

REVERSE STRESS TESTING CORRELATIONS: A REDUCED-FACTOR MODEL FRAMEWORK FOR ANALYZING TRACKING ERROR

Putthipong DARAJUNPITUK

Faculty of Commerce and Accountancy, Chulalongkorn University, Thailand;

putthipong.da@gmail.com

ARTICLE HISTORY

Received: 19 May 2026

Revised: 23 May 2026

Published: 26 June 2026

ABSTRACT

This research proposes a reduced-factor model framework for performing Reverse Stress Testing (RST) to identify the most probable asset correlation scenario that pushes a portfolio's tracking error beyond a pre-defined threshold. While the existing stress testing framework links asset correlations to economic drivers, automatically including all directly linked drivers leads to high computational time and model complexity. To address this, this study first constructs a "full model" using all directly linked drivers to serve as a baseline. It then constructs a "reduced model" by evaluating all possible combinations of economic drivers for specific target subset sizes and selecting the combination that minimizes the deviation between the stress scenarios identified by the full and reduced models. The proposed framework is empirically tested using data from the Thai equity market. The results show that retaining a higher number of economic drivers in the reduced model does not inherently result in higher accuracy. Furthermore, the optimal reduced model achieves significant operational computational time savings. Finally, the economic drivers that can be safely excluded in this framework cannot be deduced simply by observing the sector weight differences between the managed portfolio and the benchmark.

Keywords: Reverse Stress Testing, Tracking Error, Asset Correlations, Reduced-Factor Model

CITATION INFORMATION: Putthipong, D. (2026). Reverse Stress Testing Correlations: A Reduced-Factor Model Framework for Analyzing Tracking Error. *Procedia of Multidisciplinary Research*, 4(6), 13.

INTRODUCTION

For portfolio managers tasked with managing a fund against a benchmark, the risk is primarily quantified by tracking error. Tracking error is a risk measure that quantifies the difference between a managed portfolio's returns and its benchmark. Following Satchell and Hwang (2001), the ex-ante tracking error at time t is defined as:

$$TE_t = \sqrt{\Delta \mathbf{w}_t^T \boldsymbol{\Sigma}_t \Delta \mathbf{w}_t}$$

Where $\Delta \mathbf{w}_t$ is the vector of asset weight differences between the managed portfolio and the benchmark, and $\boldsymbol{\Sigma}_t$ is the covariance matrix of asset returns. To isolate the impact of correlation risk, the covariance matrix is decomposed into $\boldsymbol{\Sigma}_t = \mathbf{D}_t \cdot \mathbf{C}_t \cdot \mathbf{D}_t$, where $\mathbf{D}_t$ is the diagonal matrix of asset return volatilities and $\mathbf{C}_t$ is the correlation matrix of asset returns.

This decomposition reveals that tracking error is driven by asset weight differences between the managed portfolio and benchmark, asset return volatilities, and the correlation matrix of the asset returns. However, among these factors, correlation is particularly critical in finance. It is widely known that correlations are time-varying and respond strongly to specific events, which can cause assets that are normally uncorrelated to suddenly move together (Longin & Solnik, 2001; Ang & Bekaert, 2002). This means that a sudden shift in correlations can cause a sudden increase in tracking error. Therefore, it is necessary to understand and anticipate the impact of correlation shifts. To address this, correlation stress testing is required to simulate extreme correlation scenarios to quantify their potential impact on tracking error.

To simulate these scenarios, this study uses the framework established by Packham and Woebbecking (2022), which models the correlation matrix using risk factor coefficients that correspond to the selected economic drivers. In this framework, correlation scenarios are generated by varying these coefficients. The choice of economic drivers consists of sector and geographic indices that are directly linked to the assets held in the managed portfolio and its benchmark. For example, if the managed portfolio or its benchmark includes a German IT company, the model will include returns of the Europe and IT sector price indices as the economic drivers.

By linking correlations to economic drivers, the framework allows for reverse stress testing. Unlike traditional stress testing that applies scenarios that might represent unrealistic correlation scenarios, reverse stress testing identifies the most probable correlation scenario that pushes the tracking error to exceed a pre-defined threshold. This approach allows risk managers to identify plausible scenarios that are critical to the managed portfolio.

