LINC
LUND UNIVERSITY FINANCE SOCIETY EST 1991

TRADING & QUANTITATIVE RESEARCH REPORT

Portfolio Optimization and Shrinkage

Improving Mean-Variance Optimization Using Shrinkage of the Covariance Matrix

In collaboration with:

AP3

Analysts: Alexander Nilsson, Peder Persson, Samuel Eriksson
Supervisors: Jonas Bergroth, Elena Westerdahl

Introduction & Theory

Introduction

Modern financial markets are characterized by a high degree of complexity, uncertainty and interdependence, where asset prices are influenced not only by fundamentals but also by investors' expectations, beliefs and collective behavior. Traditional portfolio construction frameworks, most notably Mean-Variance Optimization (MVO) rests on the assumption that expected returns and risk can be reliably estimated from historical data.[1] While theoretically appealing, this assumption is often violated in practice, leading to unstable portfolios and extreme allocations that are difficult to justify. These shortcomings highlight a fundamental tension between elegant mathematical models and the noisy, sentiment-driven nature of the real world markets.

Consequently, the motivation for this research is to address the inherent fragility of the standard MVO. In practice, this optimizer often acts as an "estimation-error maximizer", allocating the most capital to assets with the most optimistic data. This fragility defines the scope of this report. Aiming to conduct a comparative analysis of standard MVO against Enhanced Portfolio Optimization (EPO)[2]. The research strives to determine whether a more advanced framework can successfully mitigate estimation errors and produce better portfolios.

Background

Prior to modern portfolio theory (MPT) investment strategies primarily relied on the standalone analysis of individual assets. In 1952 Harry Markowitz revolutionized the approach by demonstrating that the risk of a portfolio is not merely the weighted average of the individual risks of its components but also a function of the correlations between them. This insight lead to creation of the "Efficient Frontier" the combination of assets that offer the highest expected return for a predefined level of risk.

The mathematical engine behind MPT is Mean Variance Optimization. This framework treats the portfolio construction problem as a quadratic optimization task, requiring two primary inputs: a vector of expected return and a covariance matrix representing the volatility and correlations of assets. By solving the weights that minimize portfolio variance subject to a target return, MVO provides a theoretically rigorous method for capital allocation. Under ideal conditions this model produces the mathematically optimal allocation of capital.

However, this application of MVO in real-world scenarios is complicated by the stochastic nature of financial markets. Unlike physical systems governed by immutable laws, financial markets are adaptive complex systems driven by sentiment. Sentiment often causes the asset prices to deviate from their fundamental values, introducing significant "noise" into the historical data used to estimate the returns and covariance of a set of assets. Periods of extreme market sentiment can lead to structural breaks in correlations rendering historical averages poor predictors of future behavior.

This discrepancy between the theoretical stability of MVO and the dynamic reality of markets gives rise to the problem of estimation risk. It is a well-documented phenomenon. To address these limitations, the field has evolved to include robust techniques such as shrinkage estimators, exemplified by Enhanced Portfolio Optimization and Bayesian frameworks like the Black-Litterman model which attempt to stabilize inputs against noise of the market sentiment

Theory

Mean-Variance Optimization

A standard tool to optimize a portfolio was built on the insight that one can combine assets in an intelligent way to improve returns and reduce risk. The model considers assets $i \in \{1, 2, \ldots, N\}$ with returns R_i that are distributed with some expected value $E(R_i) = \mu_i$, variance $V(R_i) = \sigma_{ii}^2$ and covariance $C(R_i, R_j) = \sigma_{ij}^2$.

Let $E(R) := [E(R_1), E(R_2), \ldots, E(R_n)]^T$ be the vector of expected returns and Σ the covariance matrix where $\Sigma_{i,j} = \sigma_{ij}^2$ The goal of optimization is to construct a portfolio with weights $\mathbf{w}$ (w_i is the fraction of capital invested in asset i) to optimally trade off expected return $E(\mathbf{w}) = \mathbf{w}^T E(R)$ and portfolio variance $V(\mathbf{w}) = \mathbf{w}^T \Sigma \mathbf{w}$. There are some different approaches to optimally trade off risk and expected returns.

One approach is to maximize the utility function $\mathbf{w}^T E(R) - \frac{\lambda}{2}\mathbf{w}^T \Sigma \mathbf{w}$ for some risk-aversion parameter λ.

$$\mathbf{w} = \underset{\mathbf{w}}{\operatorname{argmax}}(\mathbf{w}^T E(R) - \frac{\lambda}{2}\mathbf{w}^T \Sigma \mathbf{w}) \tag{1}$$

To assign a specific value to the risk parameter can feel arbitrary and unclear for some investor. Therefore, this method may not be ideal for some investor.

Another approach is to maximize the expected returns such that the portfolio variance $V(\mathbf{w})$ is lower than some pre-determined value. This approach is far more intuitive and has a cap on the level of risk an investor can have. On the other hand the method does not account for the marginal gain of returns increased volatility beyond the cap which could have been beneficial.

$$\mathbf{w} = \underset{\mathbf{w}:V(\mathbf{w})\leq\sigma^2_{max}}{\text{argmax}} (\mathbf{w}^T E(R)) \quad (2)$$

Similarly, one can minimize the portfolio variance such that the expected returns is higher than some pre-determined value. Minimizing portfolio variance subject to a predetermined return target is an ideal method for institutions with fixed future liabilities. If the required return target is set higher than what the market can realistically deliver, the optimizer will force the creation of a highly concentrated portfolio, often relying on unrealistic leverage to satisfy the mathematical constraint.

$$\mathbf{w} = \underset{\mathbf{w}:\mu_{min}\leq E(R)\leq}{\text{argmin}} (\mathbf{w}^T \Sigma \mathbf{w}) \quad (3)$$

The investor can also simply maximize the Sharpe ratio according to

$$\mathbf{w} = \underset{\mathbf{w}}{\text{argmax}}\left(\frac{E(\mathbf{w})}{\sqrt{V(\mathbf{w})}}\right) \quad (4)$$

This function needs leverage to be effective since distribution of the assets is based on the Sharpe and not actual contribution to the returns. A investment opportunity may have a low volatility and normal return which results in a good Sharpe. However, the return of the portfolio will be low if the investor is not able to use leverage or the investment's liquidity is limited.

All methods find portfolios on the efficient frontier and which method the investor should use depends on what the investor wants to maximize (e.g. Sharpe, returns, risk). AP3 primarily uses method (1) and this is the method used in this report.

Boundary Conditions

In this report, we only consider portfolios with long exposure and no leverage. Thus, the portfolio has to satisfy the boundary conditions $\sum_{i=1}^{N} w_i = 1$ and $0 \leq w_i \leq 1$.

Problems with Mean-Variance Optimization

The main problem with MVO is that it gives rise to unnatural portfolios due to noise in the estimation of risk and expected returns. The portfolios that cause problems in the optimization are characterized by two things: the variance of the portfolio is underestimated and the noise in the estimation of expected returns is large relative to this variance. This causes the optimization to take large, unintuitive bets in these portfolios as the estimated returns relative to the variance is swamped by the noise from the estimation [2].

The covariance matrix, which plays a crucial role in MVO, is often determined empirically using historical data. However, when the number of assets (N) is large relative to the number of observations (T), the covariance matrix becomes unreliable causing large estimation errors. This results in portfolios that are overly concentrated in a few assets as MVO tends to place large weights on assets with the highest expected returns. This concentration of assets can lead to high out of sample risk.

Simple Enhanced Portfolio Optimization

Pedersen et al. suggests a different way to optimize portfolios, which they call simple EPO [2]. The goal of this method is to increase the estimated variance of the problematic portfolios. This is done so that the optimization doesn't allocate an unnaturally large fraction of capital into those portfolios. They achieve this by simply reducing the covariance elements in the covariance matrix by a factor of θ, while the diagonal elements (the variances) stay the same. In mathematical terms the equation for the shrunken covariance matrix $\widetilde{\Sigma}$ is simply

$$\widetilde{\Sigma} = (1-\theta)\cdot\Sigma + \theta\cdot V \quad (5)$$

where Σ is the estimated covariance matrix and $V = \text{diag}(V(R_1), V(R_2), \ldots, V(R_N))$ is the diagonal matrix with the variances. The final part of the method is simply to run the standard MVO but with the shrunken covariance matrix $\widetilde{\Sigma}$.

Sample Covariance Matrix

For any two assets, i and j, the sample covariariance estimate using monthly returns is calculated according to eqn. (6)

$$\sigma^2_{i,j} = \frac{12}{L-1}\sum_{k=1}^{L}(R_{i,k}-\bar{R}_i)(R_{j,k}-\bar{R}_j) \quad (6)$$

In this equation, $R_{i,k}$ represents the monthly return of asset i in month k, and $\bar{R}_i$ is the average monthly return over the L-month period period. The factor of 12 was applied to annualize the estimate.

