

# Regime-based portfolio optimisation: A Hidden Markov Model approach for fixed income portfolios

---

This paper introduces a regime-based approach for fixed income portfolio optimisation that identifies distinct market volatility regimes during portfolio construction to improve risk targeting and risk-adjusted returns.



Byran Taljaard

European Stability Mechanism

## Disclaimer

This working paper should not be reported as representing the views of the ESM. The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the ESM or ESM policy.

European Stability Mechanism



# Regime-based portfolio optimisation: A Hidden Markov Model approach for fixed income portfolios

---

Byran Taljaard European Stability Mechanism

## Abstract

This paper presents a methodology for incorporating a regime-based approach to portfolio optimisation, specifically for fixed income portfolios. By segmenting the market into distinct periods or "regimes" characterised by different market conditions, such as high or low volatility, investors can adjust their portfolio strategies accordingly. We propose a two-stage approach: first, applying Principal Component Analysis (PCA) to fixed income indices to represent the yield curve, and second, using the variance of the first principal component to fit a Hidden Markov Model (HMM) that identifies high and low volatility regimes. This regime-based methodology allows portfolio managers to make more informed decisions, potentially improving risk-adjusted returns. We apply this approach to portfolio optimisation, targeting a specific value-at-risk (VaR), and compare the results to a traditional approach that does not account for regime changes.

**Keywords:** Regime-based portfolio optimisation, Hidden Markov Model, Fixed income portfolios, Principal Component Analysis, Yield curve, High-volatility regime, Low-volatility regime, Value-at-Risk (VaR), US Treasury bond indices, French government bond indices, German government bond indices, EUR and USD Sovereign, Supranational, and Agency (SSA) bond indices

**JEL Codes:** G11, G12, C32, C38

---

## Disclaimer

This Working Paper should not be reported as representing the views of the ESM. The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the ESM or ESM policy. No responsibility or liability is accepted by the ESM in relation to the accuracy or completeness of the information, including any data sets, presented in this Working Paper.

© European Stability Mechanism, 2025 All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the European Stability Mechanism.

# Regime-based portfolio optimisation: A Hidden Markov Model approach for fixed income portfolios

Byran Taljaard<sup>1</sup>, European Stability Mechanism

## Abstract

This paper presents a methodology for incorporating a regime-based approach to portfolio optimisation, specifically for fixed income portfolios. By segmenting the market into distinct periods or "regimes" characterised by different market conditions, such as high or low volatility, investors can adjust their portfolio strategies accordingly. We propose a two-stage approach: first, applying Principal Component Analysis (PCA) to fixed income indices to represent the yield curve, and second, using the variance of the first principal component to fit a Hidden Markov Model (HMM) that identifies high and low volatility regimes. This regime-based methodology allows portfolio managers to make more informed decisions, potentially improving risk-adjusted returns. We apply this approach to portfolio optimisation, targeting a specific value-at-risk (VaR), and compare the results to a traditional approach that does not account for regime changes.

**Keywords:** Regime-based portfolio optimisation, Hidden Markov Model, Fixed income portfolios, Principal Component Analysis, Yield curve, High-volatility regime, Low-volatility regime, Value-at-Risk (VaR), US Treasury bond indices, French government bond indices, German government bond indices, EUR and USD Sovereign, Supranational, and Agency (SSA) bond indices

**JEL Codes:** G11, G12, C32, C38

---

<sup>1</sup> b.taljaard@esm.europa.eu. The author would like to thank Sebastien Levy for the insightful discussions on this topic that greatly enriched the development and refinement of this paper, and Kalin Anev Janse and Karol Siskind for their helpful comments and suggestions.

Disclaimer: This Working Paper should not be reported as representing the views of the ESM. The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the ESM or ESM policy. No responsibility or liability is accepted by the ESM in relation to the accuracy or completeness of the information, including any data sets, presented in this Working Paper.

## Table of Contents

Abstract.....	1
Introduction.....	3
Section 1: Principal Component Analysis.....	5
Section 2: Hidden Markov Models .....	14
Section 3: Portfolio Optimisation .....	19
Section 4: Conclusion .....	27
References.....	29
Annex A .....	30

## Introduction

In this paper, we propose a methodology to incorporate a regime-based approach to portfolio optimisation with a specific focus on fixed income portfolios. A regime-based approach involves segmenting the market into distinct periods or "regimes", characterised by different market conditions, such as high or low volatility. By identifying these regimes, investors can adjust their portfolio strategies to better align with the prevailing market environment. For example, during a high-volatility regime, a portfolio manager might increase holdings in safer assets like government bonds to mitigate risk, whereas in a low-volatility regime, they might allocate more to higher-yielding assets to enhance returns.

While regime-based approaches are not new, there is little analysis on their specific implementation within a fixed income universe. Typical approaches have focused on equities or on economic regimes across broad asset classes. Portfolios within an official institution setting require a slightly different approach, as they are typically only fixed income based and are rebalanced less frequently than the common daily frequency.

We propose a two-stage approach to regime identification on a set of fixed income indices. In the first stage, we apply Principal Component Analysis (PCA) on the universe of fixed income indices to obtain a representation of the yield curve. The first principal component and its variance (the eigenvalue) effectively capture the overall level of interest rates and their variability, providing a concise summary of yield curve movements. We then use the variance of this first principal component to fit a Hidden Markov Model (HMM), which identifies the regimes inherent in the yield curve—specifically, a high-volatility regime and a low-volatility regime.

By utilising financial market data through PCA, we avoid the difficulties of aligning economic data with more frequent financial market data, as economic indicators often have lower frequency and can introduce look-ahead bias. Focusing solely on financial market data ensures that our regime identification is timely and relevant to current market conditions. Identifying these regimes allows us to tailor the portfolio optimisation process to account for the differing risk-return profiles inherent in each regime.

We apply this approach to portfolio optimisation, targeting a specific value-at-risk (VaR), and compare the results to a more traditional approach that does not account for regime changes. Our universe of investments covers US Treasury bond indices, French government bond indices, and German government bond indices separately. We include EUR and USD

Sovereign, Supranational, and Agency (SSA) bond indices, which are used in portfolio optimisation.