However, the current approach, referred to as the “full model” in this study, uses all sector and geographic indices that are directly linked to the assets held in the managed portfolio and its benchmark. While it is theoretically robust, including all economic drivers leads to unnecessarily high computational time and model complexity. To address this, this study proposes a “reduced model”. The goal is to construct a model that closely replicates the stress correlation scenario of the full model using fewer economic drivers.

Therefore, this study aims to propose a process to identify the economic driver's composition in the reduced model that closely replicates the full model's stress scenarios while cutting computational complexity.

LITERATURE REVIEWS

Correlation Factor Models in Stress Testing

The foundation of this study relies on the correlation stress testing framework developed by Packham and Woebbecking (2022). Their approach is the modeling of the correlation matrix as a parametric function of risk factor coefficients, which corresponds to selected economic drivers consisting of sector and geographic indices that are directly linked to the portfolio's assets. By estimating these coefficients such that the modelled correlation matrix is approximately close to

the empirical correlation matrix, the framework links asset correlations and economic drivers directly.

Furthermore, the framework extends to reverse stress testing by identifying extreme but plausible correlation stress scenarios. To ensure plausibility, a probability distribution is fitted to the historical risk factor coefficients to define a region of likely outcomes. The reverse stress test then identifies a plausible combination of coefficients that represents a severe correlation scenario specifically for a portfolio's value-at-risk (VaR).

Machine Learning in Financial Stress Modeling

Beyond parametric correlation models, the broader field of financial stress testing has increasingly integrated machine learning (ML) to capture complex, non-linear market dynamics that traditional models often miss. Recent studies demonstrate that tree-based algorithms, such as Random Forests, significantly outperform standard linear approaches in forecasting severe market stress (Aldasoro et al., 2025) and projecting vulnerabilities under adverse scenarios (Moffo, 2024).

However, while these advanced ML models offer superior predictive accuracy for tail risks, they rely on high-dimensional datasets that introduce severe computational strain and complexity. Consequently, whether employing modern ML algorithms or parametric structural frameworks, the need for mathematically rigorous dimension-reduction techniques remains important. Establishing an optimal, reduced-factor model by stripping away structural noise is a critical requirement to ensure that any intensive stress testing methodology remains computationally viable.

Research Gaps and Optimal Driver Selection

While the correlation stress testing framework developed by Packham and Woebbecking (2022) provides an approach for modeling correlation risk, several research gaps remain. First, this framework has not been applied to tracking error. The existing study has focused primarily on value-at-risk. Although it is effective for assessing total loss, it does not account for the specific needs of risk managers who seek to manage the deviation of managed portfolio returns from a benchmark.

Second, the current framework lacks a secondary filter to balance model parsimony with accuracy. Regardless of the risk measure, the current approach automatically includes all sector and geographic indices that are directly linked to the assets involved. However, it does not evaluate whether certain drivers can be safely excluded. As a result, the model includes economic drivers that increase computational time and model complexity without meaningfully altering the resulting stress scenarios.

Identifying which specific economic drivers to exclude requires finding an optimal balance: maximizing computational efficiency without significantly altering the resulting stress scenarios. A logical assumption within this framework is that drivers linked to sectors with near-zero weight differences are the least relevant. Therefore, to address the gap in optimal driver selection, the following hypothesis is derived for this study:

H1: Economic drivers representing sectors with near-zero weight differences between the managed portfolio and the benchmark can be safely excluded without significantly deviating from the full model's stress correlation scenario.

RESEARCH METHODOLOGY

To evaluate the proposed framework empirically, this study focuses on the Thai equity market. While the theoretical tracking error framework applies broadly to any asset class, this empirical analysis specifically examines stocks, as the benchmark utilized is the SET High Dividend (SETHD) Index and the managed portfolio evaluated against this benchmark is the 1DIV ETF, with both weight allocations fixed as of 30 June 2025 to isolate the impact of market dynamics from active portfolio reallocation.

