Regime-based portfolio optimisation: a hidden Markov model approach for fixed income portfolios

Byran Taljaard[1]

Abstract

We present a methodology for incorporating a regime-based approach into 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 with those from a traditional approach that does not account for regime changes.

JEL classifications: G11, G12, C32, C38.

[1] European Stability Mechanism (ESM): b.taljaard@esm.europa.eu

1. Introduction

In this paper, we propose a methodology to incorporate a regime-based approach into 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 institutional setting require a slightly different approach, as they are typically based only on fixed income 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. Recognising 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 with those of 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:

1. 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.
2. 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.
3. We use this regime-based approach to perform a target VaR-based portfolio optimisation and compare the resulting portfolio with that of 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.

2. 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 data set. In the context of fixed income portfolios, PCA transforms the data into a set of uncorrelated principal components 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 two-, three-, five-, seven- 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 for fixed income because the first few principal components often have clear and intuitive interpretations (Litterman and Scheinkman (1991)):

1. The **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.
2. The **second principal component (slope factor)** captures changes in the steepness of the yield curve, reflecting differences between short-term and long-term interest rates.
3. The **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.

2.1. Illustrative example

We perform PCA on the US Treasury curve across the six-month and one-, two-, five-, seven- 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.

In Graph 2.1, the first principal component shows loadings that are all in the same direction, confirming it as the level 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 that is long in the belly of the curve and short at the wings.

US Treasury curve: first three principal component loadings by tenor (June 2014–May 2024) Graph 2.1

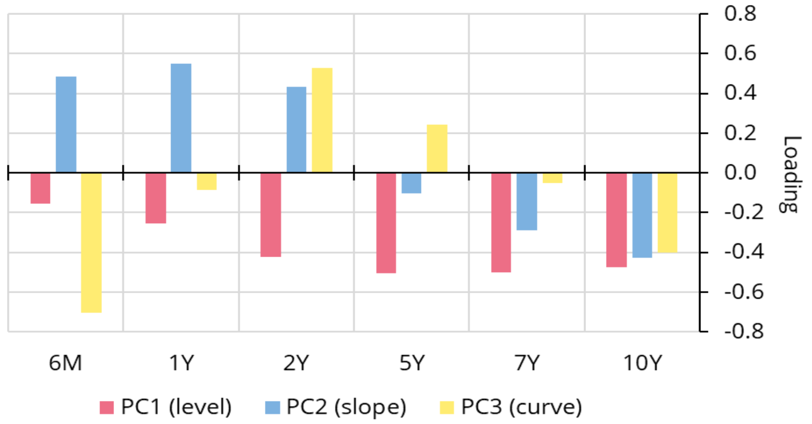

Sources: Bloomberg, ESM calculations.

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.

Proportion of variation explained Graph 2.2

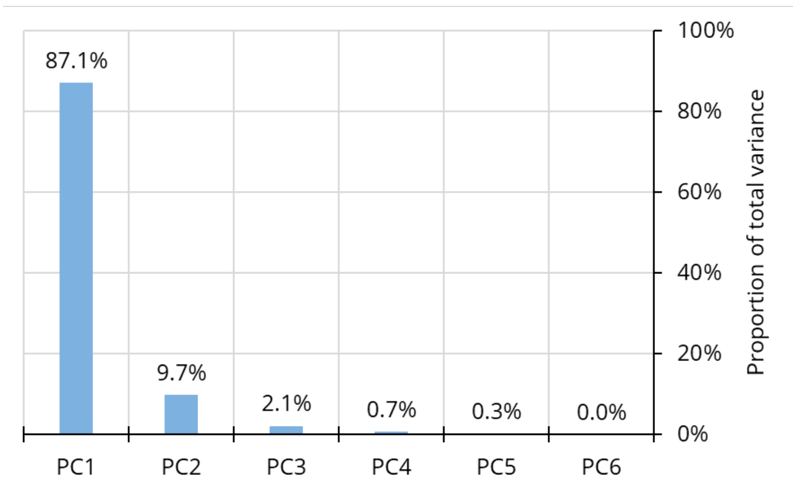

Sources: Bloomberg, ESM calculations.

