STUDENT JOURNAL OF BUSINESS AND ECONOMICS
Vol. 1, No. 5 (2026), pp. 66–74 • ISSN 3144-0363
DOI: 10.67521/sjbe.2026.045

ORIGINAL RESEARCH

Managing Volatility of Japan's Pension Fund Using Factor-Based and Constrained Mean-Variance Optimisation

Calvin Kim

The British School in Tokyo, 1-7-57 Taishido, Setagaya-ku, Tokyo, Tokyo, 8533, Japan
Corresponding author email: calvingame0312@gmail.com

Accepted: November 22, 2026

Abstract

Japan's Government Pension Investment Fund (GPIF), managing over $2 trillion under a century-long intergenerational program, relies heavily on market-capitalization-weighted indexing that concentrates portfolio risk in few dominant market factors. This study develops a quantitative framework to improve factor-level risk diversification while adhering to trading constraints and return expectations. Using daily returns for 90 continuously traded TOPIX Core30 and Large70 constituents from January 2010 through August 2026, we apply Principal Component Analysis to the standardized return correlation matrix and employ the Marchenko-Pastur theorem ($\lambda_+ = 1.3795$) to separate four structural signal eigenportfolios from 86 random noise components which we process either through clipping or truncation. We compare Top-4 Rank-4 Truncated and Top-4 Eigenvalue-Clipped Mean-Variance Optimization (MVO) against equal-weight, sample covariance, and Ledoit-Wolf shrinkage benchmarks across the 2022–2026 out-of-sample regime. Eigenvalue clipping reduces the covariance condition number by 72.9% to 173.0 compared to sample covariance and quadruples the Effective Number of Uncorrelated Bets from 1.13 to 4.44. It achieves an annualised return of 20.31%, volatility of 15.14%, and a Sharpe ratio of 1.34, outperforming both sample covariance MVO and Ledoit-Wolf MVO in risk-adjusted returns. Although the equally weighted portfolio had a higher return and Sharpe ratio during out-of-sample testing, its ENB (1.13) suggests its inferior risk diversification. The truncated model produces a substantial increase of ENB (11.72) and high condition number ($\kappa = 6.92 \times 10^{18}$), which indicates the possibility of noise-contamination but does not confirm it. The truncation may simplify covariance matrix structure but still suffers from significant asset concentration and lower DR (1.25).

Keywords: Factor-Based Investing; Mean-Variance Optimization; Random Matrix Theory; Marchenko-Pastur Theorem; Eigenvalue Clipping; Government Pension Investment Fund; Portfolio Diversification; Intergenerational Pension Program

1. Introduction

Managing over $2 trillion in assets with a century-long intergenerational dependency program, Japan's Government Pension Investment Fund (GPIF) stands as the world's largest institutional pension scheme (GPIF 2023). It aims toward securing a long-term real return of at least 1.9% annually above Japanese nominal wage growth while minimizing portfolio variance. To achieve this, the GPIF relies on a balanced

25/25/25/25 strategic asset allocation across domestic equities, international equities, domestic bonds, and international bonds, primarily implemented via market-capitalization-weighted index tracking (such as the TOPIX for domestic equities). However, this framework suffers from a critical limitation: many assets frequently share high degrees of exposure to the same macroeconomic and market risk factors, leaving the portfolio vulnerable (Markowitz 1952; Roll 1977).

This paper proposes a multi-step quantitative framework designed to optimize factor-level risk diversification while respecting existing fund requirements. Utilizing a representative liquid universe of constituents from the TOPIX, we compute the sample covariance matrix and apply Principal Component Analysis (PCA) to extract orthogonal eigenportfolios representing independent macro-market drivers. We implement the Marchenko-Pastur theorem to discard eigenvalues below the theoretical noise threshold (Marchenko and Pastur 1967; Laloux et al. 2000). Subsequently, we formulate a constrained Mean-Variance Optimization (MVO) model to solve for optimal allocation weights across these orthogonal eigenportfolios with eigenvalues above the noise band, under linear constraints that prohibit short-selling and enforce the 1.9% annual real return target. Finally, we evaluate the approach using orthogonality loss and various performance metrics, such as off-diagonal power loss, Sharpe ratio, and various other out-of-sample assessments.