Furthermore, because only the top 10 constituent weights of the benchmark are publicly reported, the SETHD weights are constructed using the top 10 constituents from the SETHD's June monthly report, with the remaining weight distributed equally among the other 20 constituents. The dataset comprises daily logarithmic returns for the 30 underlying stocks and the 9 corresponding SET sector price indices, spanning from December 2022 to June 2025. The 9 selected economic drivers are the sector price indices for Banking (BANK), Commerce (COMM), Construction Materials (CONMAT), Energy and Utilities (ENERG), Electronic Components (ETRON), Finance and Securities (FIN), Food and Beverage (FOOD), Property Development (PROP), and Transportation and Logistics (TRANS). Geographic indices are excluded because all underlying stocks operate exclusively within the Thai market, meaning a country level factor would apply uniformly across all pairs and offer no differentiation in the correlation matrix.

The methodology is structured across four primary phases: Correlation Factor Model, Risk Factor Distribution Modeling, Reverse Stress Testing (RST), and the Reduced Model Selection Process.

Correlation Factor Model

The core of this methodology relies on establishing a functional link between the selected economic drivers and the correlation matrix of stock returns. Consider a portfolio of N stocks, indexed by $i, j \in \{1, \dots, N\}$ and K economic drivers, indexed by $k \in \{1, \dots, K\}$. Following Packham and Woebbecking (2022), the correlation between stock i and stock j , denoted as $\hat{c}_{i,j}(\theta)$ is defined using a hyperbolic tangent function to ensure values remain within $[-1, 1]$. These pairwise correlations populate the full $N \times N$ modelled correlation matrix, $\hat{\mathbf{C}}(\theta)$.

The modelled correlation is driven by a vector of risk factor coefficients (θ). This vector captures a base market correlation shared by all stocks, alongside "intra-factor" and "inter-factor" coefficients for each economic driver k that are activated depending on whether both stocks, or only one stock, are exposed to economic driver k .

To determine these specific exposures, this study employs the Bayesian Variable Selection (BVS) method, which builds upon the Bayesian linear model specifications established by Fahrmeir et al. (2013) and the stochastic search variable selection developed by George and McCulloch (1997). This method determines whether a given economic driver k is a statistically relevant factor in explaining a specific stock's returns; if it is, the stock is considered exposed.

Once the exposures of all stocks to the economic drivers are mapped, the risk factor coefficients are calibrated to match historical market conditions. This is performed daily via a rolling window approach, calibrating the coefficients by minimizing the sum of squared errors (SSE) between the modelled correlation values and the empirical correlation values calculated from historical stock returns. This calibration process yields a continuous time series of historical risk factor coefficients, $\{\hat{\theta}_t\}_{t=1}^T$, where t represents a specific trading day and T is the total number of days in the observation period.

Risk Factor Distribution Modeling

To perform reverse stress testing, it is necessary to establish the statistical likelihood of these risk factor coefficient combinations. This study fits a Multivariate Normal Inverse Gaussian (NIG) distribution to the historical time series of coefficients $\{\hat{\theta}_t\}_{t=1}^T$. The parameters of this distribution, denoted as $\hat{\Phi}$, are estimated using the Expectation-Maximization (EM) algorithm described in McNeil et al. (2005).

Reverse Stress Testing (RST)

The objective of the RST framework is to identify the most probable set of risk factor coefficients that pushes the portfolio's tracking error beyond a pre-defined threshold. This is formulated as a constrained optimization problem. Specifically, the model seeks to maximize the likelihood of the coefficients, which is quantified by the probability density function of the estimated NIG distribution, while being subject to the tracking error threshold constraint and the requirement

that the resulting correlation matrix remains positive semi-definite. The optimization problem is expressed as:

$$\begin{aligned} \theta^* &= \underset{\theta}{\operatorname{argmax}} f_{\text{NIG}}(\theta \mid \hat{\Phi}) \\ \text{s. t. } \quad \text{TE}(\theta) &= \sqrt{\Delta \mathbf{w}^T \mathbf{D} \hat{\mathbf{C}}(\theta) \mathbf{D} \Delta \mathbf{w}} \geq X \\ \hat{\mathbf{C}}(\theta) &\succeq 0 \end{aligned}$$

Where θ^* represents the most probable stress scenario coefficients, f_{NIG} is the probability density function of the multivariate NIG distribution parameterized by $\hat{\Phi}$, $\Delta \mathbf{w}$ is the vector of stock weight differences, $\mathbf{D}$ is the diagonal matrix of stock return volatilities on the latest date, $\hat{\mathbf{C}}(\theta)$ is the modelled correlation matrix, and X is the pre-defined tracking error threshold.

