Economics Bulletin

Volume 46, Issue 1

Does removing the effect of short-term co-movements improve portfolio performance over monthly horizons? Dow Jones Industrial Average Analysis

Rafael B Chaves
Universidade Federal do Rio Grande do Sul

Cleiton G Taufemback
Universidade Federal do Rio Grande do Sul

Hudson S Torrent
Universidade Federal do Rio Grande do Sul

Abstract

Portfolio strategies often seek to reduce exposure to short-term market fluctuations while maintaining robust performance over defined investment horizons. This study proposes a modification to the Markowitz model that filters out short-term co-movements in asset returns, aiming to construct portfolios less sensitive to transient fluctuations. Using historical data from the Dow Jones Industrial Average, we evaluate the performance of the proposed method relative to the traditional Markowitz and Naive models over investment horizons of one, three, and six months. The results indicate that portfolios constructed with the proposed approach generally outperform the benchmark models in terms of returns and exhibit statistically significant improvements in certain periods.

Editor`s Note: This paper was originally submitted under a different manuscript number on 12/29/202 and accepted for publication on 03/16/2026. A system error required us to replace the original submission with a new manuscript with a later manuscript number.
Citation: Rafael B Chaves and Cleiton G Taufemback and Hudson S Torrent, (2026) "Does removing the effect of short-term co-movements improve portfolio performance over monthly horizons? Dow Jones Industrial Average Analysis", *Economics Bulletin*, Volume 46, Issue 1, pages 41-54
Contact: Rafael B Chaves - rafaelbechaves@gmail.com, Cleiton G Taufemback - cleiton.taufemback@ufrgs.br, Hudson S Torrent - hudsontorrent@gmail.com.
Submitted: April 29, 2026.   **Published:** March 30, 2026.

1. Introduction

In a low-interest rate environment, investors may be tempted to allocate more to equities in search of higher returns over shorter horizons. However, Heugh and Fox (2018) highlights that active share portfolios, as introduced by Cremers and Petajisto (2009), combined with disciplined investment strategies, can help investors benefit from the short-term behavior of other market participants. Adopting a strategy that focuses on more persistent return components requires patience and the willingness to tolerate a certain level of risk while waiting for the expected gains to materialize. Markowitz's modern portfolio theory (Markowitz, 1952) provides a framework for optimizing the trade-off between risk and return. Nevertheless, this approach depends heavily on the estimation of the covariance matrix of returns, which can be challenging, particularly when returns are dominated by short-term fluctuations.

A great deal of the literature on portfolio selection focuses on the covariance matrix of returns, proposing various methods to estimate it more accurately (see Pafka and Kondor (2004)). For instance, Ledoit and Wolf (2003a) and Ledoit and Wolf (2003b) introduced shrinkage techniques to improve the estimation of covariance matrices. Zhifeng and Fei (2019) propose to reduce estimation error in the mean-variance portfolio model by considering regularization in the objective function. Another branch of literature focuses on filtering with the aim of reducing noise and improving the stability of portfolio weights. Park and O'Leary (2010) proposed a Tikhonov filtering approach to regularize the covariance matrix. Daly et al. (2008) apply random matrix theory (RMT) filters to covariance matrices of financial returns, and shown improvements to the optimization of stock portfolios.

To construct an investment portfolio focused on more stable return components, Cochrane (2021) argues that investors should prioritize assets with low long-horizon volatility to capture more persistent returns. Although variable-cost assets such as stocks exhibit high volatility at daily frequencies, this volatility becomes less relevant over longer horizons, as noted by Adrian and Rosenberg (2008). In this paper, we propose a portfolio construction method that reduces the impact of short-term co-movements in returns, which may be less relevant depending on the investor's time horizon. By filtering out high-frequency components or rapid fluctuations, the covariance matrix estimated from the filtered data highlights relationships driven by lower-frequency components or more persistent trends in stock returns. This approach can be particularly beneficial when the focus is on medium-term dynamics rather than short-term noise.

Compared to our approach, Park and O'Leary (2010) regularizes the covariance matrix by shrinking its eigenvalues to reduce estimation noise and ensure invertibility. However, this method applies uniform shrinkage across all components, without accounting for the temporal dynamics embedded in asset returns. In contrast, we propose selectively removing high-frequency components from the return series, thereby filtering out short-term co-movements that may be less relevant for portfolio decisions over multi-month horizons. Additionally, frequency-based filtering offers more intuitive parameter control through a frequency cutoff that can be directly aligned with the investor's target horizon, enhancing both methodological transparency and practical applicability. Importantly, Corbae et al. (2002) and Taufemback (2023) show that relationships among variables can vary substantially across frequencies and provide formal tools to capture these frequency-dependent dynamics. Their findings reinforce the importance of modeling how different components across the frequency domain contribute to asset co-movements, further motivating our frequency-based approach.

Regarding co-movements, the Dieckelmann et al. (2024) report highlights how passive investing has increased stock price correlations - when passive funds buy or sell based on index composition rather than individual company performance, stocks move together more often. This makes it difficult to assess a company's true value, since stock prices now respond more to overall index movements than to company-specific news.

In response, contemporary research has focused on advanced filtering techniques to mitigate these distortions. Salas-Molina and Nin (2026) addresses instability in hierarchical portfolio selection by improving the signal-to-noise ratio in covariance and correlation matrix estimation, particularly when dealing with limited observations. Similarly, Bongiorno et al. (2026) develops a neural network approach that learns to regularize both eigenvalues and marginal volatilities of large covariance matrices, optimizing directly for out-of-sample portfolio variance minimization. These studies collectively underscore the necessity of the frequency-based filtering approach we propose, which aims to isolate the investment-relevant components of asset returns from short-term fluctuations.