2. Mathematical Framework

2.1 Minimum-Variance Optimization

Let R in $\mathbf{R}^{n\times 90}$ denote the return matrix of the TOPIX 90 constituents across n observations. Only 90 assets survived after filtering for discarding historical delisting of firms. Let r in $\mathbf{R}^{90\times 1}$ represent the transpose of one row of the return matrix of the assets, with expected value $\mu = \mathbf{E}[r]$ in $\mathbf{R}^{90\times 1}$. For a given portfolio weight vector x in $\mathbf{R}^{90\times 1}$, the portfolio return is given by $x^\mathsf{T} r$ and the expected portfolio return is $x^\mathsf{T}\mu$. The variance of the portfolio is found as:

$$\begin{aligned}
\mathrm{Var}(x^\mathsf{T} r) &= \mathbf{E}[(x^\mathsf{T} r - x^\mathsf{T}\mu)^2] \\
&= \mathbf{E}[(x^\mathsf{T}(r-\mu))^2] \\
&= \mathbf{E}[(x^\mathsf{T}(r-\mu))(x^\mathsf{T}(r-\mu))^\mathsf{T}] \\
&= \mathbf{E}[x^\mathsf{T}(r-\mu)(r-\mu)^\mathsf{T} x] \\
&= x^\mathsf{T}\mathbf{E}[(r-\mu)(r-\mu)^\mathsf{T}]x = x^\mathsf{T}\Sigma x
\end{aligned} \tag{1}$$

where Σ in $\mathbf{R}^{90\times 90}$ is the asset covariance matrix.

The linearly constrained minimum-variance optimization problem can be formulated as the following convex quadratic program (Boyd and Vandenberghe 2004):

$$\begin{aligned}
\min_{x\in\mathbf{R}^{90}} \quad & x^\mathsf{T}\Sigma x \\
\text{subject to:} \quad & \mathbf{1}_{90}^\mathsf{T} x = 1 \\
& x_i \geq 0 \quad \text{for all } i \text{ in } \{1,\ldots,90\} \\
& x^\mathsf{T}\mu \geq r_{\text{real}} + \gamma_{\text{wage}}
\end{aligned} \tag{2}$$

where $\mathbf{1}_{90}$ is a 90×1 vector of ones that ensures all capital is exhausted, $x_i \geq 0$ imposes the no-short-selling constraint, $r_{\text{real}} = 0.019$ represents the 1.9% real return floor above nominal wage growth, and γ_{wage} denotes expected Japanese nominal wage growth.

The nominal wage growth γ_{wage} is estimated using the baseline projection provided by the Cabinet Office of Japan. It is projected at $\gamma_{\text{wage}} = 2.2\%$ under the assumption that the total factor productivity (TFP) of Japan does not significantly alter over the observation period (Cabinet Office, Government of Japan 2026). This transforms the linear return constraint into:

$$\begin{aligned}
& x^\mathsf{T}\mu \geq r_{\text{real}} + \gamma_{\text{wage}} \\
& x^\mathsf{T}\mu \geq 0.019 + 0.022 = 0.041
\end{aligned} \tag{3}$$

This is an important nuance that affects the portfolio weights computed. Assuming higher nominal wage growth rates would have implied that Japan's TFP would increase substantially across the observation period. This research contends that utilizing a more conservative and pessimistic estimate tests the robustness of the allocation strategy should long-term growth policies fall short of historical expectations.

2.2 Covariance Estimation and Standardization

The covariance matrix of the return matrix R is defined using the centered data matrix. Let X in $\mathbf{R}^{n\times 90}$ denote the centered return matrix:

$$X = R - \mathbf{1}_n \mu^{\mathsf{T}} \tag{4}$$

where $\mathbf{1}_n$ is an $n \times 1$ vector of ones and μ^{T} is the 1×90 mean return vector. The sample covariance matrix Σ is calculated as:

$$\Sigma = \frac{1}{n-1} X^{\mathsf{T}} X \tag{5}$$

To prevent assets exhibiting high variance from being favored, the covariance structure is standardized. Let $D = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_{90})$ be the diagonal matrix of the standard deviations of individual assets. The standardized covariance matrix corresponds to the sample correlation matrix C in $\mathbf{R}^{90\times 90}$ (Goueguel 2021):