By solving this optimization, the model identifies the stress correlation matrix $\hat{\mathbf{C}}(\theta^*)$, representing the most probable correlation scenario under which the tracking error exceeds the threshold.

Reduced Model Selection Process

The final phase identifies the optimal combination of economic drivers that closely replicates the full model's stress scenario while reducing the number of economic drivers used. For a predetermined target number of drivers, denoted as m , the selection process proceeds as follows: First, the stress correlation matrix of the full model ($\hat{\mathbf{C}}_{\text{Full}}^*$) is computed using all K drivers to serve as the baseline scenario. Second, all possible combinations of economic drivers of size m are generated. For each combination, the RST is performed independently to identify its specific stress correlation matrix, generating a set of matrices $\{\hat{\mathbf{C}}_{\text{Reduced},1}^*, \hat{\mathbf{C}}_{\text{Reduced},2}^*, \dots\}$.

Finally, the deviation between the baseline full model and each reduced model is quantified using the Frobenius distance, which provides a scalar measure of the total element-wise difference between the two matrices. The framework selects the specific driver combination that minimizes this deviation:

$$\underset{y}{\operatorname{argmin}} \|\hat{\mathbf{C}}_{\text{Full}}^* - \hat{\mathbf{C}}_{\text{Reduced},y}^*\|_F$$

The economic drivers omitted from the driver combination that minimizes the deviation are classified as the "excluded drivers." By iterating this process across decreasing target sizes of m , the model identifies the specific subset of drivers that maximizes computational efficiency without suffering a significant loss of accuracy compared to the full model.

RESEARCH RESULTS

Model Parsimony and Optimal Driver Selection

To determine the optimal reduced model, the full model was evaluated for the 1DIV ETF portfolio against the SETHD Index. This evaluation was conducted on a testing date of 30 June 2025, with the tracking error threshold set at the 0.95 quantile of the empirical tracking error observed from 2 January 2024 to 30 June 2025.

The selection process iteratively reduced the target number of economic drivers (m) from 8 down to 4. For each m , all possible combinations of economic drivers were evaluated, and the combination that produced the stress correlation matrix with the least Frobenius distance from the full model was identified. Figure 1 illustrates the minimum Frobenius distance achieved at each target number of economic drivers.




Figure 1 Minimum Frobenius Distance of Stress Correlation Matrix $\hat{\mathcal{C}}(\theta^*)$ from the Full Model
 Note: The vertical axis represents the Frobenius distance, and the horizontal axis represents the target number of economic drivers (m). The full model utilizes 9 economic drivers, making Frobenius distance value at $m = 9$ a guaranteed zero.

The results reveal a non-linear relationship between the number of economic drivers and model accuracy. Retaining a higher number of economic drivers in the reduced model does not inherently result in a better replication of the full model stress correlation matrix. As shown in Figure 1, when the model was reduced by only one or two drivers ($m = 8$ and $m = 7$), the minimum Frobenius distance recorded was approximately 1.16.

However, reducing the target number of drivers to $m = 6$ resulted in a decrease in the Frobenius distance to approximately 1.04. This indicates that reduced models that use 7 or 8 economic drivers suffer from parameter over-specification, introducing structural noise and inefficiencies. By contrast, a reduced model that uses exactly 6 economic drivers strips away this unnecessary noise, replicating the stress correlation matrix of the full model more accurately. Reducing the drivers further to $m = 5$ and $m = 4$ resulted in structural collapse, with the distance increasing to 1.40 and 1.50, respectively.

To visually validate the scalar measurements provided by the Frobenius distance, Figure 2 presents the stress correlation heatmaps for the full model alongside selected reduced models. For enhanced interpretability, the axes of these heatmaps map the individual underlying stocks to their corresponding sector classifications. A sector label is positioned at the start of each contiguous block of stocks belonging to that sector.