Thus, the main goal of this research is to propose a portfolio construction method that accounts for the risk associated with assets while reducing the influence of short-term volatility in return series. Beyond methodological contributions, this study offers practical implications for asset managers and long-term investors, such as pension funds, who aim to minimize turnover costs associated with short-term volatility. By filtering out transient fluctuations, our approach aids allocators in constructing portfolios that align better with multi-month rebalancing schedules, potentially reducing transaction costs and improving risk-adjusted returns.

As detailed in Section 3, the method estimates the covariance matrix from return data filtered through a low-pass filter, which aims to attenuate short-term fluctuations in asset prices. This filtering enhances the stability and performance of portfolios over multi-month investment horizons. Beyond Section 3, the paper is organized as follows: theoretical foundations are reviewed in Section 2; data sources and processing are described in Section 4; empirical results are analyzed in Section 5; and conclusions are presented in Section 6.

2. Modern Portfolio Theory and Filtering

According to Markowitz (1952), the process of building a portfolio consists of two stages: the first is the selection of stocks, which is based on observations and experiences that reinforce beliefs about the performance of future returns. The second stage is to combine these beliefs about various assets in order to build a portfolio that can achieve the desired return with the least amount of risk.

According to Markowitz (1999), it is implied in this method that one way to reduce the variability of returns is by the diversification of assets in a portfolio, with an expected return with a low risk value. This method of dealing with the risk associated with investment has been criticized in Stoyanov et al. (2011), but it has also been shown to be effective in managing the overall risk of the portfolio.

The method of modern portfolio theory explained using matrix algebra is discussed in the literature on asset prediction and portfolio theory, as demonstrated in Cochrane (2008) and Constantinides and Malliaris (1995).

As demonstrated in Ingersoll (1987) and Cochrane (2001), given any set of time series of log-returns Y, i.e., $Y = \{y_{1,t}, y_{2,t}, \dots, y_{n,t}\}_{t=1}^{T}$, then the set of averages is given by

$\mu = \{\mu_1, \ldots, \mu_n\}$, being n the number of assets, and its covariance matrix is given by Σ.

According to Sharpe (1966), the best portfolio will be one that crosses an efficient border. According to Bailey and Lopez de Prado (2012), when defining the Efficient Frontier with Sharpe Ratio, we are looking at a set of stocks that provide the highest expected return in relation to risk.

Given a set of assets and one risk-free activity, the portfolio designed to maximize the Sharpe Ratio is known as a Markowitz-Sharpe Portfolio, according to Ingersoll (1987). The Markowitz-Sharpe weights are denoted by $\mathbf{t} = (t_1, \ldots, t_n)'$, which solves the following problem of restricted maximization.

$$\mathbf{t} = \arg\max_{\tau} \frac{\tau'\mu - r^*}{(\tau'\Sigma\tau)^{\frac{1}{2}}} \text{ s.t. } \tau'\mathbf{1} = 1 \text{ and } \tau_i > 0, \forall i, i = 1 \ldots, n, \tag{1}$$

where $\mu_{p,t} = \mathbf{t}'\mu$ e $\sigma_{p,t} = (\mathbf{t}'\Sigma\mathbf{t})^{\frac{1}{2}}$. The value of τ that solves the optimization problem (1) is:

$$\mathbf{t} = \frac{\Sigma^{-1}(\mu - r^* \cdot \mathbf{1})}{\mathbf{1}'\Sigma^{-1}(\mu - r^* \cdot \mathbf{1})}. \tag{2}$$

2.1 Low-pass Filter

Demonstrated in Christiano and Fitzgerald (2003), the low-pass filter is a linear data transformation that produces only the primary components within a predetermined frequency range, removing the remainder of the information.

According to Christiano and Fitzgerald (1998), there is a straightforward method for extracting components from a frequency band in a time series. Given the temporal series $\{y\}_{t=1}^{n}$, it can be expressed as

$$y_t = \int_0^{\pi} [a(\omega)\cos(\omega t) + b(\omega)\sin(\omega t)]d\omega, \quad \forall t,$$

or as

$$y_t = a_0 + a_1 \cos(t\omega_1) + b_1 \sin(t\omega_1) + \ldots + a_{T/2} \cos\left(t\omega_{T/2}\right) + b_{T/2} \sin\left(t\omega_{T/2}\right),$$

for $t = 1, \ldots, T$, and T being the total number of observations of the original time series and where we have

$$\omega_j = 2\pi j/T, \quad j = 1, \ldots, T/2. \tag{3}$$

Thus, if we want to filter the series $\{y_t\}_{t=1}^{T}$ by allowing only frequencies below ω_f, we can do so by defining

$$X = \begin{bmatrix} 1 & \cos(\omega_1) & \sin(\omega_1) & \cdots & \cos(\omega_f) & \sin(\omega_f) \\ 1 & \cos(2\omega_1) & \sin(2\omega_1) & \cdots & \cos(2\omega_f) & \sin(2\omega_f) \\ \vdots & \vdots & & \vdots & \vdots & \\ 1 & \cos(T\omega_1) & \sin(T\omega_1) & \cdots & \cos(T\omega_f) & \sin(T\omega_f) \end{bmatrix}.$$

and obtain the low-pass filter components by

$$\hat{\beta} = (X'X)^{-1} X'Y$$

where $\beta = \{a_0, a_1, b_1, \ldots, a_f, b_f\}'$. Therefore, the filtered time series, $\hat{Y}_f$, is estimated as

$$\hat{Y}_f = X\hat{\beta}. \tag{4}$$

3. Proposed Model

The proposed model extends the Markowitz-Sharpe methodology, as discussed in Section 2, particularly the equation (1). In this approach, the covariance matrix, Σ, is replaced with one derived from the filtered series (see equation 4). By removing the short-term covariance of returns, the portfolio weights are expected to be less sensitive to rapid and temporary fluctuations.

