbol{\Sigma}_{\{s\}} + \boldsymbol{\mu}_{\{s\}} \boldsymbol{\mu}_{\{s\}}^T\right) - \boldsymbol{\mu}_t \boldsymbol{\mu}_t^T \end{aligned} \tag{6}$$

In (6), even though the state-dependent densities corresponding to each of the m regimes in the system, characterized by $\boldsymbol{\mu}_{\{s\}}$ and $\boldsymbol{\Sigma}_{\{s\}}$, are constant over time, the conditional joint distribution of the k common factors will evolve driven by the emission probabilities of the HMM, $\gamma_{\{s\}t}$, which act as mixing weights.

On a given rebalance day, a portfolio manager may seek to estimate the joint return distribution of stocks in their investable universe, with the next rebalance date as the investment horizon. Therefore, it is crucial that the regime-aware factor loadings not only reflect the most likely current regime but also account for the probability of transitioning into other regimes during the investment holding period. Let $\boldsymbol{\gamma}_t$ be the most recent emission probabilities given $\mathcal{F}_t$ (i.e., the last row of $\boldsymbol{\Gamma}$ as of the rebalance date), and let us assume that the next portfolio rebalancing will occur in h periods ahead. Taking into account the transition probabilities from t to $t+h$, our forward-looking regime-aware factor loadings will finally be given by

$$\hat{\boldsymbol{\theta}}(h)_{i,t} = \boldsymbol{\gamma}_t \boldsymbol{\Pi}^h \hat{\boldsymbol{\Theta}}^T \tag{7}$$

where $\hat{\boldsymbol{\theta}}(h)_{i,t} \in \mathbb{R}^{1\times k}$ contains the regime-aware forward-looking sensitivity of stock i to the set of k factors, with a horizon of h periods ahead. For the sake of simplifying the notation going forward, we define

$$\boldsymbol{\Phi}(h) = \begin{pmatrix} \hat{\boldsymbol{\theta}}(h)_{1,t} & \hat{\boldsymbol{\theta}}(h)_{2,t} & \dots & \hat{\boldsymbol{\theta}}(h)_{n,t} \end{pmatrix}_{k\times n} \tag{8}$$

which contains every set of forward-looking regime-weighted factor loadings for each of the n stocks in our investable universe on each date t. Martrix $\boldsymbol{\Phi}(h)$ can be used to project the estimated regime-aware joint distribution of the five-factor model from (6) onto our investable universe, thereby estimating the expected return of the investable stocks and the systematic component of their covariance matrix. The regime-weighted idiosyncratic component of the covariance matrix of the n stocks in the investable universe can be estimated as follows:

$$\Xi_t\left(h\right) = \begin{pmatrix} \boldsymbol{\gamma}_t \boldsymbol{\Pi}^h \boldsymbol{\xi}_1 \boldsymbol{\Pi}^{hT} \boldsymbol{\gamma}_t^T & 0 & \cdots & 0 \\ 0 & \boldsymbol{\gamma}_t \boldsymbol{\Pi}^h \boldsymbol{\xi}_2 \boldsymbol{\Pi}^{hT} \boldsymbol{\gamma}_t^T & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \boldsymbol{\gamma}_t \boldsymbol{\Pi}^h \boldsymbol{\xi}_n \boldsymbol{\Pi}^{hT} \boldsymbol{\gamma}_t^T \end{pmatrix}_{n\times n}$$

where

$$\boldsymbol{\xi}_i = \begin{pmatrix} \boldsymbol{\varepsilon}_{i,\{s=1\}}^T \boldsymbol{\varepsilon}_{i,\{s=1\}} & 0 & \cdots & 0 \\ 0 & \boldsymbol{\varepsilon}_{i,\{s=2\}}^T \boldsymbol{\varepsilon}_{i,\{s=2\}} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \boldsymbol{\varepsilon}_{i,\{s=m\}}^T \boldsymbol{\varepsilon}_{i,\{s=m\}} \end{pmatrix}_{m \times m}$$

and

$$\boldsymbol{\varepsilon}_{i,\{s\}} = \boldsymbol{\mathcal{Z}} \hat{\boldsymbol{\theta}}_{i,\{s\}} - \boldsymbol{x}$$

is the vector of residuals corresponding to stock i in the investable universe, for each state $s \in (1, 2, \ldots, m)$.

Constructing regime-aware portfolios with a large number of stocks is now feasible by leveraging the equations above. We estimate the regime-conditional joint distribution of the k factors using (6) and then project this distribution onto our large investable universe using (8):

$$\begin{aligned} \boldsymbol{\psi}_{t,t+h} &= \boldsymbol{\Phi}(h)^T \boldsymbol{\mu}_t \\ \boldsymbol{\Omega}_{t,t+h} &= \boldsymbol{\Phi}(h)^T \boldsymbol{\Sigma}_t \boldsymbol{\Phi}(h) + \boldsymbol{\Xi}_t \left(h \right) \end{aligned} \tag{9}$$

In (9), the terms $\boldsymbol{\psi}_{t,t+h} \in \mathbb{R}^{n \times 1}$ and $\boldsymbol{\Omega}_{t,t+h} \in \mathbb{R}^{n \times n}$ correspond to the projected conditional vector of means and the covariance matrix, respectively, of the investable set at time t with a horizon of h periods ahead. These two terms can now be used in any mean-variance optimization framework to build regime-aware stock portfolios, regardless of the number of stocks that are part of the investable universe at time t.