A pension fund like AP3 is interested in finding long trends and stable values of the correlations. This is the reason for using monthly returns instead of daily as estimates on daily returns contains more noise.

The main advantage of using the sample covariance is that it is unbiased. However, it is usually not well-conditioned and suited for portfolio estimation, especially when the number of observations are less than the number of assets [4].

Ledoit Wolf

Olivier Ledoit and Michael Wolf developed an estimator for the covariance matrix that is more well-conditioned than the sample covariance matrix. The proposed estimator $\Sigma^* = \rho_1 S + \rho_2 I$ is a linear combination of the sample covariance matrix S and the identity matrix I. Choosing ρ_1 and ρ_2 optimally results in an estimator with less extreme eigenvalues (i.e. eigenvalues close to 0).

Spectral filtering

Another way to stabilize the covariance matrix is spectral filtering using the Marcenko-Pastur theorem from random matrix theory. The procedure is summarized as follows [5]:

1. The covariance matrix is decomposed to the correlation matrix R where V is the diagonal matrix with asset variances.

$$\Sigma = \sqrt{V} R \sqrt{V} \tag{7}$$

2. Marcenko-Pastur Threshold: Given T observations and N assets, the theoretical maximum eigenvalue for a random matrix is defined as:

$$\lambda_{max} = \left(1 + \sqrt{\frac{N}{T}}\right)^2$$

3. Eigenvalue Filtering: Eigenvalues $\lambda_i < \lambda_{max}$ for the correlation matrix R are identified as noise. These are replaced by their mean value $\bar{\lambda}_{noise}$ to preserve the trace of the matrix:

$$\tilde{\lambda}_i = \begin{cases} \lambda_i & \text{if } \lambda_i > \lambda_{max} \\ \text{avg}(\{\lambda_j\}_{j \in \text{noise}}) & \text{if } \lambda_i \leq \lambda_{max} \end{cases}$$

4. Reconstruction: The cleaned correlation matrix $\tilde{R}$ is reconstructed, its diagonal is reset to 1, and it is rescaled to the final cleaned covariance matrix: $\tilde{\Sigma} = \sqrt{V} \tilde{R} \sqrt{V}$.

Stationary Bootstrap

Politis et al. proposes a method called The Stationary Bootstrap for resampling a time-series. Instead of moving blocks with fixed length, they randomize the block-size using a geometric distribution. The main advantage of the method is that if the original time-series is stationary then the resampled pseudo-time series will also be stationary [6].

Let $\{r_1, r_2, \ldots, r_N\}$ be the original time series. Define a block $B_{i,b}$ of length b starting at index i as:

$$B_{i,b} = \{r_i, r_{i+1}, \ldots, r_{i+b-1}\}$$

For any index $j > N$, define $r_j = r_i$ where $i = j \pmod N$ (with $r_0 = r_N$), allowing blocks to wrap around the end of the series.

Let $p \in [0, 1]$ be a fixed parameter. Let I_k be a sequence of i.i.d. random variables drawn uniformly from $\{1, \ldots, N\}$, and let L_k be a sequence of i.i.d. variables drawn from a Geometric distribution:

$$L_k \sim \text{Geometric}(p), \quad P(L_k = m) = (1-p)^{m-1} p$$

The pseudo-series is created by concatenating the blocks $B_{I_1,L_1}, B_{I_2,L_2}, \ldots$ until the total length reaches N.

Since the expected length of the blocks is $E(L_k) = \frac{1}{p}$, p should be chosen to some reasonable value in order to preserve for example volatility clustering within the data.

Data & Method

Data

The Dataset used in the project was generated by AP3 and contains daily returns from October 2005 to October 2025 for various asset classes. Some of the assets are denominated in a currency other than SEK and parts of the cost of currency hedging has been incorporated in the daily returns. Although the cost of currency hedging has not been fully incorporated in the dataset, the portfolio optimization in this project treats the assets as though they are denominated in the same currency.

Equities: Sweden Large Cap, Sweden Small Cap, North America Hedged, Europe Hedged, Japan Hedged, Pacific ex Japan Hedged, EM Hedged, EM ex China Hedged.

Bonds: Sweden Rates, US Rates, Europe Rates, UK Rates, Canada Rates, Japan Rates, US Rates 10y+.

Inflation-Linked Bonds (Real Rates): Sweden Real Rates, US Real Rates, US Real Rates 10y+, Canada Real Rates, Europe Real Rates, UK Real Rates.

Credit: Sweden Mortgage Credits, Europe HY Credits, Europe IG Credits, US HY Credits, UK Credits, AUD Credits.

Alternatives & Real Assets: Sweden Real Estate, Private Equity, Infrastructure, Timberland (Forest), Commodities.

In addition, a dataset over the US 10 Year Treasury Yield from October 2005 to October 2025 was imported from Yahoo Finance as an approximation to the risk-free rate. In the few instances when the dataset was empty, the previous value was used. As the covariances were estimated based on monthly returns, this was not considered a problem.

Another dataset from AP3 was used to analyze how well the models performed in different macroeconomic regimes. Each month from October 2005 to October 2025 was associated with one of the following regimes in the dataset: boom, expansion, recovery, slowdown, contraction or recession.

Methodology

Initially the data was analyzed to get a better understanding on how and why different strategies performed in certain ways.

Firstly, a look-back period where correlation estimates were stable was found by analyzing correlations at different points in the dataset using different look-back periods.

Then the correlation against a target asset (in this case "Equities North America Hedged") was calculated using a rolling window of different stable look-back periods. Mainly length 7 years was used, representing estimation each month by looking back 7 years, similar to how the optimizers behave.

Simple EPO

This project focused on 11 specific values for the shrinkage parameter θ, ranging from 0 to 1 in increments of 0.1. For each $\theta \in \{0, 0.1, 0.2, \ldots, 1\}$, the portfolio optimization followed these steps:

1. Estimation of the Covariance Matrix

The covariance matrix was updated at the start of every three months using a rolling window with some length of L months. The estimation process followed two main steps. First, the dataset over daily-returns was converted into monthly excess returns using the US 10 Year Treasury Yield as an approximation to the risk-free rate. Then the covariances of excess returns were calculated using the standard sample covariance formula according to eqn. (6). The covariance matrix was also cleaned using the Marcenko-Pastur denoising technique.

Because the process requires a full L-month history, the covariance matrix was only estimated at points with enough previous datapoints. The shrunken covariance matrix $\bar{\Sigma}$ used in the optimization was then calculated according to eqn (5).

2. Estimation of Expected Returns

The expected excess return of each asset was assumed to be directly proportional to its volatility (standard deviation). All assets were assigned a fixed Sharpe ratio of 0.25. Therefore, the expected return for asset i was calculated as:

$$\mu_i = r_f + 0.25 \cdot \sigma_i \tag{8}$$

(where r_f is the risk-free rate at that point in time).

3. Portfolio Optimization and Monthly Recalibration

At the beginning of each three months where the covariance matrix was estimated, an optimal mean-variance portfolio was identified by solving the optimization problem (1) using a SLSQP algorithm from Python's convex optimization libraries. This calculation relied on the expected returns μ and shrunken covariance matrix $\tilde{\Sigma}$ calculated in the previous steps. In addition, the risk aversion parameter $\lambda = 2.92$ was used. This value was based on a benchmark portfolio that AP3 uses and reflect where on the efficient frontier they want to reside.

Once established, the portfolio remained unchanged for the next three months until the next rebalancing.

Since Infrastructure and Timberland are illiquid assets, their weight within the portfolio was set to a fixed percentage of 4% and 2% respectively. These constraints reflects the current portfolio composition of AP3.

Ledoit-Wolf

The portfolio strategy utilizing Ledoit-Wolf followed the same process as the simple EPO. However, the sample covariance was optimized according to the process described by Ledoit et al. [4] and this optimized matrix was used in the portfolio optimization at each recalibration date.

Benchmark Portfolio

The portfolios found using Simple EPO were compared to a benchmark portfolio composed of 55% North American Equities and 45% US Rates. The benchmark was calculated using the same dataset and recalibrated at the beginning of each three months.

Metrics Used to Compare and Analyze Portfolios

To evaluate the performance of the portfolios, several metrics were calculated. These include CAGR, volatility, and Sharpe ratio. Additionally, risk exposure was assessed using maximum drawdown, VaR 95%, and CVaR 95%.