Figure 2 Visual Comparison of Stress Correlation Matrices

Note: Heatmaps illustrate the structural differences in the stress correlation matrix between the full model and the reduced models. The axes group the underlying stocks by sector. The model at $m = 6$ exhibits the closest visual and structural alignment with the full model, whereas the model reduced to $m = 5$ demonstrates significant structural collapse.

Computational Efficiency and Operational Performance

A primary motivation for proposing a reduced-factor framework is to reduce the high computational time associated with high-dimensional stress testing. To explicitly quantify the efficiency gained by this model parsimony, a performance benchmark was conducted comparing the operational CPU runtime of the full model ($m = 9$) against the optimal reduced model ($m = 6$).

This study measured the average computational time required to execute the three operational phases of the framework: the Correlation Factor Model, the Risk Factor Distribution Modeling, and the Reverse Stress Testing (RST). The Reduced Model Selection Process is excluded, as it serves as a process to identify safely excluded economic drivers rather than a daily operational requirement. To ensure robustness, the CPU time for each phase was recorded across 10 independent runs, and the results were averaged. Table 1 presents the average computational runtime (in seconds) required for both models.


Table 1 Average Operational CPU Runtime by Model Size

Operational Phase	Full Model	Optimal Reduced Model	Time Saved (%)
Correlation Factor Model	140.43 s	130.35 s	7.18%
Risk Factor Distribution Modeling	41.64 s	1.98 s	95.24%
Reverse Stress testing (RST)	2.20 s	1.55 s	29.55%
Total Operational Time	184.27 s	133.89 s	27.34%

Note: CPU runtime represents the average duration (in seconds) across 10 independent iterations on the testing date. The Reduced Model Selection Process is excluded, as it serves as a process to identify safely excluded economic drivers rather than a daily operational requirement.

As demonstrated in Table 1, utilizing the optimal reduced model yields a 27.34% overall reduction in daily computational runtime. While the Correlation Factor Model and RST optimization phases experience moderate efficiency gains, the most profound computational saving occurs during the Risk Factor Distribution Modeling phase. In this phase, the runtime drops drastically from 41.64 seconds down to just 1.98 seconds, which is a 95.24% reduction in computational runtime.

To investigate this drastic reduction and validate the efficiency of the reduced framework, the convergence behavior of the Expectation-Maximization (EM) algorithm used to fit the Multivariate Normal Inverse Gaussian (NIG) distribution was examined. Figure 3 illustrates the log-likelihood convergence trajectory for both models.

**Figure 3** Log-Likelihood Convergence of the EM Algorithm

Note: The figure illustrates the convergence trajectories when fitting the Multivariate NIG distribution parameters. Panel (a) displays the convergence behavior of the full model, while Panel (b) displays the optimal reduced model. For both panels, the vertical axis represents the log-likelihood, and the horizontal axis represents the number of iterations during the EM algorithm. Note the different numerical scales on the axes between the two panels.

The convergence data explicitly justifies the massive time reduction. As shown in Figure 3(a), the full model operates within a highly complex, high-dimensional parameter space. Consequently, the EM algorithm requires approximately 2,600 iterations to navigate this set of risk factor coefficients to achieve log-likelihood convergence. By contrast, Figure 3(b) demonstrates that the optimal reduced model strips away unnecessary structural noise. This parsimony drastically reduces the dimensionality of the distribution's parameters, allowing the EM algorithm to efficiently converge in approximately 55 iterations.

This explicitly demonstrates that automatically including all directly linked economic drivers introduces severe, unnecessary computational bottlenecks. By identifying the optimal $m = 6$, the framework avoids these high-dimensional inefficiencies while closely replicating the full model's stress correlation matrix.

Evaluation of Excluded Economic Drivers

To understand the structural dynamics of the reduced models, it is necessary to track the specific economic drivers excluded at each target number of economic drivers (m). Table 2 details the sequential exclusion of sectors as the model size decreases.