In the next section, we present an implementation of the framework introduced in this section, in the shape of a systematic sector rotation strategy that leverages the regime-dependent behavior of stock prices in the context of a large-scale investable universe which encompasses nearly the entirety of the US stock market.

IV Building regime-aware stock portfolios

Sector rotation is a widely used portfolio allocation strategy that adjusts sector weights (e.g., Industrials, Consumer Staples) based on the current and anticipated market environment. We expect our framework to accurately identify market environments through regime modeling and adjust portfolio allocations accordingly, overweighting sectors that thrive in the current regime and underweighting those that struggle. In this empirical experiment, we aim to trade a vast number of stocks (in the thousands) on each rebalance date, covering nearly the entire stock market capitalization. We use the previously described methodology to estimate means and covariances for the investable universe and then implement the sector rotation strategy based on these estimates. Working with large investable universes in a mean-variance context presents challenges, such as singularity issues arising from ill-defined sample variance-covariance matrices and substantial estimation errors in mean and variance estimates. We anticipate that our factor-based framework can address the invertibility of the variance-covariance matrix, while the regime-aware estimates can act as a form of shrinkage for estimation errors, leading to superior risk-adjusted performance compared to a passively held portfolio.

IV.A. Portfolio Rebalancing Process

For this exercise, we set up a monthly rebalance schedule that takes place during the validation sample which starts on December 31, 2002, and ends on December 31, 2022, with the target weight calculation taking place on the last business day of the month. Given that our stock dataset is US centric, we choose the five-factor model of Fama and French (2015) as the parsimonious set of factors on which we will train the regime-switching model.

The five-factor model adds profitability (RMW) and investment (CMA) to the classical and widely accepted three-factor model of Fama and French (1993) that includes market excess returns (Mkt), size (SMB), and value (HML) as risk factors. The authors argue that the model is capable of explaining between 71% and 94% the cross-sectional variance of their dataset of companies listed on US stock exchanges. We obtain the returns for the five factors, along with the risk-free rate, from Kenneth French's website at Dartmouth. "We collect daily returns from July 1964 (the earliest available data point) to December 2022. We downsample to weekly frequency to match our stock returns dataset's frequency.

The parameter estimation process at the beginning of each rebalance date t is as follows:

1. We fit a multivariate HMM with Gaussian emissions on the five-factor model using the most recent available observations, as of t. Given that we are using weekly estimates, but we have monthly rebalances, we set $h = 4$ to obtain forward-looking estimates that also account for the probability of potentially transitioning into other market regimes during the coming month.

2. We define our investable universe as the stocks available in our validation sample that have at least five years of weekly returns as of that day, in addition to having market capitalization information and a valid GICS industry classification. The size and composition of the investable universe are expected to vary across rebalance dates due to factors such as company delistings and additions throughout the validation sample.

3. We leverage $\boldsymbol{\gamma}_t$ and $\mathbf{\Pi}$, coming from the most recent HMM estimation, and cycle through each of the n stocks in the investable universe available on t, to estimate its regime-conditional factor loadings $\mathbf{\Phi}(h)$ as described in (7).

4. Then, we combine these estimates with the most recent regime-conditional factor mo-

ments $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$ to arrive at the regime-conditional joint distribution of our investable universe.

5. Finally, we use these estimates to tilt the industry in our portfolio in a regime-aware manner.

IV.B. Market Regime Calibration and Identification

We begin by determining the optimal number of market regimes (hidden states) that best characterize the dynamics of the five-factor model. For this purpose, we set up a grid search of possible integer values ranging from 2 to 10 regimes. Modeling the dynamics of stock returns with more than 10 regimes is not only excessive and lacks widespread support in the literature but also renders estimation infeasible due to the limited number of observations relative to the parameters to be estimated. We iteratively train a multivariate Gaussian HMM on the five-factor training dataset from Kenneth French's website, varying the number of regimes in each iteration. For each iteration, we compute the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC). Figure 5 shows that both information criteria decrease when four hidden states are used in the estimation, before increasing due to the penalty for model complexity associated with additional states, particularly in the case of BIC. Based on these results, we select four as the optimal number of states to describe the dynamics in our factor sample, consistent with findings in the financial literature. This number of states remains fixed throughout the simulation.

While not essential for implementing our framework, an economic interpretation of the four regimes may be of interest. One approach is to examine the resulting regime-dependent distributions of each of the five factors. Figure 6 presents the regime-conditional Gaussian distributions for each of the five factors, across the four identified market regimes. Consider the first column, which displays the return probability densities in *Market Regime 1.* We can observe that the Mkt factor is characterized by a strong positive return, accom-

Figure 5: **Determining the Number of Market Regimes**

Figure 5 plots the information criteria resulting from a search grid designed to find the optimal number of hidden states to use in our exercise. Fitting an HMM on any dataset requires a pre-stablished number of densities to be identified from the set of observable (emmitted) variables as input. Using the training sample window, and the five factor dataset, we iteratively fit an HMM with multivariate Gaussian emissions to a grid of possible integer values ranging from 2 to 10 regimes, and register the Bayesian information criterion (BIC) and the Akaike information criterion (AIC) obtained by using that particular number of regimes. The results suggest that an optimal number of regimes to use for our dataset sits around 4 states, which seems to be in line with the financial literature.