To evaluate the practical robustness of the optimized portfolios, weight vectors are analyzed over different time periods. Specifically:

- Weight turnover between periods
- Degree of concentration or diversification in the portfolio.

A stable MVO solution should have relatively smooth adjustments rather than abrupt reallocation driven by estimation noise. Large swings in weights or extremely concentrated portfolios indicate instability and the well known sensitivity of MVO to noisy inputs.

Results

Analysis of Correlations

The first part of the analysis was to determine a look-back period where the covariance matrix was stable. In Figure 1, the correlation between "Equities North America Hedged" and "Bonds US" was calculated at different points in time using different look-back periods. This was also done for more assets, and in general, the covariance matrix was particularly stable when using a look-back window of 6-10 years.

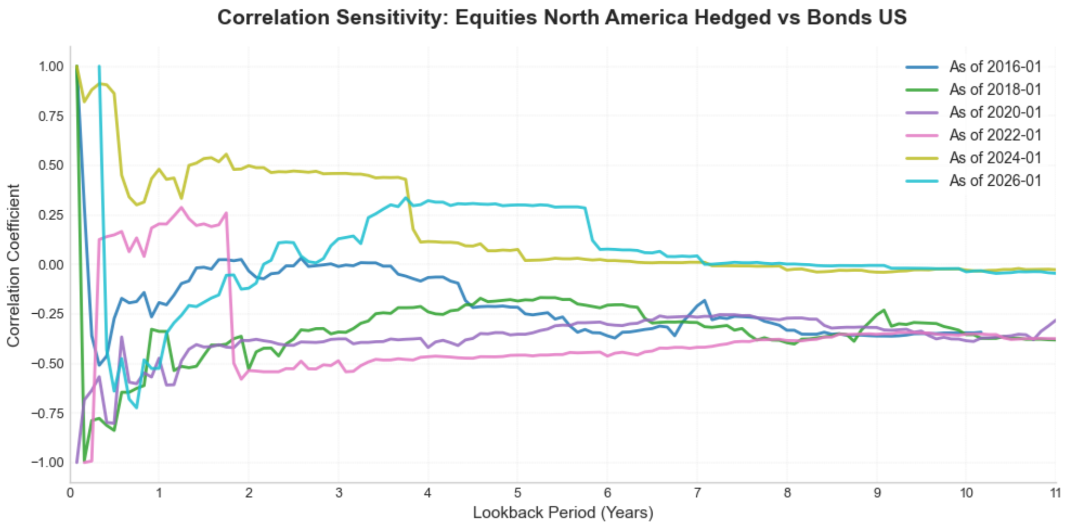

Figure 1: Correlation estimates at different dates between Equities North America and Bonds US using different lookback periods

The correlation against North America Hedged was compared with different look-back lengths and asset classes. Assets within equity markets in Figure 2 are highly correlated over the sample period, but with noticeable time variation and noise. In contrast, Figure 3 indicates that interest rates have a negative correlation with equities but still with a lot of variation and noise. An interesting exception occurs during stress periods, such as the COVID-19 crisis, when cross-equity correlations drop sharply while correlations with interest rates increase. Another important observation is that, in the most recent years, the correlation between bonds and equities has increased, and thereby reducing the diversification benefits traditionally associated with bonds.

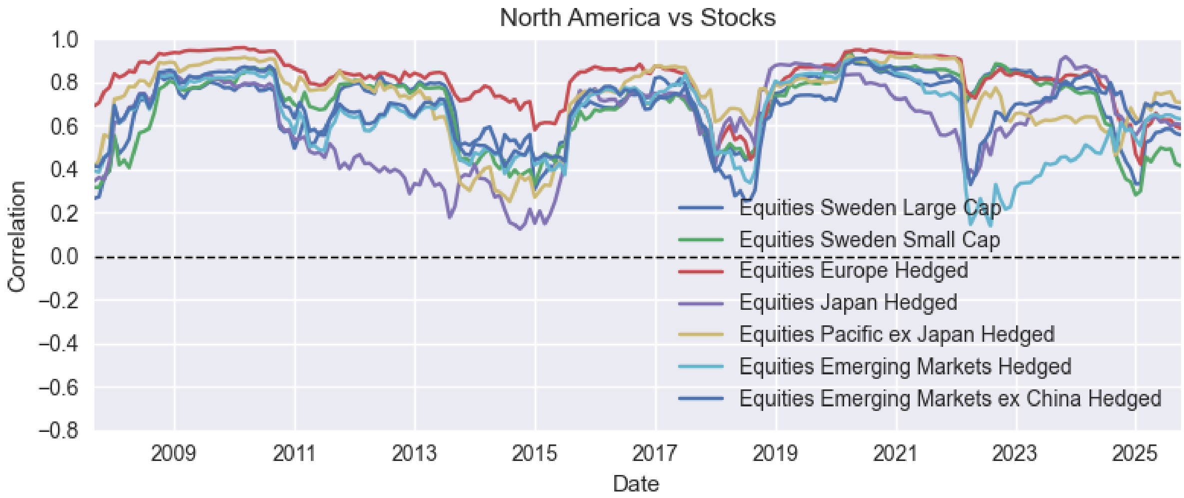

Figure 2: Rolling 24-month correlations between regional equity indices and North American hedged equities. The figure illustrates the high but time varying degree of co-movement across equity markets over the sample period.

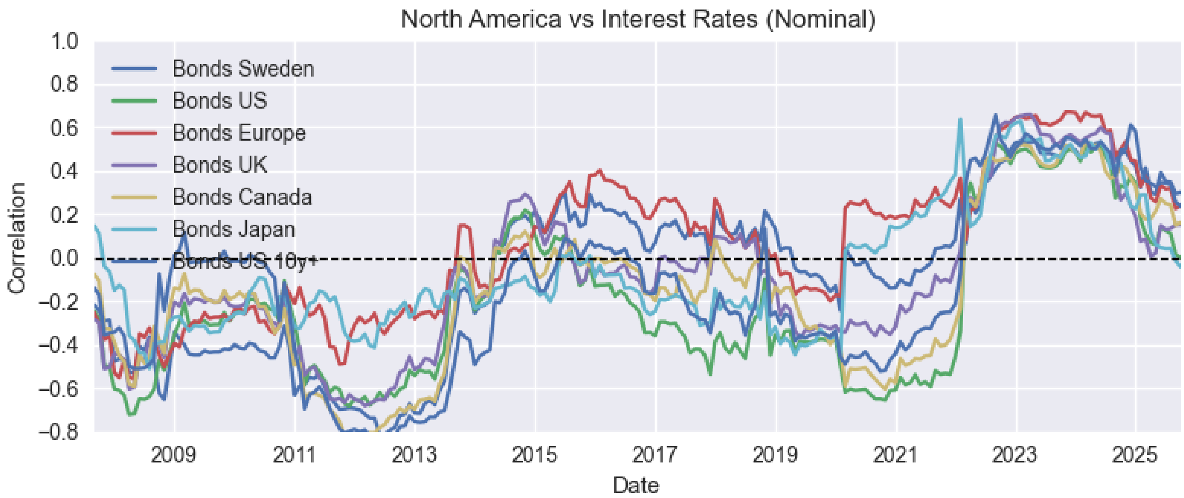

Figure 3: Rolling 24-month correlations between regional government bond yields and North American hedged equities. The figure highlights the negative, but regime dependent, correlation between interest rates and equities, with notable shifts during periods of market stress.

When instead a seven-year look-back window was, the estimated correlations became smoother. As shown in Figure 4, cross-equity correlations become tighter, while still exhibiting drops during stress periods. Similarly, Figure 5 displays the same overall pattern as before, but with clearer consistency. Three distinct regimes can be distinguished where the correlation behaves differently.