2.2. Dynamic analysis of loadings

The results cover the entire period from June 2014 to May 2024. However, if we want to get a more granular view of 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 be either all positive or all negative. In other words, while the loadings in Graph 2.1 for PC1 are all negative, they could have also been all positive. This causes an issue when performing the rolling analysis we aim to do here, as the signs of the loadings flip from positive to negative at random.

To correct for this, we explore two options based on existing literature: using the cosine similarity approach as in Hirsa et al (2023) for example; or the approach of Ogita and Aishima (2016), who use the previous eigenvector as a starting point to estimate the next one. We describe these approaches in more detail in Annex B.

In our approach we implement 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 use only the rolling eigenvalues themselves since the actual orientation of the eigenvectors is irrelevant for the size of the eigenvalues. This approach does, however, allow us to analyse the drivers of the yield curve in a more granular manner.

US Treasury curve: loadings for the first principal component on a rolling three-month basis — Graph 2.3

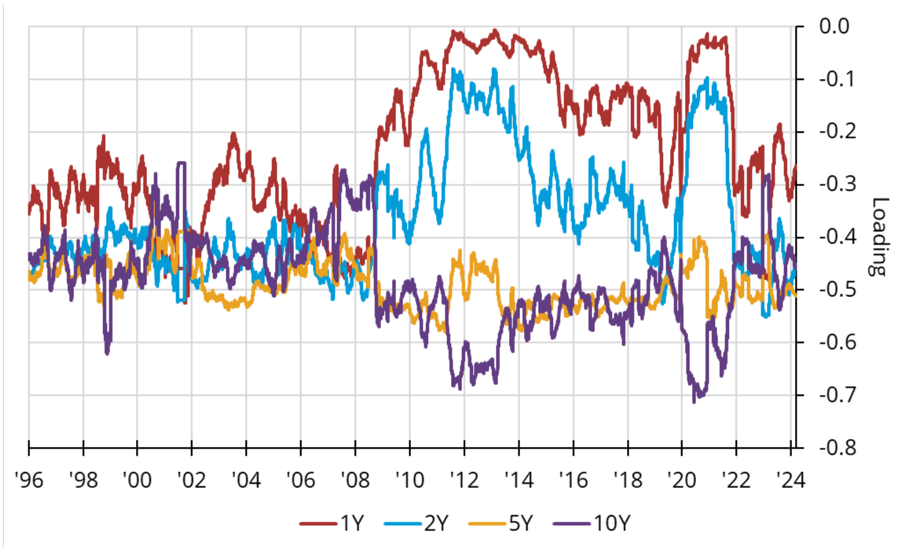

Sources: Bloomberg, ESM calculations.

Graph 2.3 shows the rolling loadings for the first principal component (PC1) for the one-, two-, five- and 10-year tenors of the US Treasury curve. The loadings exhibit interesting behaviour over time, especially when comparing pre- and post-2008. Before 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 the second phase of quantitative easing, when the 10-year point drives most of the movement in the curve, while the front end (one- and two-year points) has little influence. This happens again for a second time in 2020 following the Federal Reserve's response to the Covid-19 crisis. As the Fed hikes rates, for example in the 2016–18 period and in 2022, the front end's influence increases particularly in the two-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.

2.3. Cumulative returns of principal components

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

In Graph 2.4, we show the cumulative returns of the first three principal components. As expected, the first component has the highest volatility and its cumulative return indicates the changes in the yield curve through the cycle of hikes (2004–06, 2016–18 and 2022, for example) and cuts (2000–03, 2007–08 and 2019–20, for example). The second principal component indicates changes in the slope of the curve and, in this case, is a proxy for a flattener position given the loadings.

US Treasury curve: cumulative returns for the first three principal components[2] Graph 2.4