panied by low market volatility. On top of the previous, we also see that in this regime, in the HML factor "growth" companies tend to outperform "value" stocks, and that in the CMA factor companies that are more "aggressive" when it comes to investments outperform "conservative" companies. These are characteristic features of positive trending market states, commonly known as *bull regimes*. Conversely, *Market Regime 4*, depicted in the fourth column, represents the opposite extreme. This regime is characterized by a strong negative return and high volatility in the Mkt factor. Furthermore, value stocks and companies with conservative and robust profiles outperform their counterparts. As expected, companies with robust profitability offer some protection during this turbulent regime and outperform weaker counterparts, as indicated by the positive mean return of the RMW factor. Therefore, we label this negative trending market state as a *bear regime*. Between these extremes lie *Market Regime 2* (second column) and *Market Regime 3* (third column). In *Market Regime 2* the Mkt factor exhibits a positive mean return in a volatile environment, suggesting a stock market rebound or short rally state, which we label the *greed regime*. *Market Regime 3* is characterized by a negative mean return and relatively high volatility, suggesting a short-lived market panic, which we term the *fear regime*. This interpretation is further supported by the evolution of smoothed emission probabilities over time, as depicted in Figure 8. Our *bear regime* coincides with major recession periods, such as the dot-com bubble, the global financial crisis, and the COVID-19 pandemic. The *bull regime*, and to a lesser extent the *greed regime*, tend to coincide with periods of economic expansion. Finally, *fear regime* exhibits a very non-persistent profile, being present for short periods of time and serving as a conduit for transitioning between the aforementioned regimes. The Markov chain that governs the evolution of the process and the transition between these four regimes over time is illustrated in Figure 7. The four states are highly persistent, with a probability between 70% and 95% of staying in the same market regime within a month horizon. This, of course, varies depending on the estimation sample used during the validation sample, but the stickiness of these regimes remains througout the experiment. Within this Markovian framework, is possible to transition from any market regime into another in the next step. However, and as expected, the probability of transitioning from a *bull regime* into a *bear regime*, in only one step, is close to zero. More likely is to use one of the aforementioned transitory regimes, such as the *fear regime*, as an intermediate step before falling into a *bear regime*.

IV.C. Estimating Regime-Aware Factor Loadings

Continuing with our empirical application, we proceed to estimate the conditional regime-aware factor loadings for the stocks in our investable universe. For each weekly observation in the validation sample, we fit our RWLS to the n weekly return time series corresponding to the stocks that are available in the dataset as of that particular day. The

Figure 6: **Factor Regime-Conditional Probability Densities**

Figure 6 exhibits the resulting regime-conditional probability densities for each of the five factors, after fitting a four-state HMM with multivariate Gaussian emissions on these using the training sample window. Each column in the figure collects the distributions of the five factor during that particular market regime. For the sake of interpretability, each of the four regimes was assigned a discretionary label (i.e., Bull, Greed, Fear, and Bear) based on sylized facts of the performance of the five underlying factors during these regimes.