Let Y_f be a collection of log-return series obtained by projecting Y onto sines and cosines within the frequency range from 0 to $\omega_f = 2\pi f/T$, i.e.,

$$Y_f = \{y^f_{1,t}, \ldots, y^f_{n,t}\}_{t=1}^T,$$

where n represents the number of assets in the set, then the matrix Σ_f denotes the correlation matrix of the assets in Y_f. Finally, the portfolio is constructed using a risk-free asset and the filtered covariance matrix given by

$$\mathbf{t_f} = \arg\max_{\tau} \frac{\tau'\mu - r^*}{(\tau'\Sigma_f\tau)^{\frac{1}{2}}} \text{ given that } \tau'\mathbf{1} = 1, \tag{5}$$

where the equation (5) defines what we call the Horizon-Adjusted Markowitz or HAM.

As a result, the allocation of positions in the portfolio, given the presence of a risk-free asset, is determined by maximizing the value of $\mathbf{t_f}$, ensuring that the difference between the expected return and the risk-free rate is as significant as possible. Note that we use the expected value of the original series and not from the filtered series since the primary objective of filtering is to remove any high-frequency variation components and co-movements. Furthermore, the expected values of filtered series misrepresent the gains and losses of those series over the analyzed time.

4. Data and Methods Comparison

The data comprises all companies belonging to the Dow Jones Industrial Average, and were obtained through the Yahoo Finance for the period between January 1, 2021 and January 31, 2025. In that period, the index had 2 alterations on the companies that composed it, totalizing 33 companies being a part of the Dow Jones, please see Table IV. During the analyzed period, Amazon replaced Walgreens Boots Alliance on February 26, 2024. Subsequently, on November 8, 2024, Nvidia replaced Intel, and Sherwin-Williams replaced Dow Inc.

Following standard practice in financial literature, e.g., O'Neil and Ryan (2002); Huang et al. (2025); Mostafavi and Hooman (2025), we utilize adjusted closing prices to compute asset returns. The adjusted close price accounts for corporate actions such as stock splits and dividend distributions, ensuring that the return series reflects the total return available to an investor, rather than just capital appreciation. This distinction is relevant for long-horizon portfolio analysis, since ignoring dividend reinvestment can lead to significant downward bias in cumulative performance metrics.

The *Interest Rates, Discount Rate for the United States* series, obtained at the FED St. Louis, was used as a risk-free asset in the portfolio composition. The original series is annually divided into 252 days and converted into daily gains.

Regarding the comparison between methods, we consider the return at time t, for a given method m and horizon h, denoted by $r_{t,m,h}$. Based on this, the following comparison metrics were considered:

$$\Lambda_{m,h} = \sum_{t=1}^{T} r_{t,m,h}, \quad \hat{\mu}_{m,h} = \frac{\Lambda_{m,h}}{T}, \quad \hat{\sigma}^2_{m,h} = \frac{1}{T-1}\sum_{t=1}^{T}(r_{t,m,h} - \hat{\mu}_{m,h})^2. \tag{6}$$

We denote with $\Lambda_{m,h}$ the sum of accumulated returns for the method m given the horizon h, $\hat{\mu}_{m,h}$ is the average return value for the method m, given the horizon h and $\hat{\sigma}^2_{m,h}$ represents the variability of the method m, for the horizon h.

To evaluate the robustness of our proposed method in comparison with the benchmark models we employ the Ledoit-Wolf Test. According to Ledoit and Wolf (2008), there are two investment strategies i and j, whose additional return in comparison to a benchmark return at a point t is r_{ti} and r_{tj}, respectively. The benchmark return in this case is typically given by the return of a risk-free asset, but it might also be the ticker of another reference asset.

Given a total of T pairs of returns in t moments, as determined by the set of returns $(r_{1i}, r_{1j})', \ldots, (r_{Ti}, r_{Tj})'$ being observed. It is assumed that these observations constitute a strictly stationary temporal series, so that the multivariate distribution of returns does not change over time. This distribution has a μ median and a Σ covariance matrix.

$$\mu = \begin{pmatrix} \mu_i \\ \mu_j \end{pmatrix} \text{ and } \Sigma = \begin{pmatrix} \sigma_i^2 & \sigma_{\text{ij}} \\ \sigma_{ij} & \sigma_j^2 \end{pmatrix}.$$

The sample means and variances of the observed returns are denoted by $\hat{\mu}_i$, $\hat{\mu}_j$ and $\hat{\sigma}_i^2$, $\hat{\sigma}_j^2$, respectively. The difference between the Sharpe ratios i and j is given by

$$\Delta = \frac{\mu_i}{\sigma_i} - \frac{\mu_j}{\sigma_j}$$

and the estimated difference is given by

$$\hat{\Delta} = \frac{\hat{\mu}_i}{\hat{\sigma}_i} - \frac{\hat{\mu}_j}{\hat{\sigma}_j}.$$

Furthermore, take into account $u = \left(\mu_i, \mu_j, \sigma_i^2, \sigma_j^2\right)'$ e $\hat{u} = (\hat{\mu}_i\hat{\mu}_j, \hat{\sigma}_i^2, \hat{\sigma}_i^2)'$. Memmel (2003) calculates the error for $\hat{\Delta}$ based on the following

$$\sqrt{T}(\hat{u} - u) \xrightarrow{d} N(0; \Omega),$$

where $\xrightarrow{d}$ refers to convergence in distribution and the use of the delta method However, Ω was based on the i.i.d. return of a normal bivariate distribution, which assumes that

$$\Omega = \begin{pmatrix} \sigma_i^2 & \sigma_{\text{in}} & 0 & 0 \\ \sigma_{\text{ij}} & \sigma_n^2 & 0 & 0 \\ 0 & 0 & 2\sigma_i^4 & 2\sigma_{\text{ij}}^2 \\ 0 & 0 & 2\sigma_{\text{ij}}^2 & 2\sigma_j^4 \end{pmatrix}.$$

The above formula is invalid if the distribution of returns does not follow a normal distribution or if the observations are not correlated over time. To address this condition, Ledoit and Wolf (2008) suggests using Heteroskedasticity- and autocorrelation-consistent estimators, or HAC, for the covariance matrix, as well as bootstrap techniques.