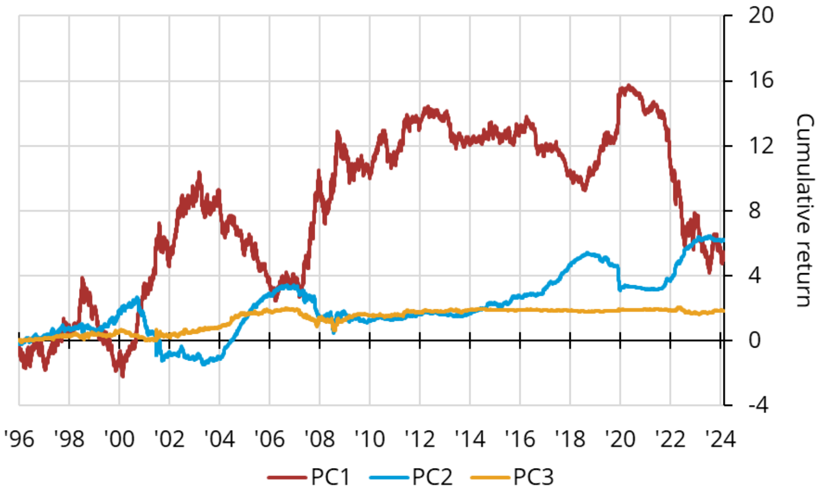

Sources: Bloomberg, ESM calculations.

We perform the same analysis on the German government bond curve (Graph 2.5) where we see a similar profile, although the front end remains a much smaller influence during the 2016–18 period when the ECB kept policy rates on hold, in contrast to the Fed.

2 It is important to compare the returns with the direction of the principal component loadings. In Graph 2.4, PC1 has negative loadings on rates (long bond position), PC2 has a positive loading on the front end and a negative loading on the long end (flattener position), while PC3 has a positive loading on the belly and negative loadings on the wings (benefits when wings outperform the belly).

German yield curve: loadings for the first principal component on a rolling three-month basis — Graph 2.5

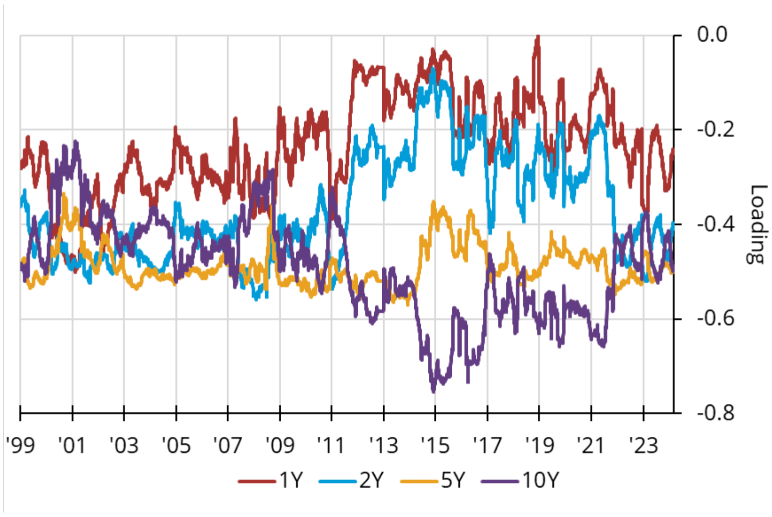

Sources: Bloomberg, ESM calculations.

We show the cumulative returns for the first three principal components in Graph 2.6. 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 United States.

German yield curve: cumulative returns for the first three principal components — Graph 2.6

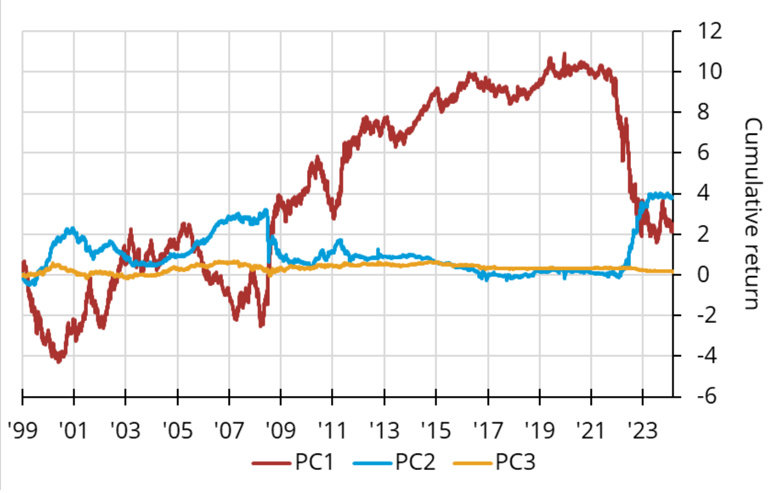

Sources: Bloomberg, ESM calculations.

2.4. 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 Graph 2.7.

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.