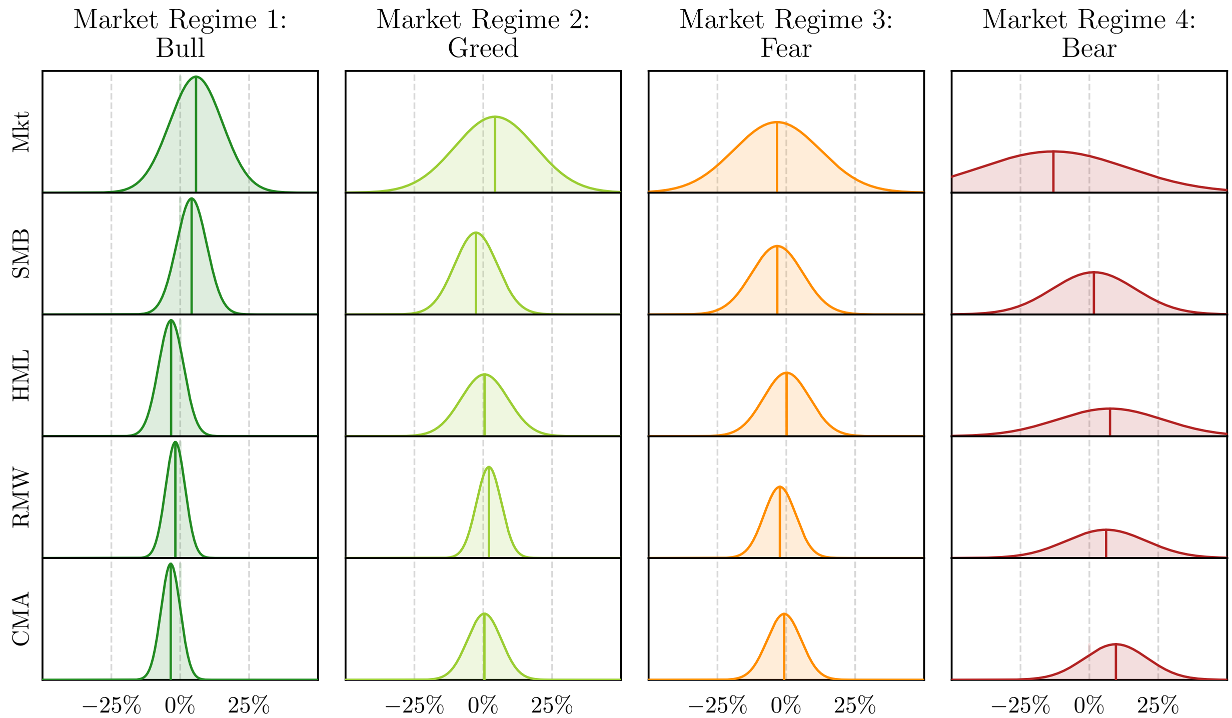

length of these returns time series comprises not only the entire training sample, but also the portion of the validation sample revealed up to time t, avoiding data snooping bias. Consequently, the amount of data available to be used in the estimation grows with the passage of time, as would be the case for a portfolio manager or researcher implementing a framework like ours in a real-life setting. Given the jagged nature of stock data due to the asynchronicity of initial public offerings, the length of the returns time series of each stock will differ, with some stocks having only a couple of observations, while others might exhibit long time series potentially starting in 1963 (the beginning of our training sample). Although emission probabilities $\mathbf{\Gamma}$ are commonly used across stocks when estimating (4), the

Figure 7: **Market Regime Transition Markov Representation**

Figure 7 represents the one-step ahead DTMC that governs the evolution of the HMM process in time. The system consist of four states which are fairly persistent (i.e., the probability of staying in the same regime over the next period is high). Transitioning from one of the extreme regimes (e.g., Bull regime) into the other extreme (i.e., Bear regime) is highly unlikely to materialize, without passing through one of the intermediate states, such as the Fear regime.

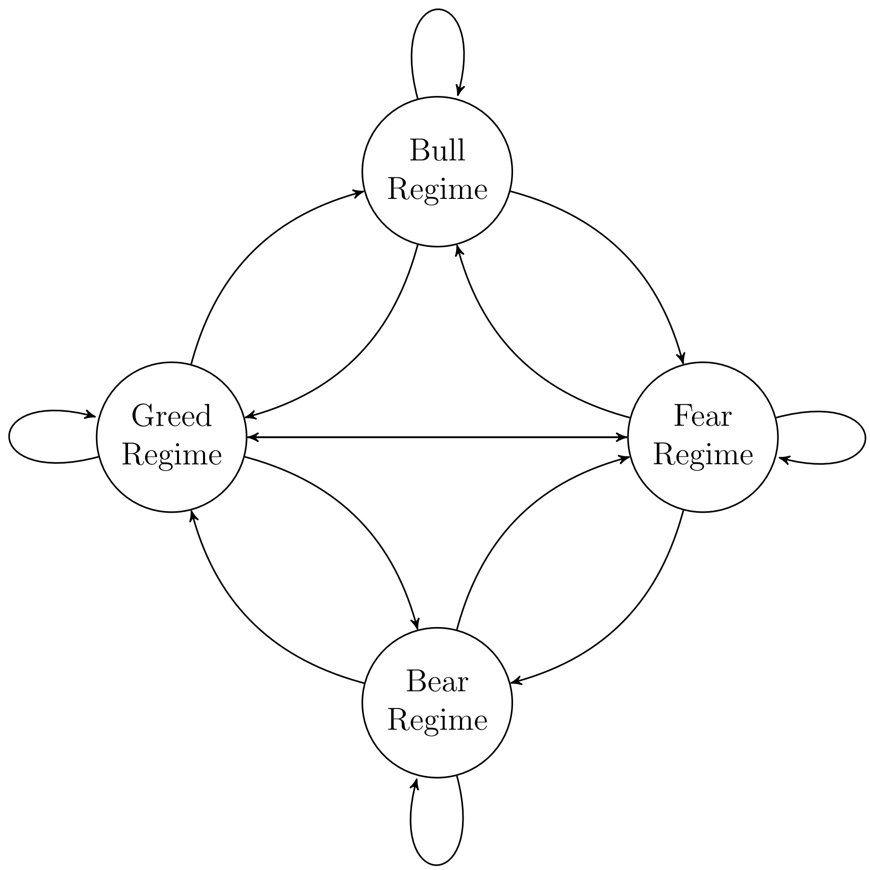

portion of this matrix ultimately used in the estimation of the RWLS of a particular stock, will depend on the length of its time series, using only the overlapping observations given the stock's time series length as of that date. We iteratively repeat this process for each date in the validation sample, extending the sample length by one observation (one week) with each iteration, as new data become available.

Figure 8: **Smoothed Emission Probabilities**

Figure 8 plots the smoothed probabilities that the observed set of realization of the performance of each of the five factor was emmitted by one of the four states in the system, using both the training and validation sample. Our bear regime coincides with major economic recession periods, while our Bull and Greed regimes align with periods of economic expansion, leaving our Fear regime as a conduit for transitioning between the aforementioned regimes