$$C = D^{-1} \Sigma D^{-1} \tag{6}$$

We can perform spectral eigendecomposition on the correlation matrix C to derive its eigenvectors V and corresponding eigenvalues placed in the diagonal matrix Λ:

$$CV = V\Lambda \tag{7}$$

2.3 Spectral Denoising via the Marchenko-Pastur Theorem

To improve the tolerance of portfolio weights to estimation errors, we attempt to regulate the condition number ratio given by $\kappa = \lambda_{\max}/\lambda_{\min}$ (Best and Grauer 1991). The theoretical upper noise threshold λ_+ derived from the Marchenko-Pastur theorem is expressed as (Marchenko and Pastur 1967; Bouchaud and Potters 2009):

$$\lambda_+ = \left(1 + \sqrt{N/n}\right)^2 \tag{8}$$

where $N = 90$ is the initial number of assets and n is the number of observations. Any eigenvalue $\lambda_i < \lambda_+$ is classified as noise and filtered out prior to optimization. The unit variance standardizes to 1 in a correlation matrix, hence is not shown.

2.4 Shrinkage of the Covariance Matrix

To mitigate sampling estimation error in the covariance structure, we implement the Ledoit-Wolf shrinkage estimator (Ledoit and Wolf 2004). The shrunk covariance matrix Σ_{LW} is defined as:

$$\Sigma_{\mathrm{LW}} = \alpha \Sigma_{\mathrm{sample}} + (1-\alpha)\Phi \tag{9}$$

where Σ_{sample} is the empirical sample covariance matrix, Φ is a structured target matrix (where we reference the scaled identity matrix), and α in $[0,1]$ is the optimal shrinkage intensity parameter that minimizes expected quadratic loss. This is used in conjunction with the long-only constraint ($x_i \geq 0$), which has been empirically confirmed to reduce sample estimation error (Jagannathan and Ma 2003).

3. Data Collection and Experimental Design

3.1 Constructing the Stock Universe

To construct a liquid and representative benchmark for the Japanese equity market, we extract constituent metadata from the Tokyo Stock Exchange (TSE) and filter for companies categorized within the TOPIX Core30 and TOPIX Large70 indices, generating an initial list of 100 TOPIX tickers. Four-digit security

tickers are formatted with the Tokyo Stock Exchange suffix (.T), and daily adjusted closing prices are retrieved from January 1, 2010 through August 8, 2026.

To prevent artificial zero-return distortion from post-2010 initial public offerings, missing prices are never backfilled and only forward-filled. Brief stops of trading are forward-filled up to a maximum of five trading days, and assets lacking price histories above this window are excluded. This filtering yields $N = 90$ continuously traded large-cap constituents. This change is adjusted for within the calculations. Daily arithmetic returns are computed as $r_{t,i} = (P_{t,i} - P_{t-1,i})/P_{t-1,i}$ and annualized assuming 252 trading days per year.

3.2 In-Sample Calibration and Out-of-Sample Split

The return dataset is partitioned into two regimes:

- **In-Sample Window (January 1, 2010 – December 31, 2021):** Comprising $T_{\text{in}} = 2{,}955$ trading days, this window is used exclusively to estimate annualized expected returns μ_{ann}, the sample correlation matrix C_{in}, the spectral eigendecomposition $C_{\text{in}}V = V\Lambda$, the Marchenko-Pastur cutoff λ_+, and the Ledoit-Wolf shrinkage intensity α.
- **Out-of-Sample Window (January 1, 2022 – August 8, 2026):** Evaluates in-sample portfolio weights across a regime marked by fundamental economic shifts like the normalisation of monetary policy, rising global interest rates, and pronounced yen volatility. This period is used to assess the model on un-seen data.