We separate the paper into three main sections:

- First Section: We take a detailed look at the results of the PCA on the underlying curves, illustrating how the principal components capture the key movements in the fixed income markets.
- Second Section: We discuss Hidden Markov Models and apply them to the PCA output from the first section, demonstrating how the HMM identifies different market regimes based on the volatility of the principal components.
- Final Section: We use this regime-based approach to perform a target VaR-based portfolio optimisation and compare the resulting portfolio with a more standard five-year historical VaR approach. This is done for a portfolio consisting only of government bond indices and another portfolio that includes an SSA index.

By adopting this regime-based methodology, portfolio managers can make more informed decisions that align with current market conditions, potentially improving risk-adjusted returns in fixed income portfolios.

## Section 1: Principal Component Analysis

PCA is a widely used technique to analyse data by reducing dimensionality while retaining as much of the variation as possible from the original dataset. In the context of fixed income portfolios, PCA transforms the data into a set of uncorrelated principal components (PCs) through eigen decomposition. This process yields eigenvectors and their associated eigenvalues. The eigenvalues represent the variance explained by each principal component, while the eigenvectors provide loadings—or weights—on the underlying data points.

For example, consider a yield curve with tenors at the 2-, 3-, 5-, 7-, and 10-year points. PCA would produce eigenvectors with loadings on each of these tenors, effectively summarising how different maturities contribute to the overall movements of the curve. This is particularly attractive for fixed income portfolios because it allows us to represent complex yield curve dynamics succinctly within a few variables, rather than analysing multiple points individually.


PCA is especially useful in fixed income because the first few principal components often have clear and intuitive interpretations (Litterman, 1991):

- **First Principal Component (Level Factor):** Represents the overall level shift in interest rates across all maturities, typically explaining over 95% of the yield curve's variation.
- **Second Principal Component (Slope Factor):** Captures changes in the steepness of the yield curve, reflecting differences between short-term and long-term interest rates.
- **Third Principal Component (Curvature Factor):** Accounts for changes in the curvature of the yield curve, indicating how medium-term rates move relative to short and long-term rates.

### *Illustrative example*

We perform PCA on the US Treasury curve across the 6-month, 1-, 2-, 5-, 7-, and 10-year maturities over the period from 1 June 2014 to 31 May 2024. By taking the covariance matrix of the daily changes and performing eigen decomposition, we extract the principal components and their associated variances.

**Figure 1: US Treasury curve: First three principal component loadings (June 2014 – May 2024)**




Source: Bloomberg, ESM

In Figure 1, the first principal component shows loadings that are all in the same direction, confirming it as the level shift factor. The second principal component displays loadings where the front-end and long-end of the curve have opposite signs, representing the slope factor. The third principal component's loadings indicate a curvature factor, being long in the belly of the curve and short at the wings.

With six points on the curve, we obtain six principal components, but with declining importance to the overall movements in the yield curve. The first principal component, for example, explains 87% of the variance of the yield curve, while the first three components together account for 99% of the total variance. We can, therefore, concentrate solely on the first three principal components to analyse the overall movements in the yield curve. This demonstrates how PCA effectively reduces the dimensionality of the data while retaining the most significant variations.



**Figure 2: Proportion of variation explained**



Source: Bloomberg, ESM

### *Dynamic Analysis of Loadings*


The results cover the entire period from June 2014 to May 2024, however, if we would like to get a more granular view at the loadings at a specific point in time, we need to perform a rolling analysis. Here we face a slight complication. The orientation of the eigenvectors is arbitrary, meaning the signs of the loadings can flip without affecting the variance explained. This can complicate temporal comparisons. That is, for the first principal component (PC1), which usually describes the level, the loadings could either be all positive or all negative. In other words, while the loadings in Figure 1 for PC1 are all negative, they could have also been all positive. This causes an issue when performing rolling analysis as we look to do here as the signs of the loadings flip from positive to negative at random.

To correct for this, we found two options based on existing literature: using the cosine similarity approach as in (Hirsa, 2023) for example, or the approach of (Ogita T. a., 2016), which uses the previous eigenvector as a starting point to estimate the next one. We describe these approaches in more detail in Box 1.

In our approach, we have implemented the Ogita-Aishima algorithm, although both approaches yield sufficiently accurate results for our purposes. It is not strictly necessary to make this adjustment if we only use the rolling eigenvalues themselves, since the actual orientation of the eigenvectors is irrelevant

for the size of the eigenvalues. It does, however, allow us to analyse the drivers of the yield curve in a more granular manner.

**Figure 3: US Treasury curve: Loadings for the first principal component on a rolling 3-month basis**




Source: Bloomberg, ESM

Figure 3 shows the rolling loadings for the first principal component (PC1) for the 1-, 2-, 5-, and 10-year tenors of the US Treasury curve. There is an interesting behaviour in the loadings over time especially when comparing pre- and post-2008. Whereas pre-2008, the loadings between the front and long end of the curve were in the same -0.3 to -0.5 region, after 2008 the long-end (10-year primarily) becomes more important to the level shift factor. This is especially true from about 2010 onwards, post QE2, when the ten-year point drives most of the movement in the curve, while the front-end (1- and 2-year points) has little influence. This happens again for a second time in 2020 following the Fed's response to the Covid crisis. As the Fed begins hiking rates, for example in the period 2016-2018 and in 2022, the front-end's influence increases particularly in the 2-year point on the curve. This dynamic view reveals which points on the curve are driving the level shifts and how their contributions change over time, especially during periods of monetary policy changes.

#### *Cumulative Returns of Principal Components*

By examining the cumulative returns of the principal components, we can observe how each component captures different aspects of the yield curve movements.

**Figure 4: US Treasury curve: Cumulative returns for the first three principal components<sup>2</sup>**



Source: Bloomberg, ESM


In Figure 4, we show the cumulative returns of the first three principal components. As expected, the first component has the highest volatility and its cumulative return is indicative of the changes in the yield curve through the cycle of hikes (2004-2006, 2016-2018, and 2022 for example) and cuts (2000-2003, 2007-2008, and 2019-2020 for example). The second principal component is indicative of changes in the slope of the curve, and in this case, is a proxy for flattener position given the loadings.

We can perform the same analysis on the German government bond curve where we see a similar profile, although the front-end remains a much lower influence during the 2016-2018 period where the ECB kept policy rates on hold, in contrast to the Fed.

---

<sup>2</sup> It is important to compare the returns to the direction of the principal component loadings. In Figure 4, PC1 has negative loadings on rates (long bond position), PC2 has a positive loading on front-end and negative loading on long-end (flattener position), while PC3 has positive loading on the belly, negative loadings on the wings (benefits when wings outperform the belly).