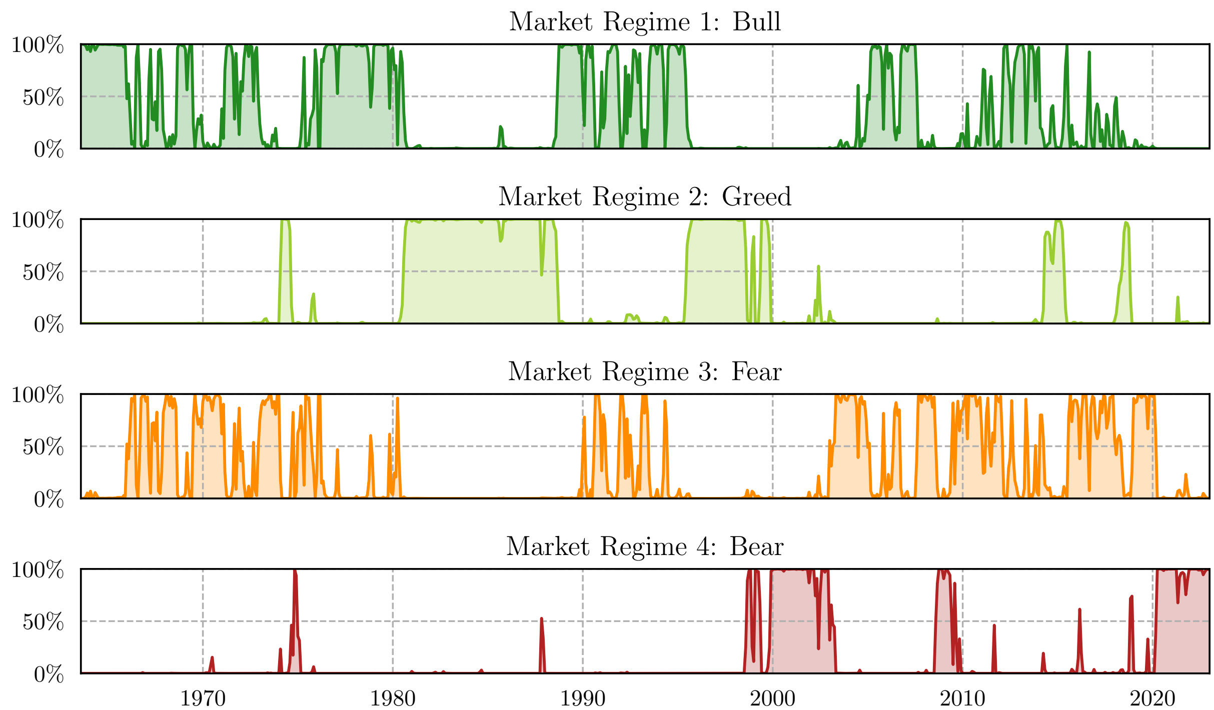

IV.D. Sector-Rotation Strategy

We now proceed to use the estimates obtained in the previous subsections to define and solve our allocation problem. Common implementations of sector rotation strategies over (under) allocate to a subset of the q stock sectors available in the investable universe (e.g., industrials, consumer staples) that the portfolio manager expects to outperform (underperform), with respect to a value-weighted benchmark, while keeping stock weights proportional to capitalization weights within each sector. In this context, the stock weights $\boldsymbol{w} \in \mathbb{R}^{n \times 1}$ to be determined by the portfolio manager can be expressed as

$$\boldsymbol{w} = \boldsymbol{\mathcal{S}}\boldsymbol{\alpha}$$

where $\boldsymbol{\mathcal{S}} \in \mathbb{R}^{n \times q}$ is a matrix that maps the market capitalization weights of each stock (along the rows) to its corresponding sector (along the columns), which is defined as follows:

$$\boldsymbol{\mathcal{S}} = diag\left(\boldsymbol{c}\right)\boldsymbol{Q}\left(diag\left(\boldsymbol{c}^T\boldsymbol{M}\right)\right)^{-1}$$

with $\boldsymbol{c} \in \mathbb{R}^{n \times 1}$ being a vector with the market capitalization of each of the n stocks that form our investable universe, and $\boldsymbol{Q} \in \mathbb{R}^{n \times q}$ is an indicator matrix that maps each stock to its corresponding sector, among the q available sector classifications. Finally, $\boldsymbol{\alpha} \in \mathbb{R}^{q \times 1}$ is a vector of multipliers, chosen by the portfolio manager, that is used to scale up or down the benchmark sector-weights to implement his/her investment bets.

The extent to which the manager can deviate the sector weights in the portfolio from its value-weighted benchmark weights vector ($\boldsymbol{\nu}$) is commonly controlled by a defined risk budget constraint ($\bar{\sigma}$), usually imposed by the investment policy. Putting everything together, we try to solve the following objective function on each rebalance day:

$$\begin{aligned} \underset{\alpha}{\text{argmax}} \quad & \boldsymbol{\psi}_{t,t+h}^T\boldsymbol{\mathcal{S}}\boldsymbol{\alpha} \\ \text{s.t.} \quad & (\boldsymbol{\mathcal{S}}\boldsymbol{\alpha} - \boldsymbol{\nu})^T\boldsymbol{\Omega}_{t,t+h}(\boldsymbol{\mathcal{S}}\boldsymbol{\alpha} - \boldsymbol{\nu}) \leq \bar{\sigma}^2 \\ & \boldsymbol{1}^T\mathcal{S}\boldsymbol{\alpha} = 1 \\ & w_i \geq 0\ ;\ \forall\ i = 1, \ldots, n \end{aligned} \tag{10}$$

Instead of simply choosing a single value for the tracking error budget constraint, we define a grid of possible annual tracking error volatility constraints and analyze the results for this range of values. Huij and Derwall (2011) study the relation between portfolio concentration and performance using a large database of global equity funds. The authors point out that the average fund in our sample has a tracking error slightly higher than 6%, with an interquartile range that goes from 4.5% to 7.5% annualized tracking error. Consistently, we define a sufficiently wide grid of possible annual tracking error volatility constraints that ranges from 2% to 8%, which should encompass the variety of risk budgets implemented by active fund managers. For each value on the grid, we run a historical simulation that implements (10) on each rebalance date.