3.3 Spectral Reconstruction and Benchmark Specifications

For $N = 90$ assets and $T_{\text{in}} = 2{,}955$ observations, Equation (8) yields a theoretical Marchenko-Pastur upper noise threshold of $\lambda_+ = 1.3795$. Again, this is adjusted for the change in asset number to 90. Exactly $k = 4$ eigenvalues exceed this cutoff ($\lambda_1 = 40.55$, $\lambda_2 = 3.87$, $\lambda_3 = 2.73$, and $\lambda_4 = 1.57$), whereas the fifth eigenvalue ($\lambda_5 = 1.36$) falls immediately below λ_+. Retaining the leading eigenvectors V_4 in $\mathbf{R}^{90\times 4}$ and eigenvalues $\Lambda_4 = \text{diag}(\lambda_1, \ldots, \lambda_4)$, we construct two spectral estimators and evaluate them alongside three benchmark specifications under the identical Sequential Least Squares Programming (SLSQP) constraints in Equations (2)–(3):

1. **Top-4 Rank-4 Truncated MVO:** Sets all 86 sub-threshold noise eigenvalues to zero ($C_{\text{rank4}} = V_4\Lambda_4 V_4^{\top}$), rescaling diagonal entries to unity before mapping back to annualized covariance space via $\Sigma_{\text{rank4}} = 252\, D C_{\text{rank4}} D$.
2. **Top-4 Eigenvalue-Clipped MVO:** Following Random Matrix Theory literature (Laloux et al. 2000; Bouchaud and Potters 2009), we replace all eigenvalues below the noise threshold with their mean $\lambda_{\text{noise}} = (1/86)\sum_{i=5}^{90} \lambda_i$, maintaining the matrix trace and strict positive definiteness to control the condition number.
3. **Benchmark Models:** (i) a $1/N$ Equal-Weight portfolio ($x_i = 1/90$), (ii) unregularized Sample Covariance MVO (Σ_{sample}), and (iii) Ledoit-Wolf Shrinkage MVO (Σ_{LW}).

4. Empirical Results and Discussion

4.1 Evaluation Metrics

Portfolio effectiveness is evaluated by two groups of metrics: out-of-sample factor orthogonality persistence and realized risk-return efficiency. The two metrics below are used to measure level of risk diversification. Measures of risk-return efficiency are provided in table 2.

Off-Diagonal Power Loss (ODPL): We project standardized out-of-sample returns Z_{out} onto the in-sample eigenvectors V_4 which yields realized factor returns $F_{\text{out}} = Z_{\text{out}}V_4$. The correlation matrix $\Omega_{\text{out}} = \text{Corr}(F_{\text{out}})$ is the identity matrix under in-sample testing period. However, as regime shifts, it induces off-diagonal leakage out-of-sample, quantified via the square of the Frobenius norm ratio, calculating absolute distance:

$$\text{ODPL}_{\text{corr}} = \frac{\|\Omega_{\text{out}} - \text{diag}(\Omega_{\text{out}})\|_F^2}{\|\Omega_{\text{out}}\|_F^2} = \frac{\sum_{i\neq j}(\Omega_{\text{out}})_{ij}^2}{\sum_{i=1}^{4}\sum_{j=1}^{4}(\Omega_{\text{out}})_{ij}^2} \tag{10}$$

where $\text{ODPL}_{\text{corr}} \to 0\%$ indicates persistent out-of-sample orthogonality.

Concentration Metrics: To measure risk dispersion, we compute the Effective Number of Stocks $N_{\text{eff}} = \left(\sum_{i=1}^{90} x_i^2\right)^{-1}$, the Diversification Ratio $\text{DR}(x) = (x^\top \sigma_{\text{out}})/\sqrt{x^\top \Sigma_{\text{out}} x}$, and Meucci's entropy-based Effective Number of Uncorrelated Bets (ENB) (Meucci 2009):

$$\text{ENB}(x) = \exp\left(-\sum_{i=1}^{90} p_i \ln p_i\right), \quad \text{where } p_i = \frac{(v_{i,\text{out}}^\top x)^2 \lambda_{i,\text{out}}}{x^\top \Sigma_{\text{out}} x} \tag{11}$$

4.2 Matrix Structure and Out-of-Sample Orthogonality

In-sample eigendecomposition illustrates factor concentration across Japanese mega-cap equities. The leading principal component ($\lambda_1 = 40.55$), likely representing broad market beta, accounts for 45.06% of total variance considering top 90 PC. Together, the four signal components clearing the Marchenko-Pastur cutoff ($\lambda_+ = 1.3795$) capture 54.13% of total variance ($\lambda_2 = 3.87$ [4.30%], $\lambda_3 = 2.73$ [3.03%], and $\lambda_4 = 1.57$ [1.74%]), whereas the remaining 86 eigenvalues ($\lambda_5 = 1.36$ down to λ_{90}) are considered noise as it falls below λ_+.