**Figure 5: German yield curve: Loadings for the first principal component on a rolling 3-month basis**



*Source: Bloomberg, ESM*

We show the cumulative returns for the first three principal components in Figure 6. There are broad similarities between the US and German principal components. Two differences, however, are the strong directional bias to the first principal component in Germany as the curve moved lower and rates stayed low for an extended period. Additionally, as rates were expected to remain low and as the ECB maintained its asset purchase programmes, the second principal component (indicative of slope) exhibits much lower volatility than that of the US.

**Figure 6: German yield curve: Cumulative returns for the first three principal components**




Source: Bloomberg, ESM

#### *Focusing on the Variance of the First Principal Component*

For our regime identification purposes, we concentrate on the variance (eigenvalue) of the first principal component, which serves as a proxy for the volatility of the underlying yield curves. This approach allows us to encapsulate the yield curve's overall volatility into a single variable, simplifying the analysis without losing critical information. We show the variance of the first principal component for both the US and German yield curve in Figure 7.

**Figure 7: German and US yield curve: Variance of first principal component (3-mo rolling)**



Source: Bloomberg, ESM

Figure 7 shows that while the German yield curve has generally been less volatile than the US curve, recent periods like the COVID-19 crisis have seen similar levels of volatility in both markets. By using the variance of the first principal component, we capture these shifts in volatility succinctly.

#### Box 1: Approaches to estimating eigenvectors on a rolling basis

##### *Using cosine similarity:*

In this approach, we take the cosine of the angle between the eigenvector from day to day to determine if the vectors are pointing in the same direction (the loadings have the same sign). This measure will oscillate between 1 and -1 depending on the relative direction of the vectors. For example, we take the first eigenvector (the loadings for the first principal component) from the previous day and compare it to the eigenvector for the current day. If the cosine similarity is less than zero, we can deduce that the two vectors are pointing in opposite directions, and we need to multiply one of them by -1. A good example of this is presented in (Hirsa, 2023).

##### *Using the Ogita-Aishima algorithm (Ogita T. a., 2016):*

This algorithm uses the previous eigenvectors as a starting point for the next eigenvector and then iteratively converges to the next set of eigenvectors based on the estimated eigenvalues. This ensures that the loadings have the same sign since if the signs were flipped on the following day, the associated eigenvector would be further away from the initial eigenvector used (the previous day's eigenvector in our case). An example of the implementation of this algorithm in financial markets is

presented in (Bilokon, 2021). An extension of this algorithm is presented in (Ogita T. a., 2019), which attempts to handle clustered eigenvectors.

Using PCA provides several advantages:

- **Simplification:** It reduces the complexity of analysing multiple yield curve points by summarising them into principal components, particularly focusing on the first component for volatility.
- **Representation of Yield Curve Dynamics:** PCA identifies which maturities contribute most to yield curve movements and how their influence changes over time.
- **Avoidance of Data Alignment Issues:** By relying solely on financial market data, we sidestep the challenges of aligning economic data—which is often published less frequently and can introduce look-ahead bias—with financial market data.
- **Foundation for Regime Identification:** The variance of the first principal component serves as an effective input for a Hidden Markov Model (HMM) to identify high and low volatility regimes inherent in the yield curve.

## Section 2: Hidden Markov Models

In the previous section, we introduced PCA and demonstrated how it can be used to decompose the yield curve into principal components. Our primary interest lies in the variance of the first principal component, which serves as a measure to identify different regimes inherent in yield curve movements. To effectively capture these regimes at each point in time, we need a methodology that can infer the underlying states driving the observed market dynamics. In this section, we propose the use of Hidden Markov Models (HMMs) to perform regime identification.

### *Understanding Hidden Markov Models*

An HMM is a statistical model that describes a system governed by unobservable (hidden) states, where each state generates observable data according to a specific probability distribution. Intuitively, we can think of the process generating each day's returns as being in one of multiple hidden states or regimes. For simplicity, we assume two states reflecting different market conditions: a high-volatility regime and a low-volatility regime. Each day, the return is generated according to the state the process is in on that day, with returns having different distributions depending on the current state.

HMMs are well-suited for identifying these unobserved regimes based on observed data, providing a robust framework to model state transitions and the probability of observing certain data given a state. In our context, the hidden states represent the volatility regimes of the yield curve, and the observed data is the variance of the first principal component derived from PCA.

### *Applying HMM to Yield Curve Variance*

We fit a two-state HMM to the variance of the first principal component of the yield curve. This approach allows us to infer the probability of being in either the high-volatility or low-volatility regime at any point in time based solely on observed market data, thus avoiding the difficulties of aligning economic data—which can introduce look-ahead bias—with more frequent financial market data.

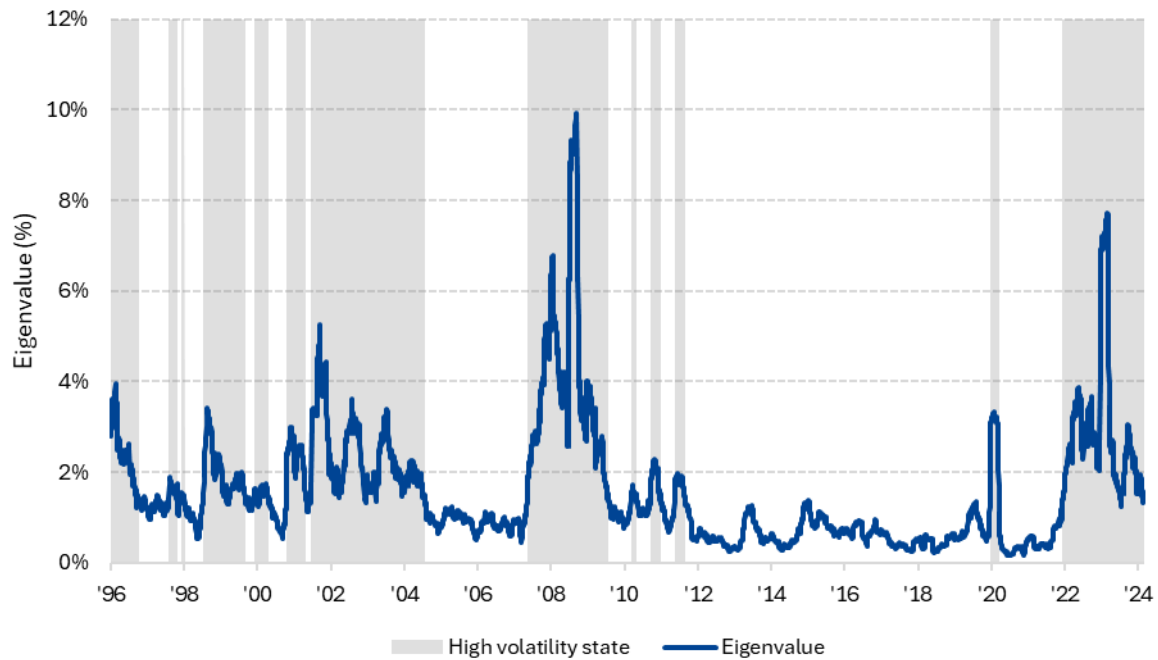
Figure 8 shows the results of fitting a two-state HMM to the variance of the first principal component over the full sample period. The HMM identifies two distinct regimes:

- **High-Volatility Regime:** Characterised by an average variance of 8%.
- **Low-Volatility Regime:** Characterised by an average variance of 2%.

By examining the distribution of daily changes in the first principal component under each regime, we can visualise the differences between them.



**Figure 8: US Treasury curve: Variance of PC1 (eigenvalue) and high volatility states**



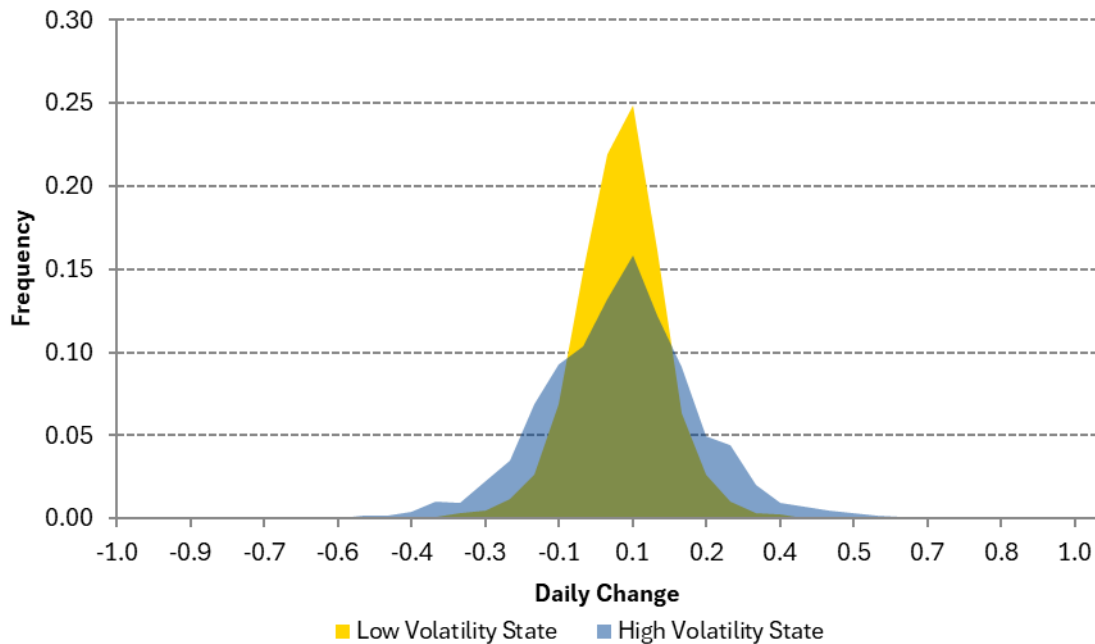
Source: Bloomberg, ESM

Figure 9 illustrates that the high-volatility regime exhibits fatter tails compared to the low-volatility regime, indicating a higher probability of extreme movements in the yield curve. This has significant implications for portfolio optimisation and construction, as the underlying fixed income indices display different covariance dynamics in each regime.

Understanding and identifying these volatility regimes is crucial for both risk management and asset allocation:

- **Risk Management:** By knowing which volatility regime the market is currently in, portfolio managers can adjust their risk exposure accordingly. In a high-volatility regime, there is a greater chance of large swings in bond prices, which can significantly impact portfolio value. Managers might choose to reduce duration or increase the quality of holdings to mitigate risk.
- **Asset Allocation:** The covariance structures of assets differ between regimes. Portfolio managers can optimise their asset allocation by using the appropriate covariance matrix corresponding to the current regime, leading to more efficient portfolios that better balance risk and return.