Figure 9 shows the information ratio measured against a value-weighted portfolio that holds every single stock in the investable universe. Interestingly enough, the information ratio progressively increases with the tracking error budget and reaches its peak around 4% before starting to decrease as the budget continues to increase. This pattern might shed some light on where part the value of this framework comes from. One possible explanation could be that, in a mean-variance setting, most of the risk-adjusted performance added by our framework comes from a more sound way to estimate the conditional risk model. At high enough risk budgets, the impact of the risk model in the resulting target portfolio weights decreases, as the introduced slack of active risk allows the return estimates to have a higher impact in the portfolio solution. Given the higher estimation error surrounding these numbers, the resulting information ratio is lower.

The remaining additional value provided by the framework seems to come from its ability to time turbulent market environments and allocate to assets that might provide protection during these environments. Figure 10 shows the portion of the simulated portfolio allocated to defensive sectors in the upper panel and the probability of the historical filtered emission probability of going through a *bear* market regime in the lower panel. We can make two main observations from this figure. Firstly, these regime probabilities tend to align with stock market crises, such as the global financial crisis of 2008 and the COVID-19 pandemic of 2020. Secondly, and potentially a more relevant finding, the resulting allocation to defensive

Figure 9: **Information Ratio at Different Tracking Error Budgets**

Figure 9 plots the information ratio of the systematic sector rotation exercise shown in this study. The calculations are down in excess of a hypothetical benchmark that holds the entire investable universe proportionally to the market capitalization of the stocks in it, at different tracking error budget levels.

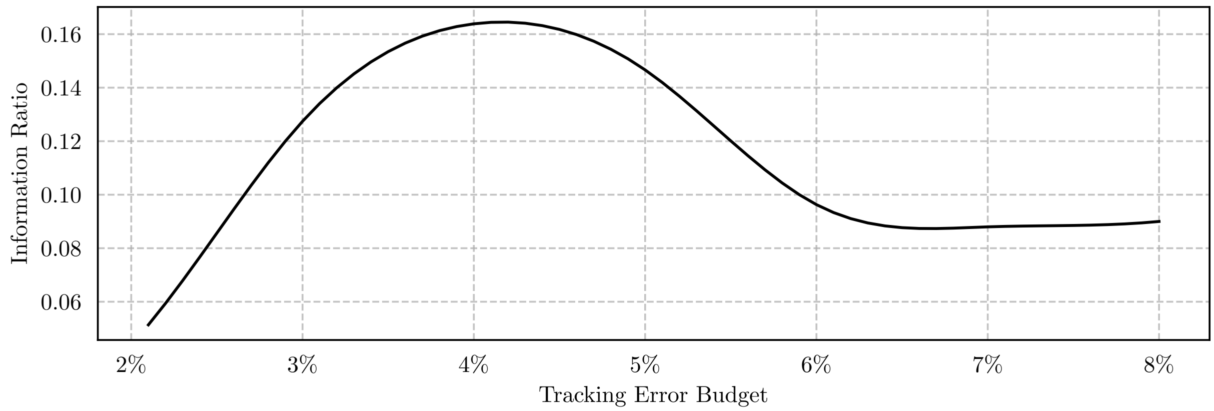

sectors, which tend to perform relatively better during challenging market environments and economic recessions, increases with the probabilities of a *bear* market regime. In light of this evidence, we attribute this outperformance to the ability of our framework of identifying the current market regime, recognizing which sectors are more suited to it, and overweighting them over the ones that do not thrive in those environments.

V Conclusions

In this paper, we introduced a novel methodology for estimating conditional means and covariance matrices that incorporates market regime dynamics and allows to handle a large number of securities. The approach proposed addresses the limitations of traditional mean-variance optimization, particularly in large-scale portfolios. Our framework showed superior risk-adjusted performance compared to passively held portfolios in both relative and absolute management settings. We construct a regime-aware sector rotation strategy that is implemented on a large-scale investable universe, encompassing nearly the entire U.S. stock market over a 20-year period. By utilizing our regime-weighted least squares, we

Figure 10: **Filtered Probability of Risk-Off Regime and Resulting Allocation to Defensive Stocks**

Figure 10 plots the resulting allocation to defensive stocks in the top panel and the filtered probability in the bottom panel. The portfolio allocations used in this figure correspond to the simulated portfolio with a 4% tracking error budget. The panels show how the proposed regime-aware framework increases the allocation to defensive sectors, such as Consumer Staples, Healthcare, and Utilities, during challenging market environments and economic recessions, indicated by the bear market regime probabilities.