Table 1 reports the out-of-sample realized correlation matrix and orthogonality loss diagnostics. Across 2022–2026, the leading four eigenportfolios exhibit an $\text{ODPL}_{\text{corr}}$ of 11.02% and a mean absolute off-diagonal correlation of 0.187, retaining nearly 89% of their diagonal variance. Notably, PC_1 and PC_4 remain virtually orthogonal ($\rho_{14} = 0.04$), while more positive correlations ($\rho_{12} = 0.29$, $\rho_{23} = \rho_{34} = 0.23$) reflect coupling between various principal components. By contrast, noise-band eigenvectors (PC_5–PC_8) exhibit an $\text{ODPL}_{\text{corr}}$ of 5.95% (mean $|\rho_{ij}| = 0.129$) because projections onto random noise remain unrelated out-of-sample.

Table 1. Out-of-Sample Realized Correlation Matrix (Ω_{out}) and Orthogonality Loss Comparison (2022–2026)

Principal Component	**PC_1**	**PC_2**	**PC_3**	**PC_4**
PC_1	**1.00**	0.29	0.18	0.04
PC_2	0.29	**1.00**	0.23	0.15
PC_3	0.18	0.23	**1.00**	0.23
PC_4	0.04	0.15	0.23	**1.00**

Factor Group	**In-Sample ODPL**	**Out-of-Sample ODPL**	**Mean $\lvert\rho_{ij}\rvert$**
Top-4 Signal (PC_1–PC_4)	0.00%	11.02%	0.187
Noise Bulk (PC_5–PC_8)	0.00%	5.95%	0.129

4.3 Analysis of Out-of-Sample Portfolio Returns, Risk Diversification, and Allocations

Table 2 summarizes the matrix conditioning, realized return, volatility, drawdown, and factor diversification metrics across all five portfolios over 2022–2026. We utilize a zero risk-free rate Sharpe ratio throughout our empirical observations (Sharpe 1963). The ENB is a risk diagnostic calculated in retrospect using the out-of-sample covariance matrix. Consequently, it serves as a retrospective diagnostic of realized portfolio risk diversification rather than a measure of information available to investors at the beginning of the evaluation period (2022).

We provide four empirical conclusions from the data above:

- **Reduced condition number and risk-adjusted efficiency of Eigenvalue Clipping:** Replacing the 86 below λ_+ eigenvalues with their mean lowers the covariance matrix condition number $\kappa(\Sigma)$ by 72.9% relative to Sample MVO (from 638.4 to 173.0) and by 70.8% relative to Ledoit-Wolf Shrinkage (593.2). Out-of-sample, Top-4 Eigenvalue-Clipped MVO generates a 20.31% annualized return at 15.14% volatility (Sharpe ratio 1.34), outperforming both Sample Covariance MVO (1.13) and Ledoit-Wolf Shrinkage MVO (1.13) by a Sharpe Ratio of 0.21 higher. It also outperforms Top-4 Rank-4 Truncated MVO (1.23) while exceeding the GPIF's 4.10% nominal return floor by

Table 2. Out-of-Sample Portfolio Performance and Diversification Metrics (2022–2026)

Model	$\kappa(\Sigma)$	**Ann. Ret.**	**Ann. Vol.**	**Sharpe**	**Max DD**	N_{eff}	**ENB**	**DR**
$1/N$ Equal-Weight Benchmark	638.4	27.55%	20.00%	1.38	−23.53%	90.0	1.13	1.67
Sample Covariance MVO	638.4	15.97%	14.14%	1.13	−15.06%	10.6	6.17	1.67
Ledoit-Wolf Shrinkage MVO	593.2	16.01%	14.15%	1.13	−15.03%	10.7	6.12	1.67
Top-4 Rank-4 Truncated MVO	6.92×10^{18}	23.70%	19.24%	1.23	−25.05%	1.8	11.72	1.25
Top-4 Eigenvalue-Clipped MVO	173.0	20.31%	15.14%	1.34	−16.78%	11.2	4.44	1.63

over 16 percentage points. It also generates a diversification ratio 1.63, comparably greater than the Top-4 Rank-4 Truncated MVO, although falling short of other spectral estimators. All together, it signals high portfolio performance both in reducing sample-error susceptibility while preserving risk-adjusted efficiency.