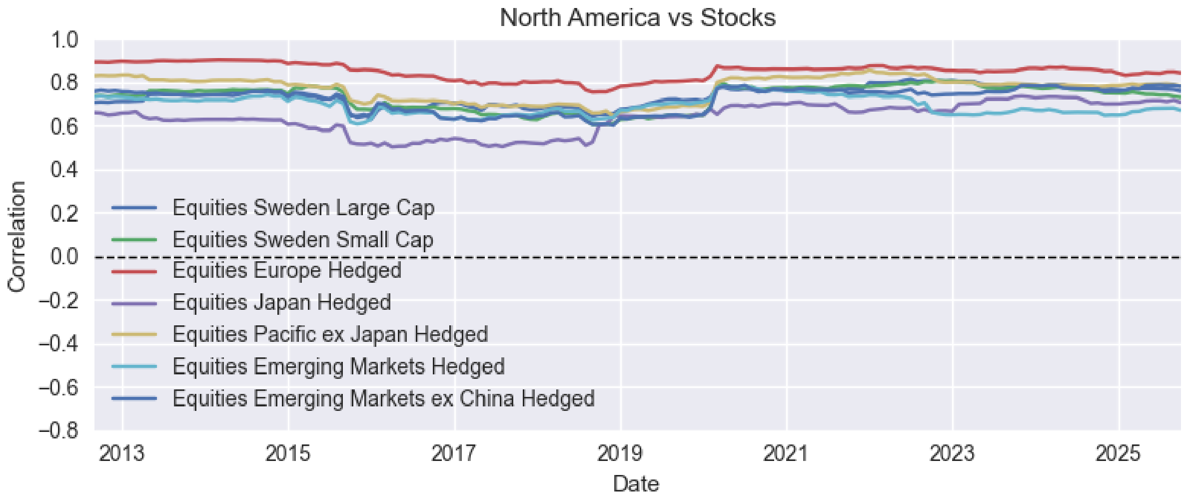

Figure 4: Rolling 7-year correlations between regional equity indices and North American hedged equities.

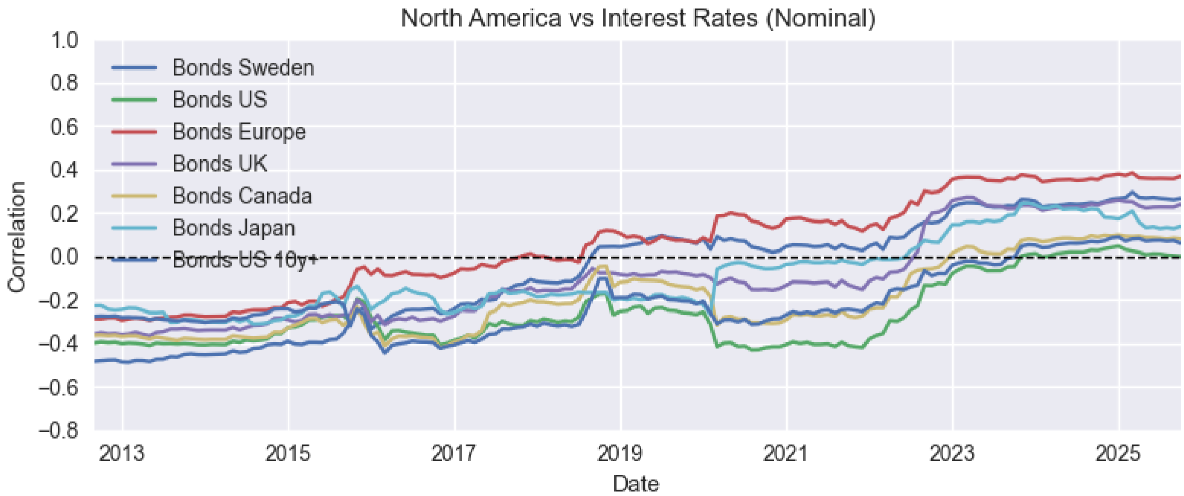

Figure 5: Rolling 7-year correlations between regional government bond yields and North American hedged equities.

Result from Strategies

Next, strategies using different shrinkage parameters, θ, were simulated over the 20 year dataset. Figure 6 compares the performance of these portfolios over the sample period. The portfolio allocation for different assets can be found in Appendix I. The most noteworthy observation is that standard MVO ($\theta = 0$) allocates a substantial part of the portfolio towards Japanese Equities.

The portfolio returns and Sharpe ratios generally improved as θ increased while the portfolio volatility increased. During periods of market downturns the portfolios experienced substantial drawdowns and generally underperformed the benchmark portfolio. Figure 7 illustrates this pattern by showing that the magnitude of the drawdowns increased during stressed market conditions. Portfolios with higher shrinkage parameters (θ close to 1) exhibit deeper and more volatile drawdowns. In contrast, lower values of θ appear to dampen downside risk, resulting in comparatively smaller drawdowns. For $\theta \leq 0.2$, performance increased by using the Marcenko-Pastur denoising technique compared to not using it. For $0.3 \leq \theta$, the increase in performance was insignificant.

The performance metrics for the different macroeconomic regimes are found in Appendix II. In some regimes like Boom and Expansion, the implemented MVO using shrinkage outperforms the benchmark. In other regimes like Recovery and Contraction, the model underperforms. This behavior can likely be explained by the fact that the model allocates more aggressively towards assets with a higher risk-profile when using shrinkage.

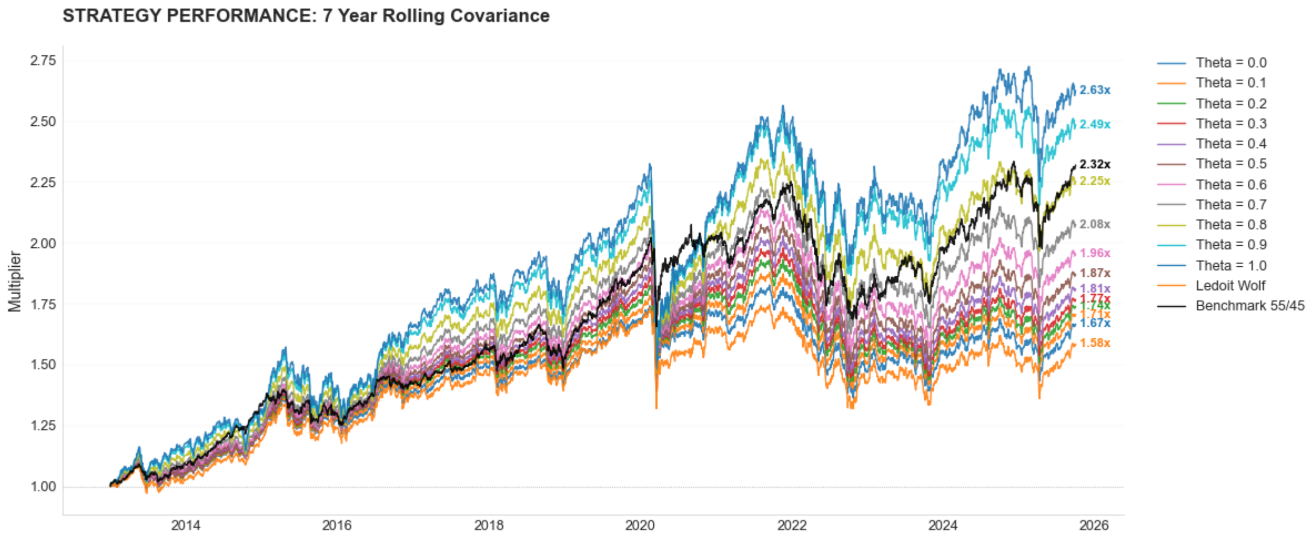

Figure 6: The cumulative performance over the original dataset using different shrinkage parameters

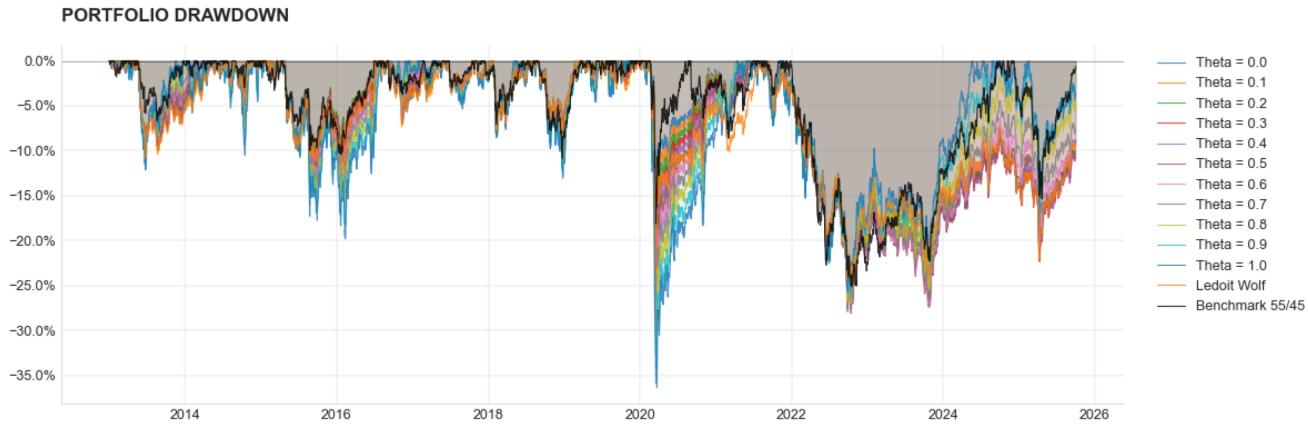

Figure 7: The drawdowns over the original dataset

Table 1 reports the full set of performance and risk metrics used to evaluate the portfolios. Consistent with the previous results, higher values of θ generally yield stronger performance. The specification with $\theta = 1.0$, together with the benchmark portfolio, exhibits the highest Sharpe ratio.