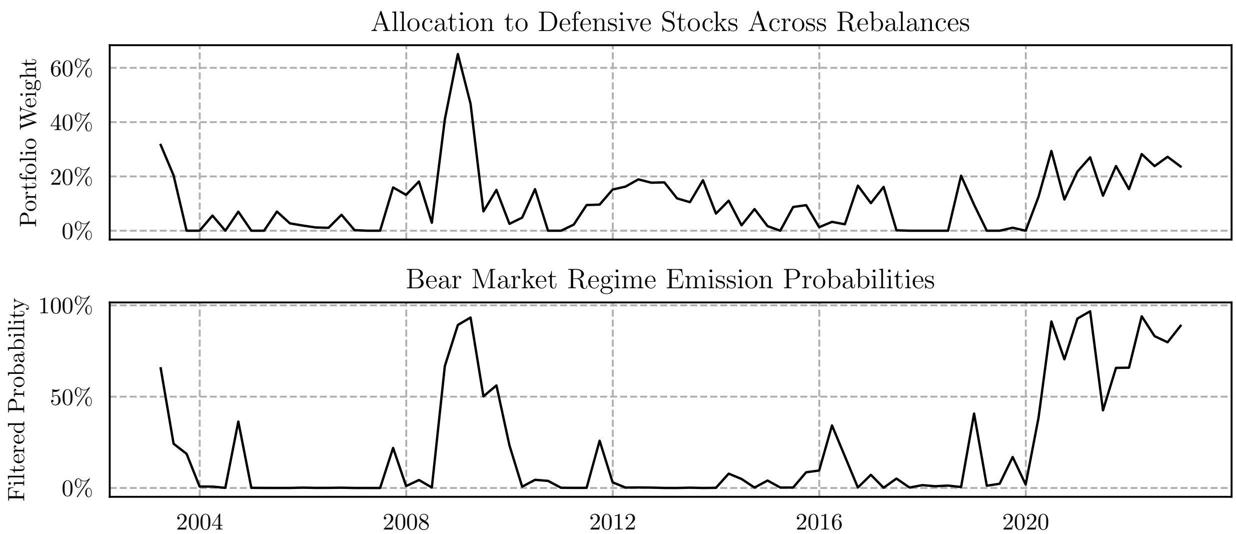

estimated forward-looking, regime-dependent factor loadings for each stock. These loadings, combined with the regime-conditional joint distribution of the five-factor model, enable us to dynamically tilt sector weights based on the expected performance of each sector in the prevailing market regime. The results of our historical simulation demonstrate that this strategy consistently outperforms a passive benchmark, particularly at a 4% tracking error budget. This outperformance is attributed to the model's ability to accurately identify the four market regimes (Bull, Greed, Fear, and Bear) and adjust sector allocations accordingly, with a notable increase in allocation to defensive sectors during turbulent market conditions, as indicated by the bear market regime probabilities.

The main findings in this paper have significant implications for both academics and practitioners, and suggest interesting further developments. It would be interesting to characterize the different features of the market regimes in different time periods. Additionally, our framework assumes exogenous switching, as the hidden state variables governing regime switches do not depend on previous innovations. A clear area of improvement for our framework would be to incorporate endogenous regime switching in the model, so that regime switches can be triggered in the presence of a large innovation, such as a sharp sudden market shock, which is not uncommon in the stock market. This can be achieved through the implementation of the multivariate Markov-switching approach to the asset classes as recently proposed in Kim and Kang (2022), who extend the univariate approach of Hwu et al. (2021). This is part of an ongoing research agenda.

References

Andrew Ang and Geert Bekaert. International Asset Allocation With Regime Shifts. *Review of Financial Studies*, 15(4):1137–1187, 7 2002a. ISSN 0893-9454. doi: 10.1093/rfs/15.4.1137.

Andrew Ang and Geert Bekaert. Regime Switches in Interest Rates. *Journal of Business & Economic Statistics*, 20(2):163–182, 4 2002b. ISSN 0735-0015. doi: 10.1198/073500102317351930.

Andrew Ang and Allan Timmermann. Regime Changes and Financial Markets. *Annual Review of Financial Economics*, 4(1):313–337, 10 2012. ISSN 1941-1367. doi: 10.1146/annurev-financial-110311-101808.

Lieven Baele, Geert Bekaert, Seonghoon Cho, Koen Inghelbrecht, and Antonio Moreno. Macroeconomic regimes. *Journal of Monetary Economics*, 70:51–71, 3 2015. ISSN 03043932. doi: 10.1016/j.jmoneco.2014.09.003.

Turan G. Bali, Robert F. Engle, and Scott Murray. *Empirical Asset Pricing: The Cross Section of Stock Returns*. Wiley, Hoboken, NJ, 2016.

Ravi Bansal and Hao Zhou. Term Structure of Interest Rates with Regime Shifts. *The Journal of Finance*, 57(5):1997–2043, 10 2002. ISSN 0022-1082. doi: 10.1111/0022-1082.00487.

Nizar Bouguila, Wentao Fan, and Manar Amayri. *Hidden Markov Models and Applications*. Springer International Publishing, Cham, 2022. ISBN 978-3-030-99141-8. doi: 10.1007/978-3-030-99142-5.

Jan Bulla, Sascha Mergner, Ingo Bulla, André Sesboüé, and Christophe Chesneau. Markov-switching asset allocation: Do profitable strategies exist? *Journal of Asset Management*, 12(5):310–321, 11 2011. ISSN 1470-8272. doi: 10.1057/jam.2010.27.

Leunglung Chan. An Exact Formula for Pricing American Exchange Options with Regime Switching. pages 211–226. 2014. doi: 10.1007/978-1-4899-7442-6_9.