German and US yield curve: variance of first principal component (three-month rolling) Graph 2.7

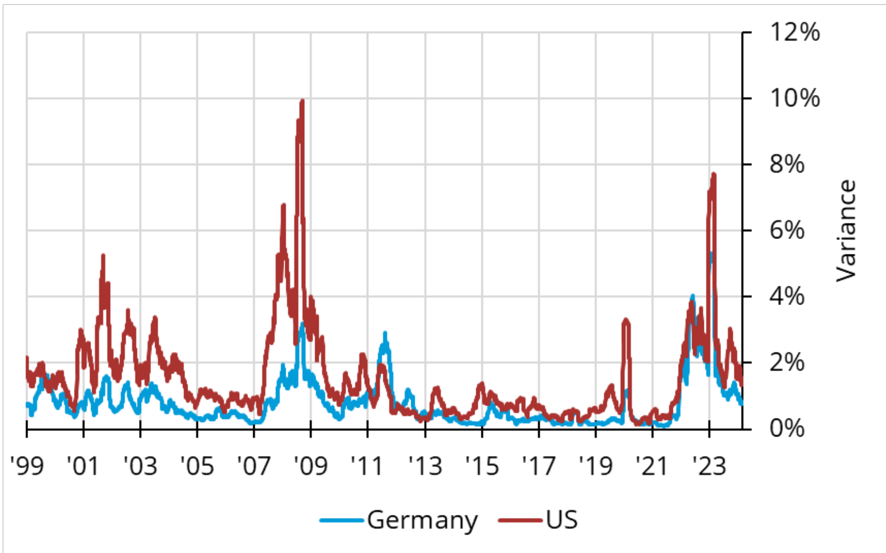

Sources: Bloomberg, ESM calculations.

Using PCA provides several advantages:

1. **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.
2. **Representation of yield curve dynamics:** PCA identifies which maturities contribute most to yield curve movements and how their influence changes over time.
3. **Avoidance of data alignment issues:** By relying solely on financial market data, we sidestep the challenges of aligning economic data – which are often less frequent and can introduce look-ahead bias – with financial market data.
4. **Foundation for regime identification:** The variance of the first principal component serves as an effective input for an HMM to identify high- and low-volatility regimes inherent in the yield curve.

3. 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 identifies 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.

3.1. 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 of 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 of the process 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 are the variance of the first principal component derived from PCA.

3.2. 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.

Graph 3.1 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:

1. A **high-volatility regime** characterised by an average variance of 8%.
2. A **low-volatility regime** charactesized 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.

Graph 3.2 illustrates that the high-volatility regime exhibits fatter tails than 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.

US Treasury curve: variance of PC1 (eigenvalue) and high-volatility states — Graph 3.1

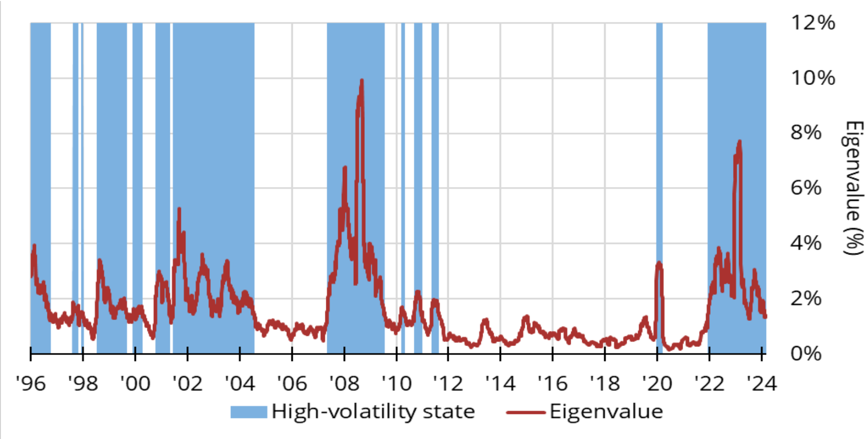

Sources: Bloomberg, ESM calculations.

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 affect portfolio value. Managers might choose to reduce the 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.