5. Empirical Study

This section presents an empirical evaluation of the proposed filtering-based portfolio optimization method, as detailed in Section 3. The objective is to assess whether incorporating frequency-domain information into the estimation of expected returns and the covariance matrix can improve the out-of-sample performance of optimized portfolios. Three investment horizons are analyzed: one month, three months, and six months.

For each re-balancing date t, the portfolios are constructed using one year of historical data to estimate the vector of expected returns and the covariance matrix. Given the assumption of 252 trading days in a year, this results in a rolling window of size $T = 252$. The dataset includes adjusted stock prices, and the assets under analysis correspond to the components of the Dow Jones index, see Section 4. The evaluation period begins at the start date of the database and ends at the last feasible date for each horizon.

The filtered expected returns are obtained by applying a low-pass filter to the return series using the method described in equation (3). Specifically, the set of frequencies is defined by $f = \{15, 16, \ldots, 100\}$, and for each frequency f, a Fourier basis composed of sine and cosine terms up to order f is constructed.

For each combination of investment horizon and frequency, the performance of three portfolio strategies is evaluated: (i) a naive equally weighted portfolio, (ii) the Markowitz-Sharpe portfolio using unfiltered inputs, and (iii) the Horizon-Adjusted Markowitz method (HAM_f, where f represent the maximum frequency). Portfolio weights are recalculated at each time step using rolling estimations, and their out-of-sample performance is assessed through the Sharpe ratio over a moving evaluation window of twenty-one days.

Although we evaluated a range of frequencies, for each investment horizon, we report only the results corresponding to the value of f that produced the best performance metrics (i.e., highest returns and Sharpe Ratios). It is important to note that the high accumulated log-returns obtained with all methods should be interpreted with caution, as the results do not account for transaction costs or the operational constraints involved in continuously reinvesting capital and maintaining overlapping positions. These factors can significantly reduce the realized profitability in practice. However, this metric remains valuable for understanding the relative performance and behavior of the strategies over time.

1 month horizon

Table I presents the descriptive statistics for the one-month horizon. The proposed method, HAM_{34}, achieved the highest average log-return of 0.0078 and the highest cumulative log-return of 5.8507, outperforming both the Naive and Markowitz-Sharpe strategies. While its standard deviation was slightly higher than that of the Naive method, it remained lower than that of the Markowitz-Sharpe approach.

Figure 1 illustrates the rolling mean of 21-day log-returns. Moments where the Sharpe Ratio (SR) showing statistically significant for each method is shown. HAM_{34} exhibited superior performance over the Naive method with 33.33% more instances. When compared to the Markowitz-Sharpe method, HAM_{34} recorded 18.75% more periods of significative higher SR.

Table I: Descriptive comparison of results for $h = 1m$.

	$\hat{\mu}_m$	$\hat{\sigma}_m$	Λ_m
Naive	0.0051	0.0416	3.8576
Markowitz-Sharpe	0.0067	0.0451	5.0405
HAM_{34}	0.0078	0.0424	5.8507

Note: $\hat{\mu}_m$, $\hat{\sigma}_m$ e Λ_m represents the average of log-returns, the standard deviation and the cumulative return using each method, see Equation (6). Results for the range between Jan 1st of 2021 to Jan 31th of 2025.

3 months horizon

As shown in Table II, the HAM_{31} method again delivered the best results across all performance metrics. It yielded an average log-return of 0.0289 and a cumulative log-return of 20.4843, both of which were notably higher than those achieved by the benchmark models. Although it had the highest standard deviation, this was offset by superior overall gains.

Figure 2 confirms the outperformance of HAM_{31}. Compared to the Naive method, it registered 100.07% more instances of higher SR. Remarkably, the Markowitz-Sharpe method did not produce any periods in which it outperformed HAM_{31} with statistical significance.

Table II: Descriptive comparison of results for $h = 3m$.

	$\hat{\mu}_m$	$\hat{\sigma}_m$	Λ_m
Naive	0.0181	0.0588	12.8335
Markowitz-Sharpe	0.0224	0.0638	15.8620
HAM_{31}	0.0289	0.0662	20.4843

Note: $\hat{\mu}_m$, $\hat{\sigma}_m$ e Λ_m represents the average of log-returns, the standard deviation and the cumulative log-return using each method, see Equation (6). Results for the range between Jan 1st of 2021 to Jan 31th of 2025.

6 months horizon

For the six-month horizon, HAM_{25} continued to outperform the benchmark strategies, as shown in Table III. It achieved the highest average log-return of 0.0503 and cumulative log-return of 32.4792, with a moderate standard deviation.

Figure 3 highlights the relative performance advantage of HAM_{25}. It exhibited 20% more instances of superior SR than the Naive method. Compared to the Markowitz-Sharpe strategy, HAM_{25} outperformed with 62.5% more periods of statistical superior SR.

Table III: Descriptive comparison of results for $h = 6m$.

	$\hat{\mu}_m$	$\hat{\sigma}_m$	Λ_m
Naive	0.0416	0.0769	26.8580
Markowitz-Sharpe	0.0470	0.0825	30.3552
HAM_{25}	0.0503	0.0789	32.4792

Note: $\hat{\mu}_m$, $\hat{\sigma}_m$ e Λ_m represents the average of log-returns, the standard deviation and the cumulative log-return using each method, see Equation (6). Results for the range between Jan 1st of 2021 to Jan 31th of 2025.