Table 2 Excluded Economic Drivers by Target Number of Economic Drivers (m)

Target Number of Economic Drivers (m)	Excluded Economic Drivers
8	ETRON
7	COMM, ETRON
6	COMM, ETRON, PROP
5	ENERG, ETRON, FOOD, PROP
4	COMM, ETRON, FOOD, PROP, TRANS

Note: Table 2 illustrates the specific sectors dropped from the full model to minimize the Frobenius distance at each target number of economic drivers. ETRON is consistently excluded across all iterations, while FOOD and ENERG are dropped precisely at $m = 5$, triggering the onset of structural collapse.

As established in Figure 1, the optimal balance of parsimony and accuracy is achieved at $m = 6$, where COMM, ETRON, and PROP are safely excluded as shown in Table 2. However, reducing the model by just one additional driver to $m = 5$ triggers the onset of a structural collapse. Table 2 reveals that this critical breaking point coincides directly with the removal of FOOD and ENERG.

To address Hypothesis 1, we investigated whether the safely excluded sectors (COMM, ETRON, and PROP) exhibited near-zero weight differences, and conversely, whether the sectors that got excluded at $m = 5$ (FOOD and ENERG) carried high weight differences. Figure 4 illustrates the absolute sector weight differences between the managed portfolio and the benchmark.




Figure 4 Absolute Sector Weight Differences between the managed portfolio (1DIV ETF) and the benchmark (SETHD)

Note: The absolute sector weight difference is calculated by aggregating the individual weight differences of all underlying stocks within a specific sector. A logical assumption is that near-zero sectors would be excluded first, but the selection process contradicts this.

The data in Figure 4 explicitly rejects Hypothesis 1. The sectors exhibiting the absolute lowest, near-zero weight differences in the entire portfolio are FOOD and ENERG. Despite their negligible weights, both are strictly necessary to maintain the model's structural integrity at $m =$

6. Conversely, PROP and COMM exhibit the third and fourth highest absolute sector weight differences, yet they are safely excluded without significantly altering the stress scenario. Given that weight differences fail as reliable predictors, the analysis was expanded to investigate whether sector volatility or the full model's stress correlation matrix can explain why certain sectors are strictly necessary while others can be safely excluded. Figure 5 presents the average stock volatility for each sector on the testing date that is used during the reverse stress test.



Figure 5 Average Sector Volatility on the Testing Date

Note: Volatility is calculated as the simple average of the standard deviations of the underlying constituent stocks within each respective sector on June 30, 2025.

An analysis of sector volatility fails to explain the exclusion dynamics. The sectors that trigger the structural collapse when removed, which are FOOD and ENER exhibit average volatilities of approximately 0.023 and 0.021, respectively. These values are indistinguishable from the safely excluded COMM and PROP sectors. In addition, the ETRON sector exhibits one of the highest volatilities in the portfolio, yet the algorithm safely targets it for exclusion across every single iteration from $m = 8$ down to $m = 4$ as shown in Table 2. This confirms that average sector volatility does not reliably predict which economic drivers can be safely excluded.

Finally, an evaluation of the full model's stress correlation matrix in Figure 2 was conducted to determine if specific visible correlation patterns justify an economic driver's exclusion. Visual inspection of the stress matrix reveals no definitive pattern. For example, FOOD lacks high-magnitude correlations with other sectors, nor does it display overwhelming intra-sector correlation. Conversely, the other strictly retained sector, ENER, displays highly pronounced intra-sector correlation. By comparison, the safely excluded PROP sector demonstrates intra-sector correlation that is visually similar to ENER. Because ENER is necessary while PROP with similar correlation patterns is safely excluded, visual inspection of the stress matrix cannot predict which economic drivers can be safely excluded.

The inability to predict which economic drivers can be safely excluded highlights the profound complexity of the RST optimization. Relying on weight differences, sector volatility, or visual correlation patterns is insufficient. The ability to accurately replicate the full model's stress scenario does not depend on any single, observable metric. Instead, it is driven by how stock weight differences, volatilities, and correlations interact simultaneously within the tracking error calculation. Retaining sectors like FOOD and ENER is mathematically necessary because their combined interactions are required by the optimization algorithm to closely replicate the full model's stress correlation matrix. Therefore, simple assumptions based on visual observations are unreliable. Risk managers must therefore rely on a selection process to successfully identify which economic drivers can be safely excluded. Specifically, they must minimize the Frobenius distance between the full and reduced stress correlation matrices for all possible combinations.