US Treasury curve: distribution of daily changes in first principal component — Graph 3.2

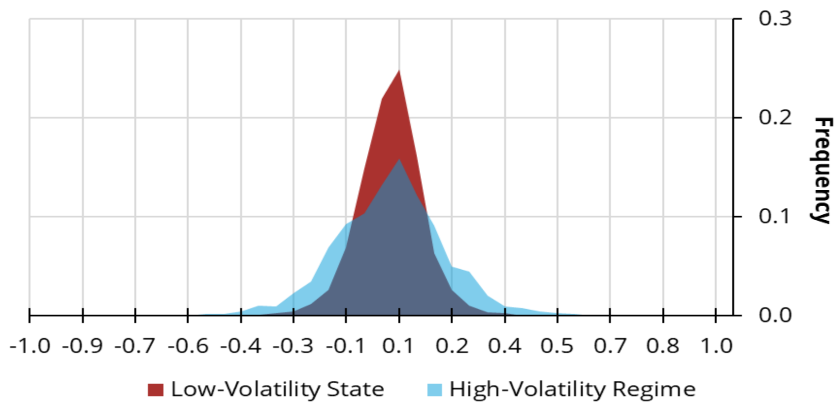

Sources: Bloomberg, ESM calculations.

3.3. 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 (one to three years, three to five years, five to seven years and seven to 10 years). 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 (Graph 3.3) 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.

Efficient frontier in high- and low-volatility states for a portfolio of US Treasury indices — Graph 3.3

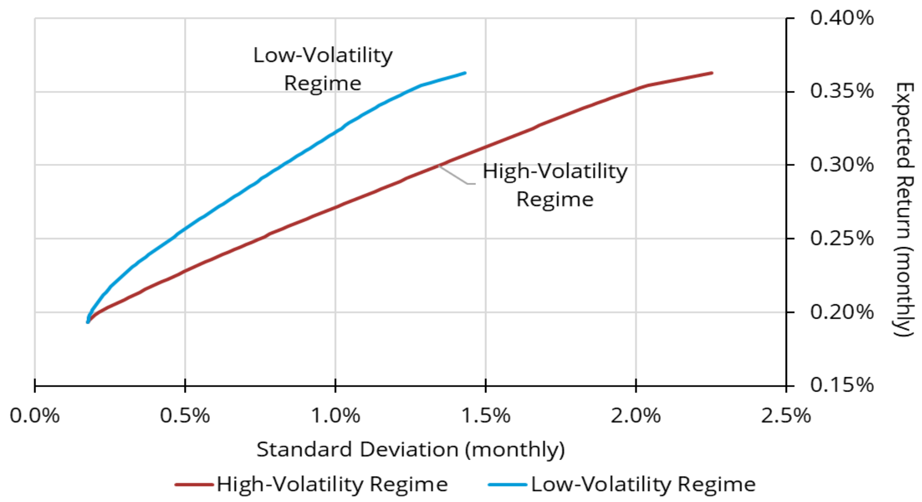

Sources: Bloomberg, ESM calculations.

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**: Recognising 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 infrequent 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.

4. Portfolio optimisation

In this section we apply the PCA and HMM approaches 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.
- **A concise 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.

4.1. Methodology

We perform this analysis separately on 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 provided in Annex A. The government bond indices cover maturities of zero to 10 years, while those used for SSA cover maturities of one to five years. 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 10 years' worth of history, meaning that in the US case the regime-based portfolio begins around 2001 and the German and French portfolios around 2008.

We have no constraints other than the value-at-risk (VaR) target; 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.

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 data set to identify the regimes, and we therefore start with an initial 10-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.[3] 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 99th 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 the 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.

4.2. Results: government bond portfolios

Table 4.1 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 VaR and reduces the number of times VaR is exceeded closer to the 1% level we would expect. However, the results are much better in the US case than in the French and especially the German where improvements are only minor.

[3] 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% x 5 = 5 basis points.

Metrics for regime-based and standard five-year historical VaR portfolio optimisation — Table 4.1

	Regime-based vs five-year historical VaR optimisation					
	United States		Germany		France	
	Regime	5-yr Historical	Regime	5-yr Historical	Regime	5-yr Historical
Average duration	1.63	1.75	2.16	2.08	2.16	2.09
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%

Sources: Bloomberg, ESM calculations.

United States start in 2002; Germany and France start in 2008.

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

In all the portfolios, the kurtosis[4] is much lower for the regime-based portfolio, indicating much lower probabilities of extreme events. However, the skewness[5] is also much lower, and in the case of the United States is slightly negative. This might indicate that the regime-based approach removes only 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 a significant reduction in negative extreme observations as well.