Table 1: Strategy Performance and Risk Metrics using 7 Year Rolling Covariance. $\theta = 0.0$ represents no shrinkage.

Strategy	CAGR (%)	Vol (%)	Sharpe	Max DD (%)	VaR 95% (%)	CVaR 95% (%)
$\theta = 0.0$	3.95	9.02	0.21	-25.43	-0.88	-1.31
$\theta = 0.1$	4.13	9.25	0.22	-25.47	-0.87	-1.32
$\theta = 0.2$	4.28	9.52	0.24	-25.97	-0.85	-1.33
$\theta = 0.3$	4.41	9.82	0.24	-26.99	-0.86	-1.35
$\theta = 0.4$	4.60	10.13	0.26	-27.71	-0.85	-1.38
$\theta = 0.5$	4.86	10.48	0.28	-28.07	-0.88	-1.42
$\theta = 0.6$	5.21	10.88	0.30	-28.37	-0.92	-1.48
$\theta = 0.7$	5.69	11.37	0.34	-29.96	-0.97	-1.57
$\theta = 0.8$	6.35	12.02	0.38	-31.86	-1.03	-1.69
$\theta = 0.9$	7.14	12.87	0.42	-34.24	-1.15	-1.86
$\theta = 1.0$	7.59	13.40	0.44	-36.33	-1.23	-1.97
Benchmark 55/45	6.58	9.46	0.48	-25.14	-0.92	-1.43
Ledoit-Wolf	3.55	9.38	0.16	-24.47	-0.90	-1.38

Result from Stationary Bootstrapping

To further assess the robustness of the results, stationary bootstrapping was employed to generate resampled return series. Table 2 presents the average performance across 1000 bootstrap iterations for different values of θ. The number of iterations were enough for the metrics to converge to some stable value. The benchmark portfolio continues to outperform the optimized strategies in terms of Sharpe ratio. Increasing θ leads to a monotonic improvement in Sharpe performance, though accompanied by higher volatility and deeper drawdowns.

Table 2: Performance Comparison using Stationary Bootstrap and 7 Year Rolling Covariance. $\theta = 0.0$ represents no shrinkage.

Strategy	CAGR (%)	Vol (%)	Sharpe	Max DD (%)	VaR 95% (%)	CVaR 95% (%)
$\theta = 0.0$	3.99	9.66	0.17	-24.88	-0.92	-1.39
$\theta = 0.1$	4.06	9.80	0.18	-25.39	-0.92	-1.40
$\theta = 0.2$	4.18	10.00	0.19	-26.08	-0.92	-1.41
$\theta = 0.3$	4.32	10.27	0.20	-26.97	-0.92	-1.43
$\theta = 0.4$	4.47	10.61	0.22	-28.07	-0.93	-1.47
$\theta = 0.5$	4.62	11.04	0.23	-29.41	-0.96	-1.52
$\theta = 0.6$	4.78	11.60	0.23	-31.06	-1.01	-1.60
$\theta = 0.7$	4.95	12.34	0.24	-33.22	-1.08	-1.72
$\theta = 0.8$	5.11	13.38	0.25	-36.17	-1.19	-1.90
$\theta = 0.9$	5.26	14.75	0.25	-39.98	-1.33	-2.14
$\theta = 1.0$	5.34	15.63	0.25	-42.38	-1.41	-2.30
Benchmark 55/45	5.46	9.06	0.34	-22.69	-0.90	-1.44
Ledoit Wolf	4.11	10.12	0.18	-26.30	-0.94	-1.44

Note: All results represent the mean performance across 1000 bootstrap iterations.

Analysis & Conclusion

Analysis

Initially, the evolving correlation structure illustrated in Figure 4 and 5 will have important implication for MVO. When correlations increase diversification benefits decrease, reducing the curvature of the efficient frontier. As the traditional negative correlation between equities and interest rates weakens, the classical diversification channel underlying balanced portfolios become less effective. This leads to the optimized portfolios, in the high-correlation regime, increasingly resemble concentrated allocations toward the assets with the strongest stand alone Sharpe ratios. Although diversification benefits still exists, the practical ability to identify and exploit them may be obscured by estimation noise and parameter uncertainty.

The optimal value of the shrinkage parameter is not definitive for all investors, as can be deduced from the performance metrics for both the historical simulation in Table 1 and the stationary bootstrap simulation in Table 2. A higher value results in a higher CAGR and a better Sharpe ratio. However, a higher value also increases the maximum drawdowns and results in a worse VaR/CVaR. Thus, if an investor wants to utilize simple EPO for a similar application, their risk appetite becomes the deciding factor for which shrinkage parameter to choose.

The optimizer attempts to maximize the Sharpe ratio by utilizing the covariance between assets. When $\theta = 1$, all cross-asset correlations are removed, and the optimizer can no longer utilize diversification benefits. This is likely one reason why the portfolios with the most shrinkage also experience the worst risk metrics.

This can also be seen by analyzing the portfolio compositions. When the shrinkage θ increases, the allocation toward bonds is reduced, and the allocation toward more volatile assets like equities is increased. One possible reason for this allocation toward more volatile assets relates to the utility function and the assumption regarding future returns. The assumption in this model was that the excess return of an asset is proportional to its volatility. As a consequence, the utility function favors volatile assets as long as there is enough diversification to significantly reduce the second term of the utility function.

Furthermore, the assumption of excess returns in Eq. (11) causes an issue when optimizing the allocation toward credits. Credits generally exhibit low historical volatility, resulting in low estimated expected returns, yet they maintain a positive correlation with equities. As a consequence, credits appear strictly inferior to other asset classes. Forcing a uniform Sharpe ratio across all asset classes systematically disadvantages assets that earn their risk premium from asymmetric risk spreads rather than from standard daily volatility.

Diversification works best when different assets take turns performing well and the correlations between them change over time. During the last 20 years—the period on which the models in this paper are based—the market has exhibited strong momentum. In such environments, concentrated exposures tend to outperform, while diversified portfolios trade potential returns for reduced risk. This may partly explain why the benchmark portfolio performs well relative to the optimized portfolios, as its larger exposure to equities (especially US Equities) allows it to benefit more from the persistent upward trend in equity markets.

It was observed that plain MVO ($\theta = 0$) allocated a substantial part of the portfolio toward Japanese equities. One explanation for this behavior could be that the currency hedging of the asset returns is somewhat misleading. For example, the performance of Japanese equities was based on two indices, only one of which was currency-hedged. This could have caused the data used in the optimization to be misleading.

Another observation was that the sample correlations were particularly stable when using an estimation period of 6–10 years. This suggests that the sample covariance matrix with an estimation period of 7 years used by the optimizer should be relatively robust. For a pension fund like AP3, there is another advantage to using a long look-back window, as the portfolio does not experience too much turnover during turbulent periods. If the portfolio is reallocated too much in a turbulent period, there is a significant risk that it will not be optimal once market conditions return to normal.

When applying the Marchenko-Pastur denoising technique to the sample covariance matrix before optimizing the portfolio, performance improved somewhat. This suggests that there is still some inherent instability within the 7-year sample covariance matrix, even though the correlations appear to be stable. However, when first shrinking the sample covariance matrix and then applying the technique for $0.3 \leq \theta$, performance did not improve as much.

This suggests that a shrunken covariance matrix is more robust than the plain sample matrix. This can also be seen from the portfolio allocations, as higher levels of shrinkage caused less portfolio turnover.

However, this does not necessarily mean that the shrunken covariance matrix is better for portfolio optimization, as the improved returns come with greater risks.

The purpose of the Ledoit-Wolf technique is to create a robust covariance matrix when the number of observations is less than or similar to the number of assets. This is not the case in this model, as 84 observations are used to estimate the covariance between 30 assets. Hence, the Ledoit-Wolf matrix is similar to the sample matrix, as it is already relatively robust, and the performance is similar to using plain MVO.

Conclusion

This report examined whether the simple EPO can mitigate estimation errors and produce better portfolios than standard MVO. The results confirm that plain MVO ($\theta = 0$) generates unintuitive portfolios with a higher degree of turnover. This result might have been impacted by the fact that not all assets were fully currency-hedged, causing the optimizer to overweight Japanese equities.