6. Conclusion

This study introduced an alternative to the traditional Markowitz-Sharpe portfolio optimization framework by incorporating a short-term fluctuation filtering approach aimed at reducing return variability. The analysis was conducted using investment periods of one, three, and six months. The U.S. basic interest rate served as the benchmark for the risk-free asset. The proposed methodology demonstrated consistent improvements over the traditional approach in return and risk-adjusted performance across all horizons, highlighting the potential of signal filtering techniques in enhancing portfolio allocation strategies.

This study demonstrates that filtering short-term co-movements improves portfolio performance over monthly horizons, offering practical benefits for institutional investors with quarterly rebalancing schedules by reducing noise-driven volatility. Nevertheless, some limitations require attention: the choice of frequency cutoff remains somewhat arbitrary and may require adjustment across different market conditions, and our reliance on historical U.S. equity data limits generalizability to other markets or asset classes. Future work could explore the robustness of this methodology in emerging markets, cryptocurrency portfolios, or mixed-asset environments, and investigate whether adaptive filtering techniques - such as those employing machine learning to dynamically select cutoff frequencies - outperform the fixed-horizon approach proposed here. Additionally, comparing frequency-domain filtering with alternative noise reduction methods including hierarchical clustering approaches, such Salas-Molina and Nin (2026), would provide valuable insights into the relative effectiveness of different covariance matrix cleaning strategies.

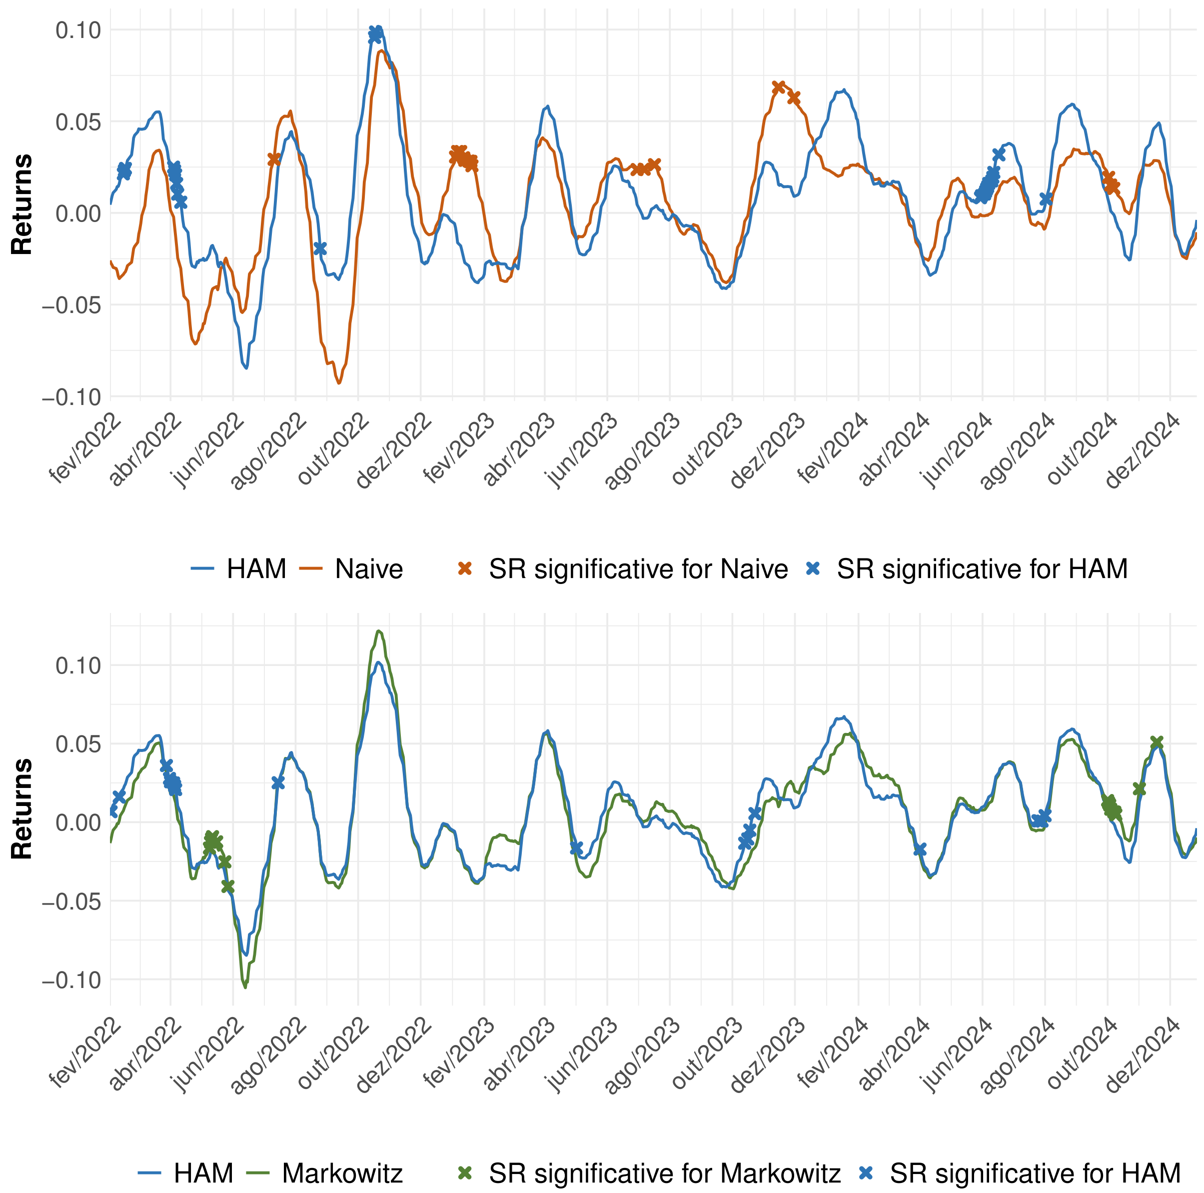

Figure 1: Rolling mean returns of twenty one days for one month horizon, HAM_{34} and Naive methods at top and HAM_{34} and Markowitz-Sharpe methods at bottom.