In general, however, this approach appears to be much more effective in the US Treasury market than in the government bond markets of France or Germany.

In Graph 4.1 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 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, 2005–06). This is likely because there are a number of

4 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".

5 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.

high-volatility periods in short succession which remain within the five-year window, leading to lower risk-taking.

Duration of the US regime-based portfolio to standard five-year historical VaR portfolio

Graph 4.1

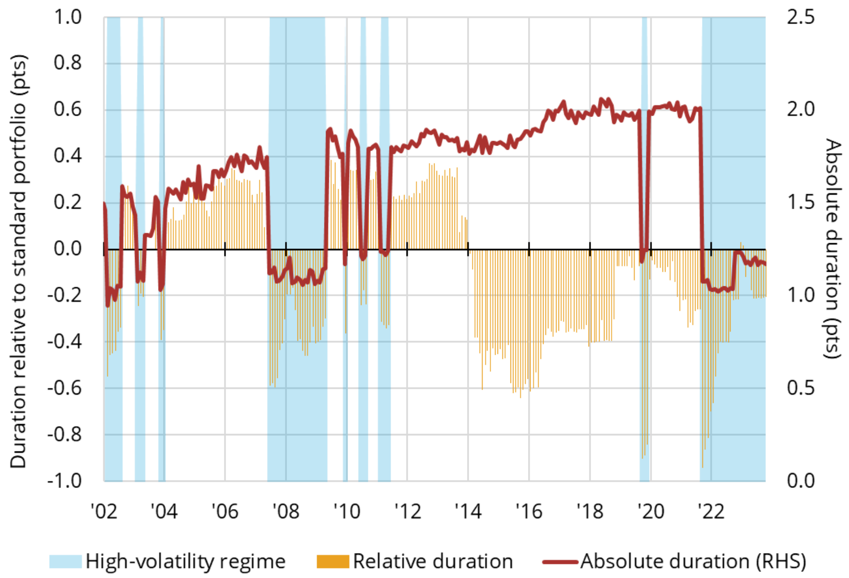

Sources: Bloomberg, ESM calculations.

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-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.

Graph 4.2 shows the same active duration for the German indices. Here we have a smaller period to compare, although the results are similar to the US case. After 2011 and 2012, the regime-based portfolio maintains a higher duration as it moves into the low-volatility regime well before the standard five-year historical VaR portfolio which converges in duration terms in 2017. After this, the five-year historical VaR portfolio is generally of longer duration, specifically into the 2019 and 2022 higher-volatility periods where the regime-based portfolio adjusts its duration lower.

Duration of the German regime-based portfolio to standard five-year historical VaR portfolio | Graph 4.2

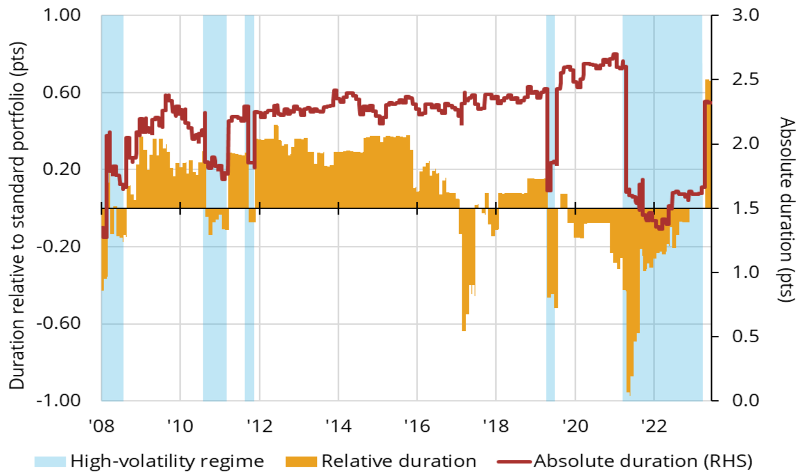

Sources: Bloomberg, ESM calculations.

4.3. Expanding to include SSA

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

Metrics for regime-based and standard five-year historical VaR portfolio optimisation including SSA | Table 4.2

	Regime-based vs five-year historical VaR optimisation					
	United States		Germany		France	
	Regime	5-yr Historical	Regime	5-yr Historical	Regime	5-yr Historical
Average duration	1.61	1.85	2.48	2.65	2.38	2.60
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%

Sources: Bloomberg, ESM calculations.