These corner portfolios are caused by estimation errors in the expected returns and the covariance matrix. To be more precise, a certain combination of assets appears to have far lower risk in-sample than it actually does. This means that the portfolio optimizer allocates a large fraction of the portfolio to those particular assets.

Although the sample covariance matrix was relatively robust due to the long estimation period (7 years), shrinkage still improved robustness, as evidenced by the degree to which spectral filtering further improved the matrix. This could also be seen from the portfolio allocations, as a higher level of shrinkage reduced portfolio turnover.

Even though returns and the Sharpe ratio improved when using shrinkage, risk metrics did not. The conclusion of this report is therefore that the simple EPO model improves the robustness of the covariance matrix, but it is not a sufficient model for this particular application when using a sample covariance matrix based on monthly returns with a long estimation period and using the assumption that expected excess returns are proportional to asset volatility.

While the simple EPO might not be well-suited for this application, it might be better suited for more systematic trading applications. Pedersen et al. found that the model worked well when using a systematic trading strategy with leverage and an exponentially weighted covariance matrix based on daily returns and a 150-day center of mass [2]. The conclusions are therefore strictly related to the application outlined in this report.

References

References

[1] Frank J. Fabozzi, Harry M. Markowitz, Francis Gupta (1952). *Portfolio Selection*. The Journal of Finance Volume 7, Issue 1. https://onlinelibrary.wiley.com/doi/full/10.1111/j.1540-6261.1952.tb01525.x

[2] Pedersen, L. H., Babu, A., & Levine, A. (2021). Enhanced Portfolio Optimization. Financial Analysts Journal, 77(2), 124–151. https://doi.org/10.1080/0015198X.2020.1854543

[3] Black, F.,Litterman, R. *Global Asset Allocation With Equities, Bonds, and Currencies*, Goldman Sachs Fixed Income Research, 1991.

[4] Ledoit, O., Wolf, M. *A well-conditioned estimator for large-dimensional covariance matrices*, Anderson Graduate School of Management, UCLA; Department of Economics and Business, Universitat Pompeu Fabra, 2004.

[5] López de Prado, M. Machine Learning for Asset Managers, Cambridge University Press; Elements in Quantitative Finance, 2020.

[6] Politis, D. N., and Romano, J. P. (1994). The Stationary Bootstrap. *Journal of the American Statistical Association*, Vol. 89, No. 428, pp. 1303–1313. Taylor & Francis on behalf of the American Statistical Association.