- **Non-Binding Return Floor Constraint:** Comparing unconstrained Global Minimum Variance (GMV) portfolios against return-constrained MVO solutions ($x^{\mathsf{T}}\mu \geq 0.041$) reveals identical weights for both estimators. Top-4 Eigenvalue-Clipped GMV achieves an in-sample expected return of 12.67% (16.39% sample volatility) and Top-4 Rank-4 GMV achieves 9.65% (19.65% sample volatility). This result is replicated by the out-of-sample results illustrated in table 2. Because both clear the 4.10% threshold under TFP projections, the return constraint is entirely non-binding.
- **Failure of Rank-4 Truncation:** Zeroing out all 86 noise eigenvalues ($\lambda_{i>4} = 0$) makes Σ_{rank4} numerically singular ($\kappa = 6.92 \times 10^{18}$) because one of its eigenvalues is near zero, close to the limit of computer rounding. The optimizer collapses into merely 3 active stocks, producing high out-of-sample volatility (19.24%), a low Diversification Ratio (1.25), and the deepest Maximum Drawdown (−25.05%) out of all models. Its Effective Number of Stocks (1.8) or high asset concentration indicates its physical inability to exploit inter-asset correlation, justifying its low Diversification Ratio. Notably, the model produces an out-of-sample ENB of 11.72. This apparent contradiction is a byproduct of using the out-of-sample covariance matrix, allowing the value to exceed 4, the initial number of eigenvalues beyond noise threshold. The substantial increase isn't definitive proof of noise contamination yet indicates the possibility of it. Overall, the model may simplify covariance matrix structure but at the critical expense of overall portfolio robustness and effective diversification.
- **Inadequacy of Equally Weighted Benchmark:** Although the equally weighted benchmark produces the highest annualised return (27.55%), it also suffers from the highest volatility (20.00%) out of all models evaluated. Although its Sharpe ratio (1.38) is notable, GPIF holds superior preference towards risk-aversion after the return target is met. With a Maximum DrawDown rate (−23.53%) that falls just second to that Top-4 Rank-4 Truncated MVO (25.05%), its volatility must be considered. Both the equally weighted index and the market capitalization portfolio closely parallels each other. They are exposed to systematic betas and share roughly the same ENB (1.13). The results can therefore serve as a simple comparison point for the GPIF's market capitalization approach.

5. Policy Implications for Japan's Government Pension Investment Fund

These empirical findings carry direct implications for Japan's Government Pension Investment Fund (GPIF):

- **Escaping Single-Factor Concentration:** GPIF's reliance on market-capitalization-weighted indexing leaves its domestic equity dominated by a single market mode (ENB = 1.13), exposing the fund to macroeconomic shocks (GPIF 2023). Reweighting constituents using eigenvalue-clipped correlation matrices lifts factor diversification nearly fourfold (ENB = 4.44) across 20 active mega-cap holdings without sacrificing principles of risk diversification. This indicates far greater levels of risk diversification that aligns with the risk-averse motif of the GPIF.
- **Matrix Stability and Turnover Control:** Managing over $2 trillion in assets exposes the GPIF to substantial market-impact costs during portfolio rebalancing (GPIF 2023). By lowering $\kappa(\Sigma)$ to 173.0 and rendering the return floor non-binding, eigenvalue clipping suppresses weight sensitivity to sample estimation noise, promoting lower turnover.
- **Intergenerational Sustainability Under Statutory Governance Rules:** With a realized volatility of 15.14% and limiting drawdowns to -16.78% during the 2022–2026 period, the results directly support the GPIF's century-long intergenerational sustainability horizon while clearing the conservative 4.1% nominal return floor ($r_{\text{real}} + \gamma_{\text{wage}}$) (Cabinet Office, Government of Japan 2026). Crucially, all optimizations strictly enforce $\mathbf{1}^{\top} x = 1$ and $x_i \geq 0$, the framework requires no short-selling, leverage, or derivatives to achieve the designated return targets. It has the potential to be deployed as a systematic and coherent overlay within existing statutory rules.