US, Germany and France start in 2007.

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

Table 4.2 shows the same metrics as in Table 4.1 but includes the respective one- to five-year SSA index. Here the regime-based portfolio performs well across all three regions and arguably better than 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 United States 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. Graph 4.3 shows the monthly weights for the US regime-based portfolio with SSA, while Graph 4.4 shows the weights without SSA.

US regime-based portfolio weights including SSA — Graph 4.3

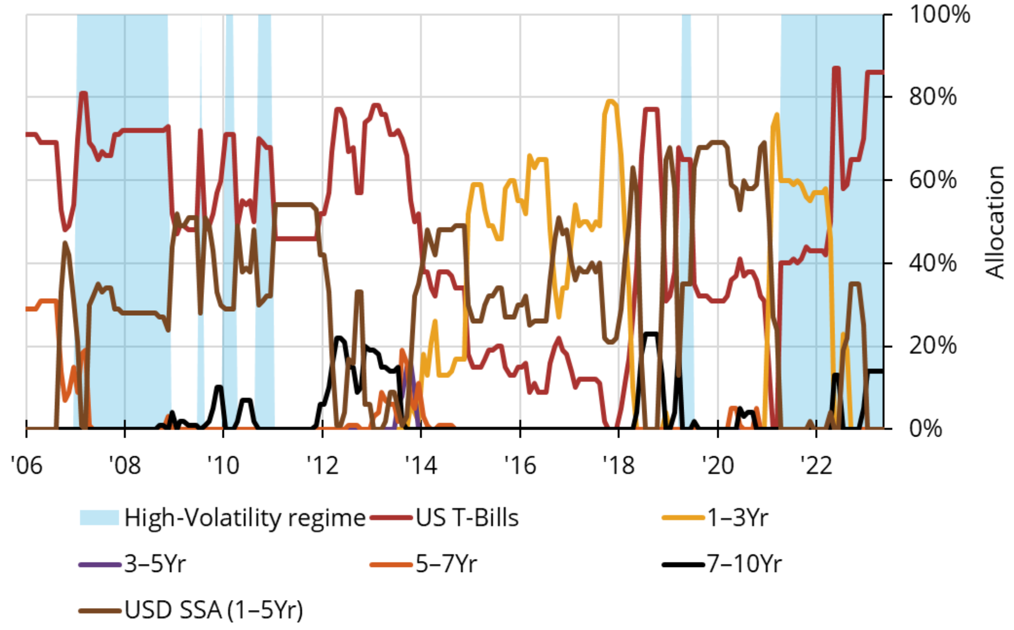

Sources: Bloomberg, ESM calculations.

In the case without SSA, during low-volatility periods, there is a mixture of five- to seven-year and seven- to 10-year allocations together with a large positioning in US Treasury bills. Including SSA, the weights to the five- to seven-year and seven- to 10-year indices are essentially zero, except for a short period in 2013–14 when there is a small allocation to the seven- to 10-year indices. 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 Treasury bills and SSAs.

One exception is the 2015–18 period when there is a mixture of one- to three-year indices and SSA. This contrasts with the simulation without SSA where there is a 100% allocation to one- to three-year indices during the same period.

US regime-based portfolio weights without SSA

Graph 4.4

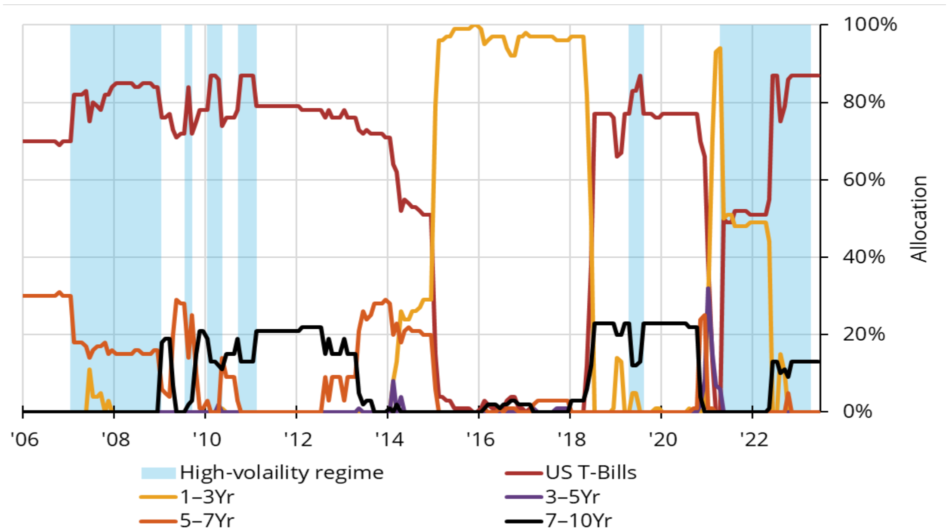

Sources: Bloomberg, ESM calculations.