Appendix I - Historical Asset Allocations

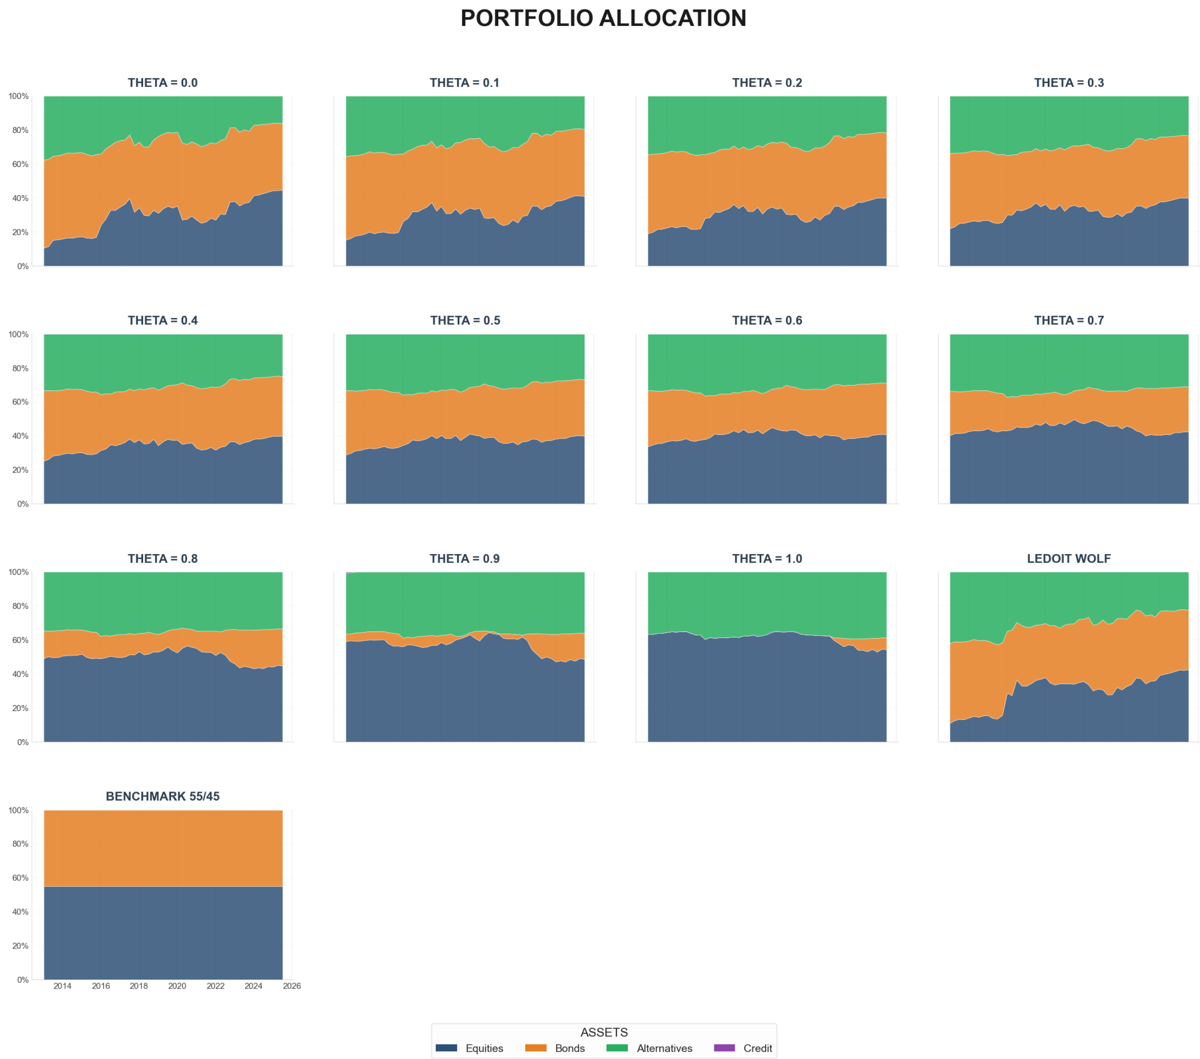

Figure 8: Allocation to Equities, Bonds, Credit and Alternatives

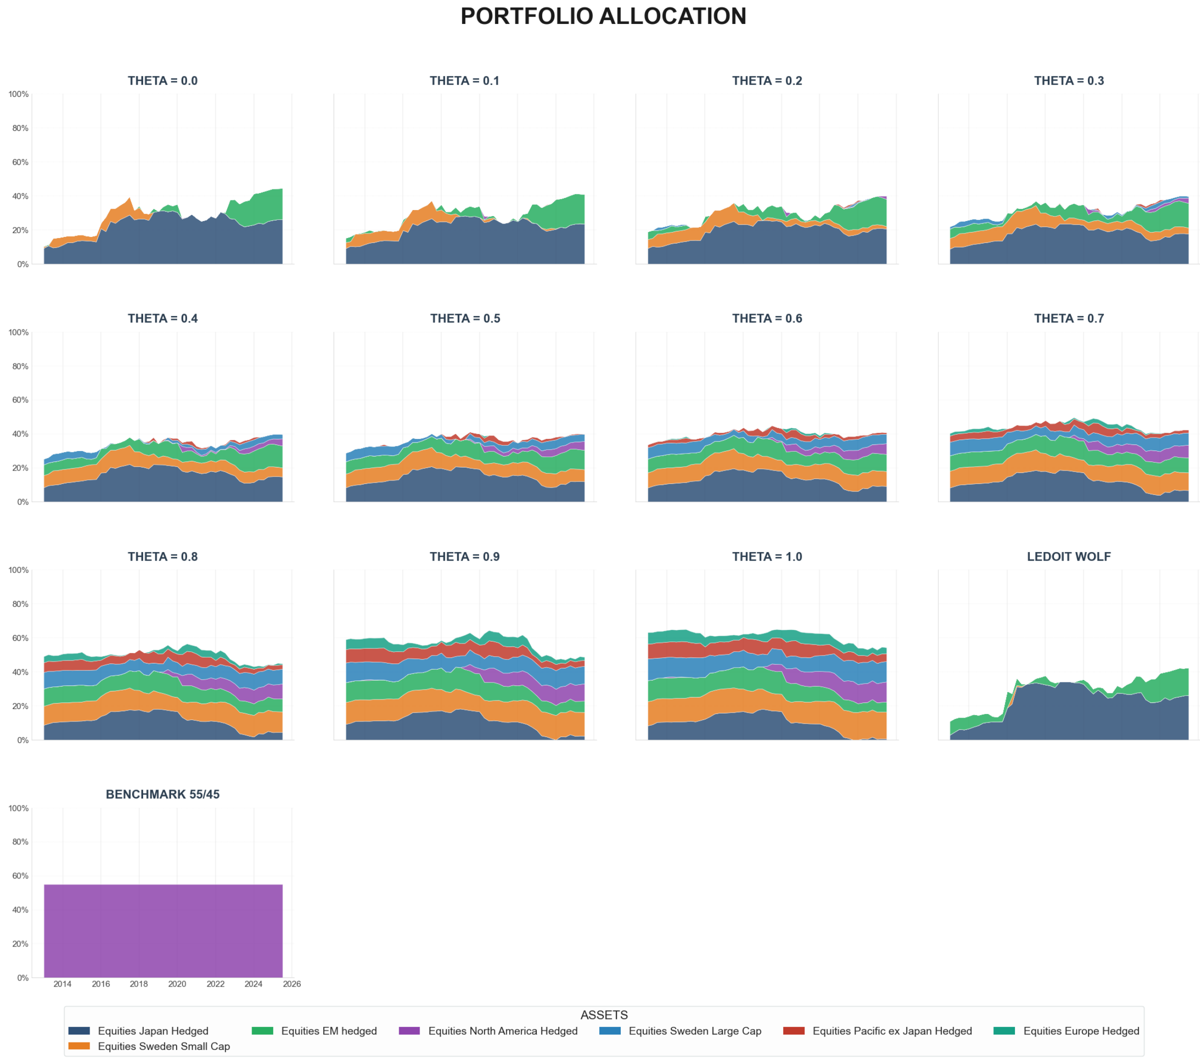

Figure 9: Allocation to Equities for the different strategies

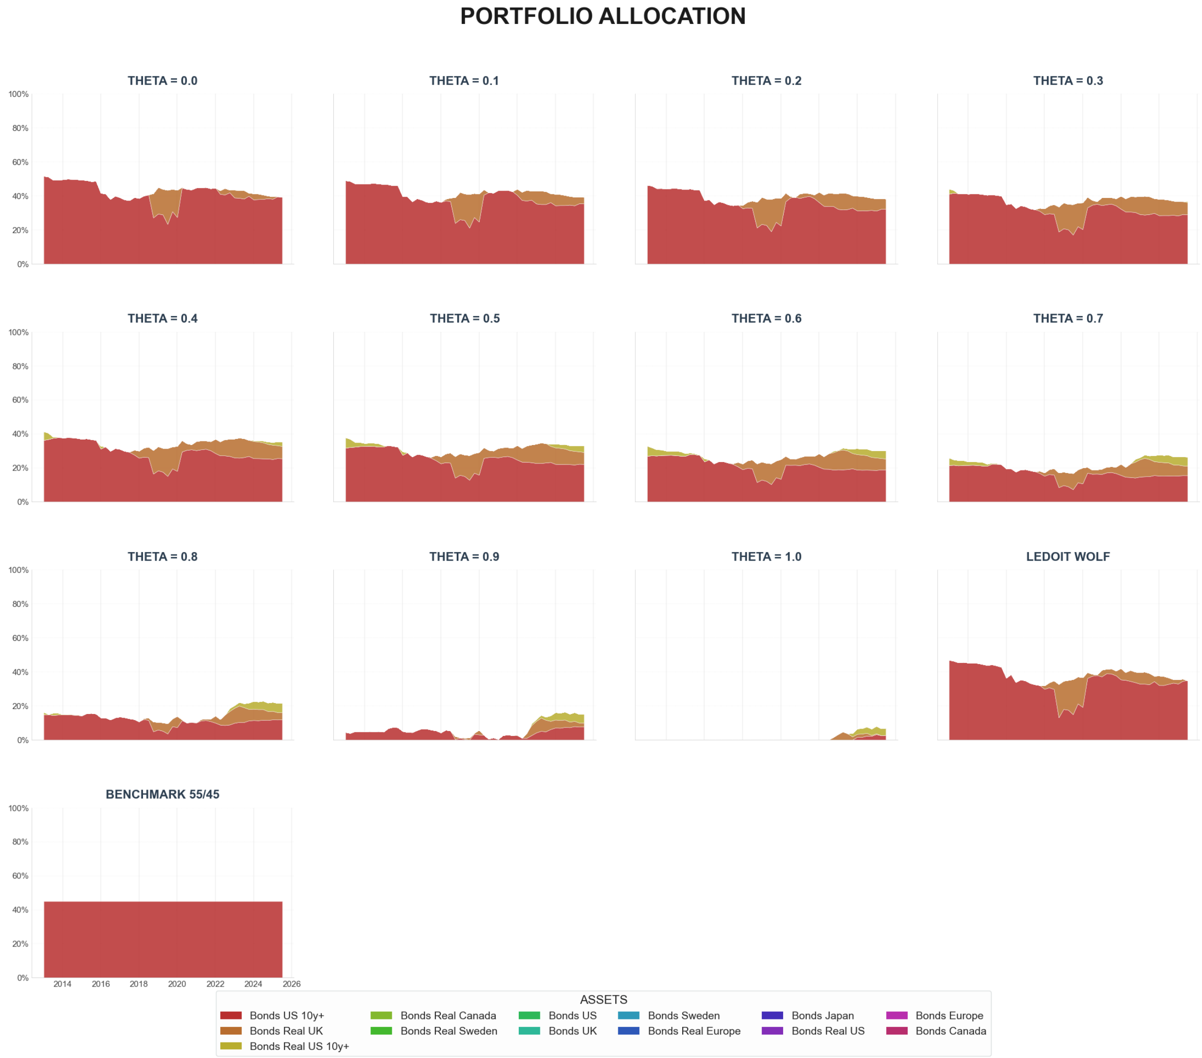

Figure 10: Allocation to Bonds for the different strategies

Figure 11: Allocation to Credit for the different strategies

Figure 12: Allocation to Alternatives for the different strategies

Appendix II - Performance Across Macroeconomic Regimes

Table 3: Performance Metrics across different Macroeconomic Regimes.
$\theta = 0.0$ represents no shrinkage.

Strategy	CAGR (%)	Vol (%)	Sharpe	Max DD (%)	VaR 95% (%)	CVaR 95% (%)
Regime: Entire Dataset						
$\theta = 0.0$	3.95	9.02	0.21	-25.43	-0.88	-1.31
$\theta = 0.1$	4.13	9.25	0.22	-25.47	-0.87	-1.32
$\theta = 0.2$	4.28	9.52	0.24	-25.97	-0.85	-1.33
$\theta = 0.3$	4.41	9.82	0.24	-26.99	-0.86	-1.35
$\theta = 0.4$	4.60	10.13	0.26	-27.71	-0.85	-1.38
$\theta = 0.5$	4.86	10.48	0.28	-28.07	-0.88	-1.42
$\theta = 0.6$	5.21	10.88	0.30	-28.37	-0.92	-1.48
$\theta = 0.7$	5.69	11.37	0.34	-29.96	-0.97	-1.57
$\theta = 0.8$	6.35	12.02	0.38	-31.86	-1.03	-1.69
$\theta = 0.9$	7.14	12.87	0.42	-34.24	-1.15	-1.86
$\theta = 1.0$	7.59	13.40	0.44	-36.33	-1.23	-1.97
Benchmark 55/45	6.58	9.46	0.48	-25.14	-0.92	-1.43
Ledoit-Wolf	3.55	9.38	0.16	-24.47	-0.90	-1.38
Regime: Boom						
$\theta = 0.0$	4.92	3.21	-0.09	-7.24	-0.81	-1.06
$\theta = 0.1$	5.45	3.42	-0.03	-7.19	-0.78	-1.07
$\theta = 0.2$	5.88	3.65	0.02	-7.16	-0.78	-1.09
$\theta = 0.3$	6.28	3.85	0.05	-7.24	-0.79	-1.10
$\theta = 0.4$	6.85	4.05	0.10	-7.34	-0.82	-1.13
$\theta = 0.5$	7.54	4.25	0.15	-7.49	-0.85	-1.17
$\theta = 0.6$	8.33	4.45	0.20	-7.65	-0.86	-1.23
$\theta = 0.7$	9.24	4.66	0.25	-8.06	-0.87	-1.32
$\theta = 0.8$	10.36	4.92	0.31	-8.59	-0.95	-1.44
$\theta = 0.9$	11.72	5.24	0.38	-9.24	-1.00	-1.58
$\theta = 1.0$	12.46	5.36	0.42	-9.39	-1.04	-1.64
Benchmark 55/45	12.49	4.69	0.47	-8.63	-0.76	-1.10
Ledoit-Wolf	4.19	3.33	-0.15	-10.06	-0.80	-1.17
Regime: Expansion						
$\theta = 0.0$	8.72	5.17	0.40	-9.98	-0.66	-0.97
$\theta = 0.1$	9.12	5.25	0.43	-9.76	-0.65	-0.96
$\theta = 0.2$	9.50	5.31	0.46	-9.61	-0.65	-0.96
$\theta = 0.3$	9.88	5.34	0.49	-9.78	-0.63	-0.97
$\theta = 0.4$	10.30	5.39	0.52	-10.51	-0.64	-0.98
$\theta = 0.5$	10.78	5.49	0.55	-11.33	-0.67	-1.00
$\theta = 0.6$	11.39	5.65	0.59	-12.31	-0.71	-1.03
$\theta = 0.7$	12.20	5.88	0.63	-13.66	-0.75	-1.09
$\theta = 0.8$	13.30	6.25	0.67	-15.35	-0.85	-1.18
$\theta = 0.9$	14.58	6.73	0.71	-17.52	-0.94	-1.30
$\theta = 1.0$	15.05	6.95	0.72	-19.14	-1.00	-1.36
Benchmark 55/45	9.74	4.71	0.54	-7.76	-0.65	-0.97
Ledoit-Wolf	8.77	5.36	0.40	-11.90	-0.67	-1.02
Regime: Recovery						
$\theta = 0.0$	-1.44	3.19	-1.02	-20.82	-0.86	-1.16
$\theta = 0.1$	-1.24	3.15	-1.03	-21.51	-0.82	-1.17
$\theta = 0.2$	-0.95	3.16	-1.01	-22.55	-0.81	-1.18
$\theta = 0.3$	-0.58	3.20	-0.98	-23.61	-0.79	-1.18