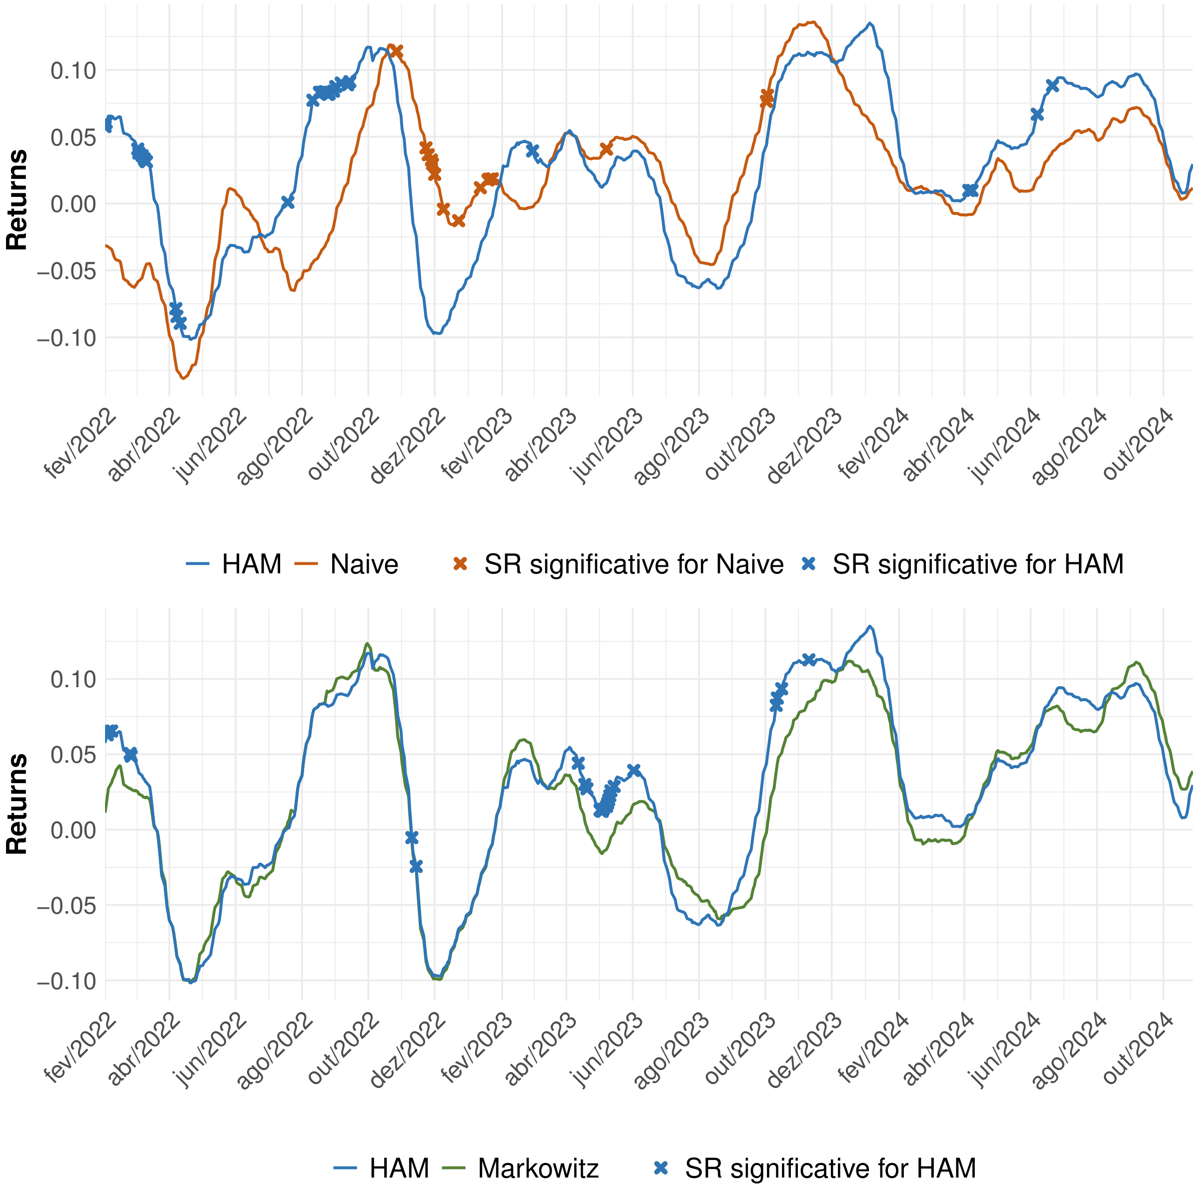

Figure 2: Rolling mean returns of twenty one days for three months horizon, HAM_{31} and Naive methods at top and HAM_{31} and Markowitz-Sharpe methods at bottom.