DISCUSSION & CONCLUSION

This study proposed a reduced-factor model framework to address the high computational complexity associated with reverse stress testing correlations for tracking error. By iteratively evaluating subset combinations of economic drivers, the framework successfully identifies reduced models that accurately replicate the extreme correlation scenarios of the full model.

Empirical evaluation of the 1DIV ETF against the SETHD Index demonstrated a non-linear relationship between the number of retained drivers and model accuracy. Reducing the model to exactly 6 economic drivers ($m = 6$) provided the closest replication of the full model's stress correlation matrix by stripping away unnecessary structural noise. Conversely, reduced models that use 7 or 8 economic drivers suffered from parameter over-specification, while subsets smaller than 6 resulted in structural collapse.

Crucially, the optimal reduced model effectively resolves the primary challenge of computational complexity. By avoiding the high-dimensional parameter spaces that cause severe bottlenecks during the Risk Factor Distribution Modeling phase, the $m = 6$ framework reduced overall daily operational CPU runtime by over 27%, achieving an approximately 95% reduction in the distribution fitting phase alone.

Furthermore, an analysis of the specific drivers excluded at $m = 6$ rejected the assumption that sector weight differences alone can determine driver relevance. Sectors with near-zero weight differences (FOOD and ENERG) were necessary to the reverse stress test, whereas sectors with high weight differences (COMM and PROP) were safely excluded. This confirms that while sector weights are a fundamental component of tracking error, weight differences alone are insufficient for identifying which economic drivers can be safely reduced in a reverse stress testing correlations framework.

It is important to acknowledge the empirical limitations of this study regarding its scope. The analysis is currently constrained to a single market which is the Thai equity market. While this provides a robust initial case study, the generalizability of the observed non-linear relationship and the specific optimal reduction point at $m = 6$ remains unverified for other geographic regions or asset classes. The optimal driver count of $m = 6$ should therefore be interpreted as a portfolio-specific finding rather than a universal constant.

Consequently, future research should conduct robustness checks across diverse benchmarks and managed portfolios to assess whether the behavior of the Frobenius distance and the optimal driver count remain consistent. Furthermore, because isolated metrics such as weight differences, sector volatility, or visual correlation patterns are unreliable predictors, risk managers must currently rely on selection processes to identify which drivers to exclude. Future studies should focus on identifying the underlying mathematical characteristics shared by these safely excluded drivers. If a reliable predictive characteristic can be established, the framework could pre-filter combinations by exclusively dropping drivers that exhibit this specific characteristic, leading to a significant reduction in the computational time required for the selection process itself.

REFERENCES

- Aldasoro, I., Hördahl, P., Schrimpf, A., & Zhu, X. S. (2025). *Predicting financial market stress with machine learning* (BIS Working Papers No. 1250). Bank for International Settlements.
- Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *Review of Financial Studies*, 15(4), 1137–1187.
- Fahrmeir, L., Kneib, T., Lang, S., & Marx, B. (2013). *Regression: Models, Methods and Applications*. Springer.
- George, E. I., & McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica*, 7, 339–373.

- Longin, F., & Solnik, B. (2001). Extreme correlation of international equity markets. *The Journal of Finance*, 56(2), 649–676.
- McNeil, A. J., Frey, R., & Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques, and Tools*.
- Moffo, A. M. F. (2024). A machine learning approach in stress testing US bank holding companies. *International Review of Financial Analysis*, 95, 103476.
- Packham, N., & Woebeking, F. (2022). Correlation scenarios and correlation stress testing. *Journal of Economic Behavior & Organization*, 205, 55–67.
- Satchell, S., & Hwang, S. (2001). Tracking error: Ex ante versus ex post measures. *Journal of Asset Management*, 2(3), 241–246.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Conflicts of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.



Copyright: © 2026 by the authors. This is a fully open-access article distributed under the terms of the Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0).