**Figure 9: US Treasury curve: Distribution of daily changes in first principal component**



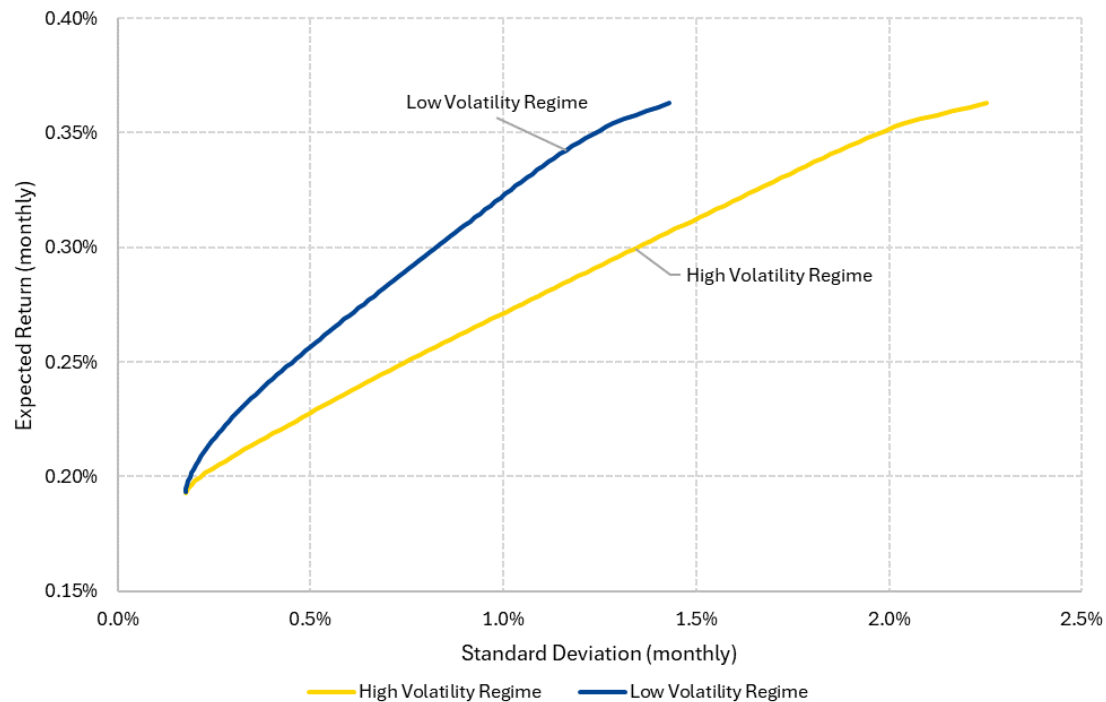
Source: Bloomberg, ESM

#### *Example: Constructing Efficient Frontiers Under Different Regimes*

To illustrate the practical application, we construct efficient frontiers for a portfolio of US Treasury indices under both volatility regimes. We use Intercontinental Exchange (ICE) indices covering US Treasury bills and bonds of various maturities (1-3 year, 3-5 year, 5-7 year, and 7-10 year). We assume that the expected returns for each index remain the same in both regimes but adjust the covariance matrix according to the historical covariances observed in each regime.

The efficient frontiers reveal that, for the same level of expected return, the low-volatility regime offers a much lower standard deviation than the high-volatility regime. This difference is due to the flatter efficient frontier in the high-volatility regime, reflecting the higher variances and covariances among the indices.

**Figure 10: Efficient frontier in high and low volatility states for a portfolio of US Treasury indices**



Source: Bloomberg, ESM

Assumes (simplistically) that returns are the same for each state, but that the covariance matrix follows the historical covariance matrix in each state. Realistically the high state returns could be higher (but more volatile) as US Treasuries tend to rally in high volatility states.

Incorporating regime identification through HMMs offers tangible benefits:

- **Enhanced Portfolio Optimisation:** By adjusting the optimisation process to account for the current volatility regime, portfolios can be constructed to better align with market conditions, potentially improving risk-adjusted returns.
- **Dynamic Risk Assessment:** Identifying the regime allows for dynamic risk assessment and proactive adjustments to the portfolio, rather than relying on static models that may not account for regime shifts.
- **Efficient Use of Market Data:** Using financial market data exclusively ensures timely regime identification without the complications of aligning economic indicators, which can be sporadic and subject to revisions.
- **Mitigation of Look-Ahead Bias:** By avoiding reliance on economic data that may introduce look-ahead bias, the methodology provides a more accurate and reliable basis for decision-making.

By integrating HMMs with PCA-derived variance measures, we provide a robust methodology for identifying market regimes in fixed income portfolios. This approach enables portfolio managers to make informed decisions based on prevailing market conditions, ultimately enhancing portfolio performance and risk management.

## Section 3: Portfolio Optimisation

In this section we apply the PCA and HMM approaches described previously to a portfolio optimisation over a set of fixed income indices. The combination of PCA and HMM is strategic for effective regime identification and portfolio optimisation:

- **PCA Simplifies Complex Data:** PCA reduces the dimensionality of the data by transforming multiple correlated index returns into uncorrelated principal components. The first principal component captures the most significant variance in the data—often associated with the overall level of interest rates.
- **Variance as a Volatility Measure:** The variance of the first principal component serves as a concise and robust measure of market volatility. This single variable encapsulates the collective movements of multiple indices, providing a clear signal for regime shifts.
- **HMM Identifies Hidden Regimes:** HMM uses the variance extracted from PCA as an observable input to model the hidden states of the market. It probabilistically determines whether the market is in a high-volatility or low-volatility regime based on the observed data.
- **Integration Enhances Responsiveness:** By feeding PCA-derived volatility measures into the HMM, we create a responsive system that adapts to changes in market dynamics. This integrated approach allows us to detect and respond to regime shifts more effectively than using either method alone.

### *Methodology*

We perform this analysis on separate universes of US Treasury indices, German government bond indices, and French government bond indices, before combining each with an SSA fixed income index in the same currency. The set of indices used are given in Annex A. The government bond indices used cover the 0-10 year maturities, while those used for SSA cover the 1-5 year maturities. Our sample period starts in 1991 for the US indices and in 1998 for the German and French indices. The period ends on 31 May 2024. We require at least ten-years' worth of history, meaning that in the US case the regime-based portfolio begins around 2001 and in the German and French portfolios around 2008.

We have no constraints other than the value-at-risk (VaR) target and, therefore, the optimisation would yield both a duration and curve position based on the resulting weights. The returns for the optimisation step are estimated using the respective points of the yield curve at the time, for each index. This is done for both the historic approach and the regime-based portfolio.

While in Section 1 we performed PCA on the yield curve, we find using the indices directly provides moderately better results overall although the interpretation of the first three principal components (level, slope, and curve) remains the same. The PCA is performed on a rolling three-month basis, and

the resulting eigenvalue (the variance of the first principal component) is used in the HMM step of the model.

In our analysis, up until now, we have looked at the full history when fitting the HMM. In this analysis we fit the HMM on an expanding window basis, fitting the model on all the output from the PCA step as at that point in time. The regime-based portfolio requires a large initial dataset to identify the regimes and we, therefore, start with an initial ten-year window and expand this as we step through time. Comparisons between the approaches are done over common periods (that is, from 2010 onwards).

The rebalance is performed monthly and within the return optimisation step, we include a penalty term to ensure that the change in weights is not too volatile from one month to the next<sup>3</sup>. The penalty term used throughout is 10% of the sum of the squared differences in weights from one month to the next. This ensures a smoother transition in weights relative to the change in VaR, which is important for larger portfolios and for managing execution costs. This penalty term is used throughout all optimisations.

Our goal is to construct a portfolio with a VaR target, specifically, a 3% VaR at the 99<sup>th</sup> percentile over a one-year horizon by considering the estimated regime at the time of rebalance. In our regime-based approach, we separate the history up until the rebalance date into two regimes based on the PCA and HMM model. We then use the data from the current regime at rebalance date to estimate a VaR and construct the portfolio. We compare our results with a typical historical VaR with a five-year lookback window approach.