Conclusion

We have presented a methodology that integrates principal component analysis (PCA) and hidden Markov models (HMMs) 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 HMMs 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 show in Section 2 that by using this approach, we can obtain two distinct distributions of returns with the high-volatility regime displaying fatter tails, as we would expect. We use 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 to be flatter than in the low-volatility regime. This causes the volatility to be higher in the high-volatility regime for the same expected return.

We use this approach to backtest a portfolio strategy where we target a VaR level of 3% at the 99th 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 with a standard five-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 those of Germany and France where the improvement is less significant.

We extend this by adding in a one- to five-year SSA index to the universe. The results are better in all cases, including those of Germany and France, with higher returns and better risk statistics in general.

Overall, we find this methodology promising for including 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:

- 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). Perhaps including other assets would allow for fitting HMMs with more than two regimes that can better distinguish between these two findings.
- Using a machine learning approach could add further value by pre-empting regime changes allowing the portfolio to adjust sooner to future regimes.

References

Altun, Y, I Tshochantaridis and T Hofmann (2003): "Hidden Markov support vector machines", *Proceedings of the twentieth international conference on machine learning (ICML-03)*, Washington DC.

Bailey, S (2012): "Principal component analysis with noisy and/or missing data", *Publications of the Astronomical Society of the Pacific*, vol 124, no 919, pp 1015–23.

Bilokon, P and D Finkelstein (2021): "Iterated and exponentially weighted moving principal component analysis", *arXiv:2108.13072*.

Crépey, S, N Lehdili, N Madhar and M Thomas (2022): "Anomaly detection in financial time series by principal component analysis and neural networks", *Algorithms*, vol 15, no 10, 385.

Fiszeder, P and W Orzeszko (2021): "Covariance matrix forecasting using support vector regression", *Applied Intelligence*, vol 51, pp 7029–42.

Hirsa, A, F Klinkert, S Malhotra and R Holmes (2023): "Robust rolling PCA: managing time series and multiple dimensions", *SSRN*.

Li, Y (2019): "Improve orthogonal GARCH with hidden Markov model", *arXiv:1909.10108*.

Litterman, R and J Scheinkman (1991): "Common factors affecting bond returns", *Journal of Fixed Income*, vol 1, no 1, summer, pp 54–61.

Melnikov, O, L Raun and K Ensor (2016): "Dynamic principal component analysis: identifying the relationship between multiple air pollutants", *arXiv:1608.03022*.

Meucci, A (2010): "Fully flexible views: theory and practice", *arXiv:1012.2848*.

Ogita, T and K Aishima (2016): "Iterative refinement for symmetric eigenvalue decomposition adaptively using higher-precision arithmetic", *Mathematical and Engineering Technical Reports*, no 2016-11, June.

Ogita, T and K Aishima (2019): "Iterative refinement for symmetric eigenvalue decomposition II: clustered eigenvalues", *Japan Journal of Industrial and Applied Mathematics*, vol 36, no 3, pp 435–59.

Annex A: List of indices

List of indices used as underlying securities for portfolio optimisation — Table A.1

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 Sovereign	EUR SSA	EUR
ES82	ICE BofA 3-5 Year excluding Aa3/AA- & Lower Euro Supranationals & Foreign Sovereign	EUR SSA	EUR
ES86	ICE BofA 5-10 Year excluding Aa3/AA- & Lower Euro Supranationals & Foreign Sovereign	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

Annex B: 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 whether 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 with 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 et al (2023).

Using the Ogita-Aishima algorithm (Ogita and Aishima (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 and Finkelstein (2021). An extension of this algorithm is presented by Ogita and Aisima (2019), who attempt to handle clustered eigenvectors.