Continued on next page...

Table 3 – continued from previous page

Strategy	CAGR (%)	Vol (%)	Sharpe	Max DD (%)	VaR 95% (%)	CVaR 95% (%)
$\theta = 0.4$	-0.10	3.24	-0.95	-24.20	-0.77	-1.20
$\theta = 0.5$	0.44	3.31	-0.90	-24.41	-0.77	-1.24
$\theta = 0.6$	0.97	3.40	-0.86	-24.19	-0.82	-1.31
$\theta = 0.7$	1.63	3.50	-0.81	-23.49	-0.86	-1.40
$\theta = 0.8$	2.44	3.64	-0.74	-22.30	-0.97	-1.52
$\theta = 0.9$	3.29	3.83	-0.68	-20.54	-1.04	-1.69
$\theta = 1.0$	4.31	4.01	-0.61	-22.24	-1.12	-1.79
Benchmark 55/45	5.64	3.29	-0.70	-19.72	-0.82	-1.29
Ledoit-Wolf	-2.95	3.19	-1.09	-20.66	-0.91	-1.21
Regime: Slowdown						
$\theta = 0.0$	-4.56	5.27	-0.63	-20.23	-1.08	-1.38
$\theta = 0.1$	-4.46	5.42	-0.61	-20.25	-1.11	-1.38
$\theta = 0.2$	-4.51	5.61	-0.59	-20.54	-1.16	-1.40
$\theta = 0.3$	-4.74	5.83	-0.58	-21.10	-1.12	-1.43
$\theta = 0.4$	-5.06	6.08	-0.56	-21.59	-1.14	-1.49
$\theta = 0.5$	-5.29	6.33	-0.55	-21.92	-1.17	-1.55
$\theta = 0.6$	-5.37	6.59	-0.53	-22.14	-1.21	-1.63
$\theta = 0.7$	-5.33	6.94	-0.49	-22.20	-1.26	-1.74
$\theta = 0.8$	-5.16	7.40	-0.45	-22.26	-1.31	-1.89
$\theta = 0.9$	-4.96	8.00	-0.41	-22.28	-1.48	-2.11
$\theta = 1.0$	-5.06	8.25	-0.40	-22.54	-1.53	-2.23
Benchmark 55/45	-8.56	5.29	-0.83	-22.71	-1.15	-1.67
Ledoit-Wolf	-3.86	5.60	-0.56	-19.72	-1.14	-1.41
Regime: Contraction						
$\theta = 0.0$	1.78	5.13	-0.41	-25.43	-1.04	-1.74
$\theta = 0.1$	1.51	5.33	-0.40	-25.47	-1.03	-1.77
$\theta = 0.2$	1.29	5.54	-0.40	-25.97	-1.07	-1.80
$\theta = 0.3$	1.01	5.78	-0.39	-26.99	-1.05	-1.86
$\theta = 0.4$	0.81	6.02	-0.38	-27.71	-1.06	-1.92
$\theta = 0.5$	0.71	6.27	-0.36	-28.07	-1.12	-1.98
$\theta = 0.6$	0.67	6.56	-0.35	-28.37	-1.18	-2.05
$\theta = 0.7$	0.66	6.91	-0.33	-29.96	-1.27	-2.14
$\theta = 0.8$	0.69	7.35	-0.30	-31.86	-1.29	-2.28
$\theta = 0.9$	0.79	7.94	-0.27	-34.24	-1.36	-2.49
$\theta = 1.0$	0.73	8.41	-0.25	-36.33	-1.51	-2.68
Benchmark 55/45	2.95	5.22	-0.35	-25.14	-1.19	-1.96
Ledoit-Wolf	1.43	5.44	-0.40	-24.47	-1.09	-1.87
Regime: Recession						
$\theta = 0.0$	12.69	5.24	0.06	-23.85	-0.91	-1.38
$\theta = 0.1$	13.01	5.42	0.08	-24.51	-0.85	-1.36
$\theta = 0.2$	13.14	5.61	0.09	-25.52	-0.82	-1.34
$\theta = 0.3$	13.38	5.79	0.10	-26.56	-0.85	-1.32
$\theta = 0.4$	13.76	5.98	0.11	-27.16	-0.85	-1.32
$\theta = 0.5$	14.08	6.18	0.13	-27.40	-0.94	-1.34
$\theta = 0.6$	14.52	6.38	0.14	-27.25	-1.02	-1.41
$\theta = 0.7$	15.13	6.57	0.16	-26.62	-1.03	-1.51
$\theta = 0.8$	15.97	6.77	0.19	-25.68	-1.18	-1.65
$\theta = 0.9$	17.26	7.00	0.24	-28.64	-1.24	-1.82
$\theta = 1.0$	18.60	7.15	0.29	-30.55	-1.33	-1.95
Benchmark 55/45	23.06	6.87	0.46	-23.37	-1.00	-1.26
Ledoit-Wolf	11.12	5.14	-0.02	-23.63	-0.85	-1.43

Disclaimer

Disclaimer

These analyses, documents, and any other information originating from LINC Research & Analysis (henceforth "LINC R&A") are created for information purposes only, for general dissemination, and are not intended to be advisory. The information in the analysis is based on sources, data, and persons which LINC R&A believes to be reliable. LINC R&A can never guarantee the accuracy of the information. The forward-looking information found in this analysis is based on assumptions about the future, and is therefore uncertain by nature, and using information found in the analysis should therefore be done with care. Furthermore, LINC R&A can never guarantee that the projections and forward-looking statements will be fulfilled to any extent. This means that any investment decisions based on information from LINC R&A, any employee or person related to LINC R&A, are to be regarded to be made independently by the investor. These analyses, documents, and any other information derived from LINC R&A are intended to be one of several tools involved in investment decisions regarding all forms of investments regardless of the type of investment involved. Investors are urged to supplement with additional relevant data and information, as well as consult a financial adviser before any investment decision. LINC R&A disclaims all liability for any loss or damage of any kind that may be based on the use of analyses, documents, and any other information derived from LINC R&A.

Conflicts of Interest and Impartiality

To ensure LINC R&A's independence, LINC R&A has established compliance rules for analysts. In addition, all analysts have signed an agreement in which they are required to report any conflicts of interest. These terms have been designed to ensure that COMMISSION DELEGATED REGULATION (EU) 2016/958 of 9 March 2016, supplementing Regulation (EU) No 596/2014 of the European Parliament and of the Council concerning regulatory technical standards for the technical arrangements for objective presentation of investment recommendations or other information recommending or suggesting an investment strategy and for disclosure of particular interests or indications of conflicts of interest.

Other