In summary, at each rebalance date (monthly), the regime-based approach involves the following multi-step process:

- Perform a rolling (three-month) PCA on the underlying index returns up to rebalance date,
- Extract the rolling eigenvalue of the first principal component (the variance),
- Fit a two-state HMM using the eigenvalue as a feature,
- Separate the history of index returns and proposed portfolio returns according to the regime classification,
- Estimate VaR (historic) for the current regime as per the estimate from the HMM,
- Optimise the portfolio weights such that annual VaR is 3% at the 99th percentile, and with the 10% penalty term to ensure smooth weight.

#### *Results: Government bond portfolios*

Table 1 shows some of the metrics for the portfolio optimisations across the US, German and French government bond indices. In all three cases, the regime-based approach improves the actual portfolio

---

<sup>3</sup> By including the 10% penalty, we ensure that the turnover generated needs to be compensated by additional return. For example, assuming a 5% turnover in portfolio weights, the additional return would need to exceed  $10\% \times 5 = 5\text{bps}$ .

VaR and reduces the number of VaR exceedances closer to the 1% level we would expect. However, the results are much better in the US case than France and especially Germany, where there are only minor improvements.

**Table 1: Metrics for regime-based and standard five-year historical VaR portfolio optimisation**

	United States		Germany		France	
	Regime	5-yr Historical	Regime	5-yr Historical	Regime	5-yr Historical
Average Duration	1.6258	1.7464	2.1586	2.0849	2.1581	2.0933
Actual Portfolio VaR	-3.04%	-3.56%	-3.35%	-3.67%	3.19%	-3.73%
Comp. Annual Growth Rate	1.89%	1.86%	0.72%	0.65%	0.95%	0.80%
Annualised Volatility	1.57%	1.76%	1.42%	1.43%	1.42%	1.49%
Sharpe Ratio	0.046	0.040	0.028	0.024	0.037	0.028
CAGR/unit duration	1.163	1.066	0.334	0.314	0.441	0.384
CAGR/unit VaR	0.039	0.032	0.013	0.011	0.018	0.013
Skewness	-0.081	0.236	0.189	0.454	0.117	0.426
Kurtosis	3.876	5.666	3.932	7.053	4.328	8.954
VaR Exceedances	1.07%	1.77%	1.58%	1.58%	1.24%	1.81%

# Annual VaR for actual portfolio returns

\* Short-term bill yield used as risk free rate

US start 2002, Germany and France start 2008

In all the portfolios the kurtosis<sup>4</sup> is much lower for the regime-based portfolio, indicating much lower probabilities of extreme events. However, the skewness<sup>5</sup> is also much lower, and in the case of the US is slightly negative. This might be an indication of the regime-based approach only removing extreme positive observations. Despite this, however, the regime-based portfolio improves the Compound Annual Growth Rate (CAGR) and the Sharpe ratios and lowers the VaR of the portfolio's return, indicating that there is a significant reduction of negative extreme observations as well.

In general, however, it does appear this approach is much more effective in the US Treasury market than that of France or Germany.

In Figure 11 we show the duration of the regime-based portfolio relative to the five-year historical VaR portfolio for the US indices. We note again, for this analysis, that the only difference between the two


<sup>4</sup> This measures the "tailedness" of a distribution, or the likelihood of extreme events. A higher kurtosis indicates a higher probability of extreme positive or negative returns, often referred to as "fat tails."

<sup>5</sup> This measures the degree of asymmetry in a distribution. A negative skewness indicates a higher probability of negative deviations from the mean and vice versa.

portfolios is the historic returns used in the VaR calculation, the expected returns are set to the same values for each portfolio.

As expected, during high volatility periods, the regime-based portfolio has a shorter duration than the historical VaR approach, however, pre-2008, during low volatility periods, the regime-based approach has a higher duration than the historical VaR portfolio (for example, 2005-2006). This is likely because there are a number of high volatility periods in short succession which remain within the five-year window leading to lower risk-taking.

**Figure 11: Duration of the US regime-based to standard five-year historical VaR portfolio**




Source: Bloomberg, ESM

After 2008 however, we enter an extended period of low volatility and since the regime-based portfolio uses the full history, it still embeds some of the higher low volatility periods from the beginning of the sample. Therefore, despite increasing duration, it becomes shorter than the historical VaR portfolio by as much as 60 points in duration in 2016.



**Figure 12: Duration of the Germany regime-based to standard five-year historical VaR portfolio**



Source: Bloomberg, ESM

Figure 12 shows the same active duration for the German indices. Here we have a smaller period to compare, although the results are similarly to the US case. After 2011 and 2012, the regime-based portfolio maintains a higher duration as it moves into the low volatility regime before the standard five-year historical VaR portfolio catches in duration terms in 2017. After this, the five-year historical VaR portfolio is generally longer duration and specifically into the 2019 and 2022 higher volatility periods where the regime-based portfolio adjusts its duration lower.

#### *Expanding to include SSA*

In this section, we expand our analysis by including the SSA indices in Annex A into the US, Germany and France government portfolios. The approach is the same as in the previous section but including the relevant currency's SSA index (USD or EUR).