Qiang Dai, Kenneth J. Singleton, and Wei Yang. Regime Shifts in a Dynamic Term Structure Model of U.S. Treasury Bond Yields. *Review of Financial Studies*, 20(5):1669–1706, 9 2007. ISSN 0893-9454. doi: 10.1093/rfs/hhm021.

Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 2 1993. ISSN 0304405X. doi: 10.1016/0304-405X(93)90023-5.

Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 4 2015. ISSN 0304405X. doi: 10.1016/j.jfineco.2014.10.010.

Andrew Filardo. Choosing information variables for transition probabilities in a time-varying transition probability Markov switching model. 1998.

Massimo Guidolin and Allan Timmermann. Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11):3503–3544, 11 2007. ISSN 01651889. doi: 10.1016/j.jedc.2006.12.004.

James Hamilton. What's Real About the Business Cycle? Technical report, National Bureau of Economic Research, Cambridge, MA, 2 2005.

James D. Hamilton. Regime switching models. In *Macroeconometrics and Time Series Analysis*, pages 202–209. Palgrave Macmillan UK, London, 2010. doi: 10.1057/9780230280830_23.

Joop Huij and Jeroen Derwall. Global equity fund performance, portfolio concentration, and the fundamental law of active management. *Journal of Banking & Finance*, 35(1): 155–165, 1 2011. ISSN 03784266. doi: 10.1016/j.jbankfin.2010.07.032.

Shih-Tang Hwu, Chang-Jin Kim, and Jeremy Piger. AN ¡i¿N¡/i¿ -STATE ENDOGENOUS MARKOV-SWITCHING MODEL WITH APPLICATIONS IN MACROECONOMICS AND FINANCE. *Macroeconomic Dynamics*, 25(8):1937–1965, 12 2021. ISSN 1365-1005. doi: 10.1017/S1365100519000920.

Bruce I. Jacobs, Kenneth N. Levy, and Harry M. Markowitz. Portfolio Optimization with Factors, Scenarios, and Realistic Short Positions. *Operations Research*, 53(4):586–599, 8 2005. ISSN 0030-364X. doi: 10.1287/opre.1050.0212.

Chris Kelliher, Avishek Hazrachoudhury, and Bill Irving. A Novel Approach to Risk Parity: Diversification across Risk Factors and Market Regimes. *The Journal of Portfolio Management*, 48(4):73–90, 2 2022. ISSN 0095-4918. doi: 10.3905/jpm.2022.1.327.

Young Min Kim and Kyu Ho Kang. Bayesian Inference of Multivariate Regression Models with Endogenous Markov Regime-Switching Parameters. *Journal of Financial Econometrics*, 20(3):391–436, 6 2022. ISSN 1479-8409. doi: 10.1093/jjfinec/nbaa021.

Mark Kritzman, Sébastien Page, and David Turkington. Regime Shifts: Implications for Dynamic Strategies (corrected). *Financial Analysts Journal*, 68(3):22–39, 5 2012. ISSN 0015-198X. doi: 10.2469/faj.v68.n3.3.

Rogemar S. Mamon and Robert J. Elliott, editors. *Hidden Markov Models in Finance*, volume 209. Springer US, Boston, MA, 2014. ISBN 978-1-4899-7441-9. doi: 10.1007/978-1-4899-7442-6.

Peter Nystrup, Bo William Hansen, Henrik Madsen, and Erik Lindström. Regime-Based Versus Static Asset Allocation: Letting the Data Speak. *The Journal of Portfolio Management*, 42(1):103–109, 10 2015. ISSN 0095-4918. doi: 10.3905/jpm.2015.42.1.103.

Peter Nystrup, Stephen Boyd, Erik Lindström, and Henrik Madsen. Multi-period portfolio selection with drawdown control. *Annals of Operations Research*, 282(1-2):245–271, 11 2019. ISSN 0254-5330. doi: 10.1007/s10479-018-2947-3.

Michael T. Owyang and Garey Ramey. Regime switching and monetary policy measurement. *Journal of Monetary Economics*, 51(8):1577–1597, 11 2004. ISSN 03043932. doi: 10.1016/j.jmoneco.2004.01.004.

Andre F. Perold. Large-Scale Portfolio Optimization. *Management Science*, 30(10):1143–1160, 10 1984. ISSN 0025-1909. doi: 10.1287/mnsc.30.10.1143.

Abdullah Z. Sheikh and Jianxiong Sun. Regime Change: Implications of Macroeconomic Shifts on Asset Class and Portfolio Performance. *The Journal of Investing*, 21(3):36–54, 8 2012. ISSN 1068-0896. doi: 10.3905/joi.2012.21.3.036.

Tyler Shumway. The Delisting Bias in CRSP Data. *The Journal of Finance*, 52(1):327–340, 3 1997.

Christopher A Sims and Tao Zha. Were There Regime Switches in U.S. Monetary

Policy? *American Economic Review*, 96(1):54–81, 2 2006. ISSN 0002-8282. doi: 10.1257/000282806776157678.

Tak Kuen Siu. A Hidden Markov-Modulated Jump Diffusion Model for European Option Pricing. pages 185–209. 2014. doi: 10.1007/978-1-4899-7442-6_8.

Wei Yang. Business Cycles, Regime Shifts, and Return Predictability. *Journal of Financial and Quantitative Analysis*, 58(7):3058–3084, 11 2023. ISSN 0022-1090. doi: 10.1017/S0022109023000261.