6. Methodological Limitations and Future Work

Several empirical and institutional limitations provide consideration when translating this framework into real-world applications:

1. **Survivorship Bias and the Immobile Covariance Matrix:** Constructing the investment universe from current TOPIX Core30 and Large70 constituents and backtesting to 2010 introduces survivorship bias by discarding historical delistings of firms. Furthermore, holding a static 2010–2021 covariance estimate fixed across 2022–2026 does not adapt to changing levels of correlation between principal components. Future work should evaluate rolling estimation windows to adjust the changing nature of correlation.
2. **Single Asset Emphasis:** Analyzing domestic equities in isolation omits cross-asset covariance with currency exposures, Japanese Government Bonds (JGBs), and global equities across the GPIF's full 25/25/25/25 strategic asset allocation. Further investigation should consider methods to achieve relatively fixed allocation while preserving characteristics of this investigation.
3. **Statutory Ownership:** Although the long-only constraint $x_i \geq 0$ distributes weights across 20 active stocks ($N_{\text{eff}} = 11.2$), an unconstrained upper bound ($x_i \leq 1$) allows individual allocations up to 14.68% (9432.T), which would exceed statutory ownership and cause market disruption. Institutions should enforce specific upper weight caps ($x_i \leq u_i$) to prevent this.
4. **Sensitivity to Expected Return Estimates (μ):** While the 4.10% return floor was non-binding in the 2010–2021 calibration window, tighter return targets can reintroduce sensitivity to sample mean estimation. Incorporating Black-Litterman equilibrium would be a natural step ahead.

Funding Sources
The authors declare that there were no funding sources related to this work.

Conflicts of Interest
The authors declare that there are no conflicts of interest related to this work.

Acknowledgements
The author acknowledges Junying Chi, Douglas Williams, and Remy Hewavitharana for peer-reviewing and providing substantial feedback to the paper.

AI Usage Acknowledgment

During the preparation of this work, the author used Manus AI to verify calculations from author-supplied data, improve language readability, produce LaTeX typesetting, and assist with the design of tables and figures. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

Best, M. J., and Grauer, R. R. 1991. "On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in the Means of Individual Assets." *The Review of Financial Studies* 4 (2): 315–342.

Bouchaud, J.-P., and Potters, M. 2009. "Financial Applications of Random Matrix Theory: A Short Review." In *Handbook of Financial Time Series*, 849–865. Berlin, Heidelberg: Springer.

Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge: Cambridge University Press.

Cabinet Office, Government of Japan. 2026. *Economic and Fiscal Projections for Medium to Long Term Analysis*. https://www5.cao.go.jp/keizai3/projection-e/projection2607.pdf.

Goueguel, C. L. 2021. "Covariance vs Correlation in PCA: What's the Difference?" Christian L. Goueguel Personal Blog, March 15, 2021. https://christiangoueguel.com/blog/posts/post9/.

Government Pension Investment Fund (GPIF). 2023. *Annual Report Fiscal Year 2022*. Tokyo, Japan: Government Pension Investment Fund.

Jagannathan, R., and Ma, T. 2003. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *The Journal of Finance* 58 (4): 1651–1684.

Laloux, L., Cizeau, P., Bouchaud, J.-P., and Potters, M. 2000. "Random Matrix Theory and Financial Correlations." *International Journal of Theoretical and Applied Finance* 3 (3): 391–397.

Ledoit, O., and Wolf, M. 2004. "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices." *Journal of Multivariate Analysis* 88 (2): 365–411.

Marchenko, V. A., and Pastur, L. A. 1967. "Distribution of Eigenvalues for Some Sets of Random Matrices." *Mathematics of the USSR-Sbornik* 1 (4): 457–483.

Markowitz, H. 1952. "Portfolio Selection." *The Journal of Finance* 7 (1): 77–91.

Meucci, A. 2009. "Managing Diversification." *Risk* 22 (5): 74–79.

Roll, R. 1977. "A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory." *Journal of Financial Economics* 4 (2): 129–176.

Sharpe, W. F. 1963. "A Simplified Model for Portfolio Analysis." *Management Science* 9 (2): 277–293.