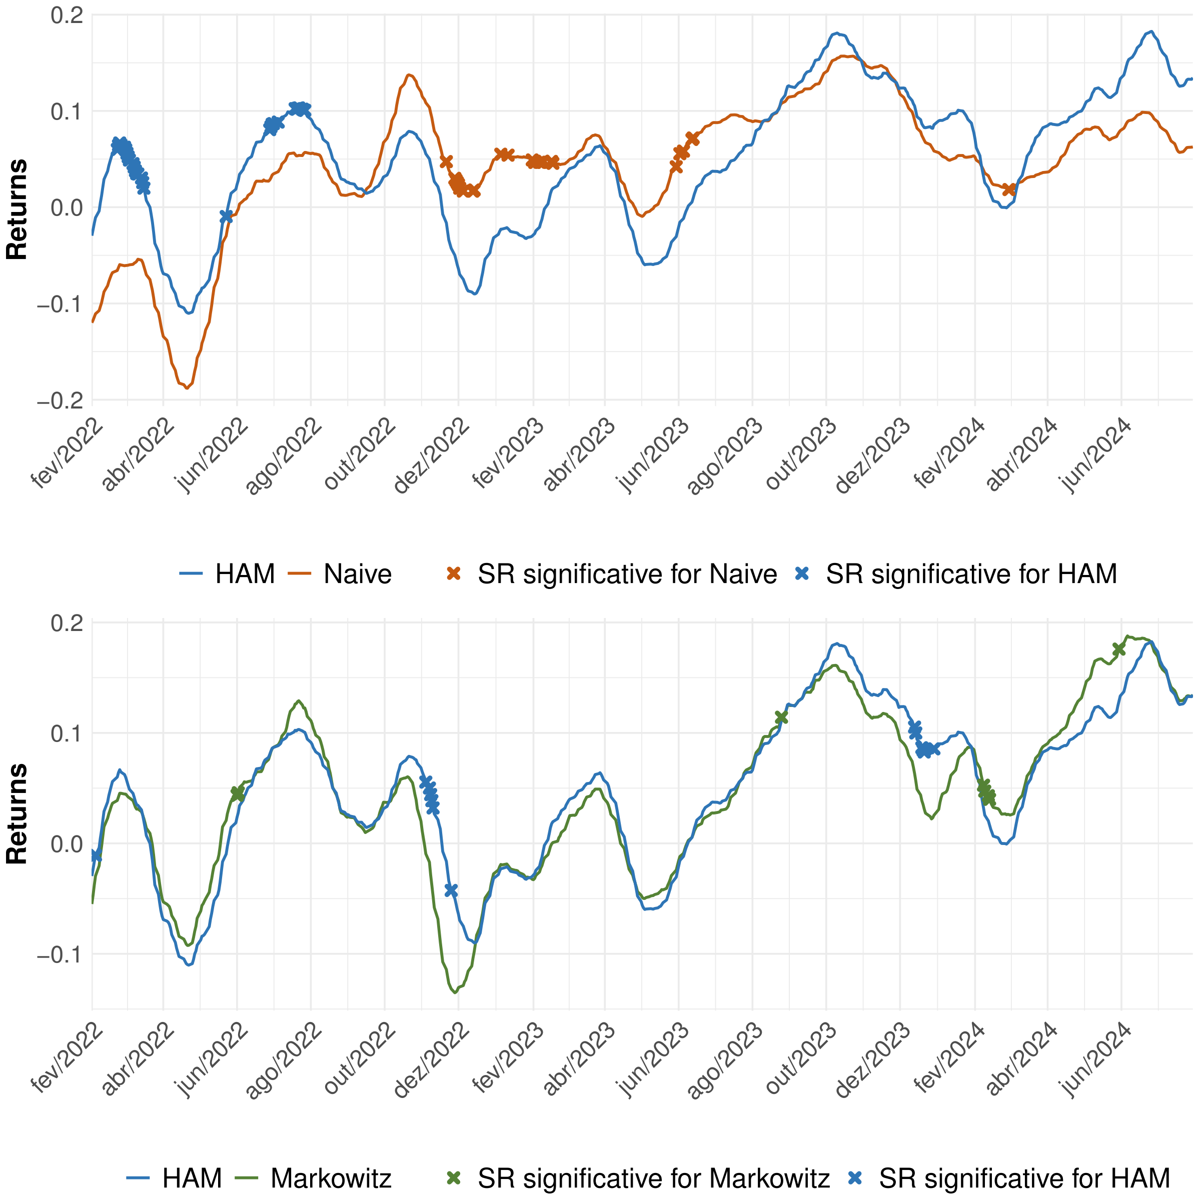

Figure 3: Rolling mean returns of twenty one days for six months horizon, HAM_{25} and Naive methods at top and HAM_{25} and Markowitz-Sharpe methods at bottom.

Table IV: Companies present in the Dow Jones Industrial Average Index, between 2021 and 2025.

Industry	**Ticker**	**Added in**	**Removed in**
3M Company	MMM	Aug 09 1976	-
The Goldman Sachs Group, Inc.	GS	Sep 23 2013	-
Nike, Inc.	NKE	Sep 23 2013	-
American Express Company	AXPB	Aug 30 1982	-
The Home Depot, Inc.	HOME	Nov 01 1999	-
The Procter & Gamble Company	PG	Aug 15 1933	-
Amgen Inc.	AMGN	Aug 31 2020	-
Honeywell International Inc.	HON	Aug 31 2020	-
salesforce.com, inc.	CRM	Aug 31 2020	-
Apple Inc.	AAPL	Mar 19 2015	-
Intel Corporation	ITLC	Nov 01 1999	Nov 08 2024
The Travelers Companies, Inc.	TRV	Jun 08 2009	-
The Boeing Company	BA	Mar 12 1987	-
IBM Corporation	IBM	May 26 1932	-
UnitedHealth Group Incorporated	UNH	Sep 24 2012	-
Caterpillar Inc.	CAT	May 06 1991	-
Johnson & Johnson	JNJ	Mar 17 1997	-
Verizon Communications Inc.	VERZ	Apr 08 2004	-
Chevron Corporation	CVX	Feb 19 2008	-
JPMorgan Chase & Co.	JPM	Mar 17 1997	-
Visa Inc.	V	Sep 23 2013	-
Cisco Systems, Inc.	CSCO	Jun 08 2009	-
McDonald's Corporation	MCD	Oct 30 1985	-
Walgreens Boots Alliance, Inc.	WBA	Jun 26 2018	Fev 26 2022
The Coca-Cola Company	KO	May 26 1932	-
Merck & Co., Inc.	MRK	Jun 29 1979	-
Walmart Inc.	WMT	Mar 17 1997	-
Dow Inc.	DOW	Apr 02 2019	Nov 08 2024
Microsoft Corporation	MSFT	Nov 01 1999	-
The Walt Disney Company	DIS	Mar 17 1997	-
Amazon	AMZN	Fev 26 2022	-
NVIDIA	NVDA	Nov 08 2024	-
Sherwin Williams	SHW	Nov 08 2024	-