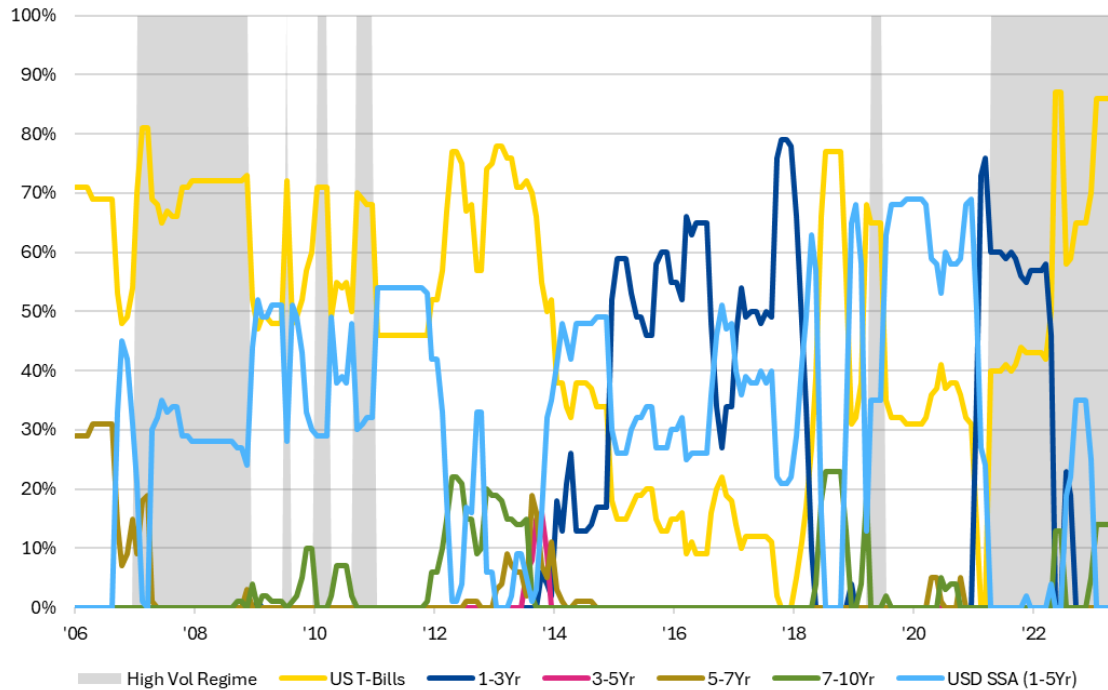
**Table 2: Metrics for regime-based and standard five-year historical VaR portfolio optimisation including SSA**

	United States		Germany		France	
	Regime-based	5-yr Historical	Regime-based	5-yr Historical	Regime-based	5-yr Historical
Average Duration	1.6146	1.8515	2.4823	2.6535	2.3835	2.6013
Actual Portfolio VaR	-2.98%	-3.67%	-3.01%	-3.68%	-3.13%	-3.96%
Comp. Annual Growth Rate	1.82%	1.86%	1.18%	1.04%	1.29%	1.07%
Annualised Volatility	1.36%	1.61%	1.31%	1.48%	1.43%	1.54%
Sharpe Ratio	0.053	0.468	0.055	0.041	0.054	0.039
CAGR/unit duration	1.128	1.005	0.477	0.392	0.542	0.412
CAGR/unit VaR	0.038	0.031	0.024	0.018	0.025	0.017
Skewness	0.024	0.009	0.011	0.091	-0.061	0.042
Kurtosis	6.488	6.799	4.507	8.092	3.554	6.721
VaR Exceedances	0.99%	1.58%	1.04%	1.73%	1.26%	1.98%

Table 2 shows the same metrics as in Table 1 but including the respective 1-5Yr SSA index. Here the regime-based portfolio performs well across all three regions, and arguably better than in the case without SSA indices. Realised portfolio VaR's are close to the target of 3% and VaR exceedances are at the 1% level in the case of the US and Germany, with France improving from 1.98% to 1.26%. Furthermore, the regime-based portfolios all show improved CAGRs, lower volatility, and higher Sharpe ratios.

While the duration profile of the regime-based portfolios is similar to the simulation without SSA, the weight allocation between the two simulations is slightly different. Figure 13 shows the monthly weights for the US regime-based portfolio with SSA, while Figure 14 shows the weights without SSA.

**Figure 13: US Regime-based portfolio weights including SSA**

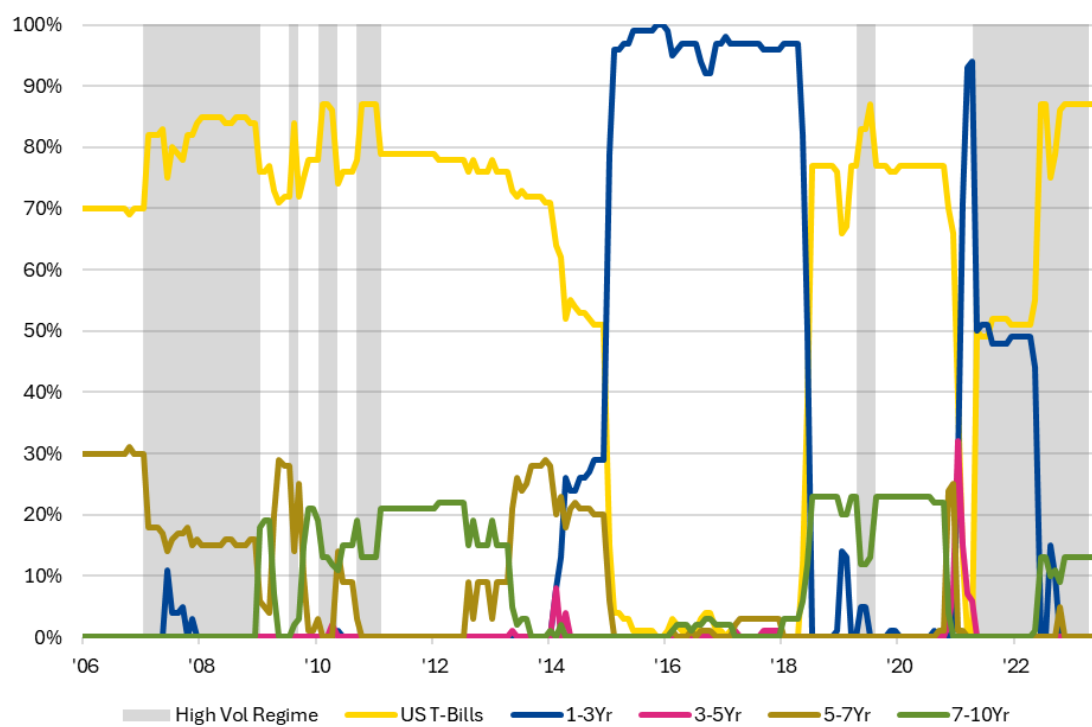


Source: Bloomberg, ESM

In the case without SSA, during low volatility periods there is a mixture of 5-7Yr and 7-10Yr allocation together with a large positioning in US T-Bills. Including SSA, the weights to the 5-7Yr and 7-10Yr indices are essentially zero, except for a short period in 2013/2014 where there is a small allocation to 7-10Yr. Mostly, however, these allocations are replaced by allocations to SSA and during most of low volatility periods the regime-based portfolio with SSA is a mixture of T-Bills and SSAs.

One exception is the 2015-2018 period, where there is a mixture of 1-3Yr and SSA. This contrasts with the simulation without SSA where there is a 100% allocation to 1-3Yr during the same period.