References

ADRIAN, T. AND J. ROSENBERG (2008): "Stock returns and volatility: Pricing the short-run and long-run components of market risk," *The journal of Finance*, 63, 2997–3030.

BAILEY, D. H. AND M. LOPEZ DE PRADO (2012): "The Sharpe ratio efficient frontier," *Journal of Risk*, 15, 13.

BONGIORNO, C., E. MANOLAKIS, AND R. N. MANTEGNA (2026): "End-to-end large portfolio optimization for variance minimization with neural networks through covariance cleaning," *The Journal of Finance and Data Science*, 12, 100179.

CHRISTIANO, L. J. AND T. J. FITZGERALD (1998): "The business cycle: it's still a puzzle," *Economic Perspectives-Federal Reserve Bank Of Chicago*, 22, 56–83.

——— (2003): "The Band Pass Filter," *International Economic Review*, 44, 435–465.

COCHRANE, J. H. (2001): *Asset pricing*, Princeton, NJ: Princeton University Press.

——— (2008): "A mean-variance benchmark for intertemporal portfolio theory," *Manuscript, University of Chicago.*

——— (2021): "Portfolios for long-term investors," *SSRN Electronic Journal.*

CONSTANTINIDES, G. M. AND A. G. MALLIARIS (1995): "Portfolio theory," *Handbooks in operations research and management science*, 9, 1–30.

CORBAE, D., S. OULIARIS, AND P. C. B. PHILLIPS (2002): "Band Spectral Regression with Trending Data," *Econometrica*, 70, 1067–1109.

CREMERS, K. J. M. AND A. PETAJISTO (2009): "How Active Is Your Fund Manager? A New Measure That Predicts Performance," *The Review of Financial Studies*, 22, 3329–3365.

DALY, J., M. CRANE, AND H. J. RUSKIN (2008): "Random matrix theory filters in portfolio optimisation: A stability and risk assessment," *Physica A: Statistical Mechanics and its Applications*, 387, 4248–4260.

DIECKELMANN, D., E. SICILIANO, AND A. SOWIŃSKI (2024): "Passive investing and its impact on return co-movement, market concentration and liquidity in euro area equity markets," Financial stability review, European Central Bank.

HEUGH, K. AND M. FOX (2018): "Long-Term Conviction in a Short-Term World," Tech. rep., Morgan Stanley Investment Management, active Fundamental Equity — Global Opportunity Team — Fund Investment Insight.

HUANG, X., D. P. NEWTON, E. PLATANAKIS, AND C. SUTCLIFFE (2025): "Single-stage portfolio optimization with automated machine learning for M6," *International Journal of Forecasting*, 41, 1450–1460.

INGERSOLL, J. E. (1987): *Theory of financial decision making*, Rowman and Littlefield Studies in Financial Economics, Lanham, MD: Rowman & Littlefield.

Ledoit, O. and M. Wolf (2003a): "Honey, I shrunk the sample covariance matrix," *UPF Economics and Business Working Paper.*

——— (2003b): "Improved estimation of the covariance matrix of stock returns with an application to portfolio selection," *Journal of empirical finance*, 10, 603–621.

——— (2008): "Robust performance hypothesis testing with the Sharpe ratio," *Journal of Empirical Finance*, 15, 850–859.

Markowitz, H. (1952): "Portfolio Selection," *The Journal of Finance*, 7, 77–91.

Markowitz, H. M. (1999): "The early history of portfolio theory: 1600–1960," *Financial analysts journal*, 55, 5–16.

Mostafavi, S. M. and A. R. Hooman (2025): "Impact of global indices on forecasting the S&P 500 index," *Machine Learning with Applications*, 22, 100750.

O'Neil, W. J. and C. Ryan (2002): *How to make money in stocks: A winning system in good times or bad*, McGraw-Hill New York.

Pafka, S. and I. Kondor (2004): "Estimated correlation matrices and portfolio optimization," *Physica A: statistical mechanics and its applications*, 343, 623–634.

Park, S. and D. P. O'Leary (2010): "Portfolio Selection Using Tikhonov Filtering to Estimate the Covariance Matrix," *SIAM Journal on Financial Mathematics*, 1, 932–961.

Salas-Molina, F. and J. Nin (2026): "Risk mitigation through noise reduction in hierarchical portfolio selection," *Expert Systems with Applications*, 299, 130304.

Sharpe, W. F. (1966): "Mutual Fund Performance," *The Journal of Business*, 39, 119–138.

Stoyanov, S. V., S. T. Rachev, B. Racheva-Yotova, and F. J. Fabozzi (2011): "Fat-Tailed Models for Risk Estimation," *The Journal of Portfolio Management*, 37, 107–117.

Taufemback, C. G. (2023): "Asymptotic Behavior of Temporal Aggregation in Mixed-Frequency Datasets," *Oxford Bulletin of Economics and Statistics*, 85, 894–909.

Zhifeng, D. and W. Fei (2019): "Sparse and robust mean-variance portfolio optimization problems," *Physica A: Statistical Mechanics and its Applications*, 523, 1371–1378.