**Figure 14: US Regime-based portfolio weights without SSA**



Source: Bloomberg, ESM

## Section 4: Conclusion

In this paper, we have presented a methodology that integrates Principal Component Analysis (PCA) and Hidden Markov Models (HMM) to implement a regime-based approach in portfolio optimisation for fixed income portfolios. This was done to separate the historical observations into two regimes, namely, a high volatility and low volatility regime. Our approach leverages the strengths of both PCA and HMM to identify underlying market regimes based solely on financial market data, thereby avoiding the pitfalls of aligning economic data with higher-frequency market data and the risk of look-ahead bias.

We showed in Section 2 that using this approach, we can obtain two distinct distributions of returns with the high volatility regime displaying fatter tails as we would expect. We used these distributions to build a simple efficient frontier. As expected, the variances and covariances of the relevant indices increase in the high volatility regime, causing the efficient frontier in this regime to be flatter than the low volatility regime. This causes the volatility to be higher in the case of the high volatility regime for the same expected return.

We used this approach to backtest a portfolio strategy where we target a level of 3% of VaR at the 99<sup>th</sup> percentile. In the regime-based portfolio, we use the historical observations from the regime that the model estimates we are in at each rebalance date and compare this to a standard 5-year historical VaR.

Using a universe of government bond indices, we find that this approach reduces the realised portfolio VaR closer to the 3% level and reduces the VaR exceedances closer to the expected 1% level. Furthermore, the portfolio has a higher return and lower volatility profile than the more standard approach. The results, however, are better in the US case than that of Germany and France where the improvement is less significant.

We extend this by adding in a 1-5Yr SSA index to the universe. The results in this case are better in all cases, including Germany and France with higher returns and better risk statistics in general.

Overall, we find this methodology sufficiently promising to apply it in a regime-aware approach in fixed income portfolio optimisation. Integrating PCA and HMM into fixed income portfolio optimisation provides a robust framework for incorporating regime-based strategies. By capturing the essential movements of the yield curve and identifying hidden volatility states, portfolio managers can make more informed decisions that are aligned with current market dynamics.

Some areas where we think further research could add value include:

- In some cases, we find that high volatility leads to higher returns for government bonds (2008 for example), whereas other high volatility periods are detrimental to government bond returns

(2022 for example). It is possible that including other assets would allow for fitting HMM's with more than two regimes that can better distinguish between these two.

- Pre-empting the regime, using a machine learning approach could add further value by allowing the portfolio to adjust sooner to future regimes.

# References

- Altun, Y. a. (2003). Hidden Markov support vector machines. *Proceedings of the 20th international conference on machine learning (ICML-03)*, 3-10.
- Bailey, S. (2012). Principal component analysis with noisy and/or missing data. *Publications of the Astronomical Society of the Pacific*, 1015.
- Bilokon, P. a. (2021). Iterated and exponentially weighted moving principal component analysis. *arXiv preprint arXiv:2108.13072*.
- Crépey, S. a. (2022). Anomaly detection in financial time series by principal component analysis and neural networks. *Algorithms*, 385.
- Fiszeder, P. a. (2021). Covariance matrix forecasting using support vector regression. *Applied intelligence*, 7029-7042.
- Hirsa, A. a. (2023). Robust Rolling PCA: Managing Time Series and Multiple Dimensions. *SSRN*.
- Li, Y. (2019). Improve Orthogonal GARCH with Hidden Markov Model. *arXiv preprint arXiv:1909.10108*.
- Litterman, R. (1991). Common factors affecting bond returns. *Journal of fixed income*, 54-61.
- Melnikov, O. a. (2016). Dynamic principal component analysis: identifying the relationship between multiple air pollutants. *arXiv preprint arXiv:1608.03022*.
- Meucci, A. (2010). Fully flexible views: Theory and practice. *arXiv preprint arXiv:1012.2848*.
- Ogita, T. a. (2016). Iterative refinement for symmetric eigenvalue decomposition adaptively using higher-precision arithmetic.
- Ogita, T. a. (2019). Iterative refinement for symmetric eigenvalue decomposition II: clustered eigenvalues. *Japan Journal of Industrial and Applied Mathematics*, 435-459.

# Annex A

List of indices used as underlying securities for portfolio optimisation

Ticker	Security Name	Asset Class	Currency
G0FB	ICE BofA French Treasury Bill Index	France	EUR
G1F0	ICE BofA 1-3 Year France Government Index	France	EUR
G2F0	ICE BofA 3-5 Year France Government Index	France	EUR
G3F0	ICE BofA 5-7 Year France Government Index	France	EUR
G4F0	ICE BofA 7-10 Year France Government Index	France	EUR
G0DB	ICE BofA German Treasury Bill Index	Germany	EUR
G1D0	ICE BofA 1-3 Year German Government Index	Germany	EUR
G2D0	ICE BofA 3-5 Year German Government Index	Germany	EUR
G3D0	ICE BofA 5-7 Year German Government Index	Germany	EUR
G4D0	ICE BofA 7-10 Year German Government Index	Germany	EUR
ES81	ICE BofA 1-3 Year excluding Aa3/AA- & Lower Euro Supranationals & Foreign Sovereigns	EUR SSA	EUR
ES82	ICE BofA 3-5 Year excluding Aa3/AA- & Lower Euro Supranationals & Foreign Sovereigns	EUR SSA	EUR
ES86	ICE BofA 5-10 Year excluding Aa3/AA- & Lower Euro Supranationals & Foreign Sovereigns	EUR SSA	EUR
G0BA	ICE BofA US Treasury Bill Index	US	USD
G1O2	ICE BofA 1-3 Year US Treasury Index	US	USD
G2O2	ICE BofA 3-5 Year US Treasury Index	US	USD
G3O2	ICE BofA 5-7 Year US Treasury Index	US	USD
G4O2	ICE BofA 7-10 Year US Treasury Index	US	USD
DS2V	ICE BofA 1-5 Year AAA-AA Developed Markets US Foreign Government & Supranational Index	USD SSA	USD



European Stability Mechanism



6a Circuit de la Foire Internationale  
L-1347 Luxembourg

Tel: +352 260 292 0

[www.esm.europa.eu](http://www.esm.europa.eu)

[info@esm.europa.eu](mailto:info@esm.europa.eu)

