
PHYSICS-INFORMED SINGULAR-VALUE LEARNING FOR CROSS-COVARIANCES FORECASTING IN FINANCIAL MARKETS

Efstratios Manolakis

Università di Catania, Dipartimento di Fisica e Astronomia “Ettore Majorana”
Catania, Italy
efstratios.manolakis@phd.unict.it

Christian Bongiorno

Université Paris-Saclay, CentraleSupélec,
Mathématiques et Informatique pour la Complexité et les Systèmes,
91190, Gif-sur-Yvette, France
christian.bongiorno@centralesupelec.fr

Rosario N. Mantegna

Università degli Studi di Palermo, Dipartimento di Fisica e Chimica "Emilio Segrè"
Palermo, Italy
Complexity Science Hub
Vienna, Austria
rosario.mantegna@unipa.it

ABSTRACT

A new wave of work on covariance cleaning and nonlinear shrinkage has delivered asymptotically optimal analytical solutions for large covariance matrices. Building on this progress, these ideas have been generalized to empirical cross-covariance matrices, whose singular value decomposition identifies canonical comovement modes between two asset sets, with singular values quantifying the strength of each mode and providing natural targets for shrinkage. Existing analytical cross-covariance cleaners are derived under strong stationarity and large-sample assumptions, and they typically rely on mesoscopic regularity conditions such as bounded spectra; macroscopic common modes (e.g., a global market factor) violate these conditions. When applied to real equity returns, where dependence structures drift over time and global modes are prominent, we find that these theoretically optimal formulas do not translate into robust out-of-sample performance. We address this gap by designing a random-matrix-inspired neural architecture that operates in the empirical singular-vector basis and learns a nonlinear mapping from empirical singular values to their corresponding cleaned values. By construction, the network can recover the analytical solution as a special case, yet it remains flexible enough to adapt to non-stationary dynamics and mode-driven distortions. Trained on a long history of equity returns, the proposed method achieves a more favorable bias-variance trade-off than purely analytical cleaners and delivers systematically lower out-of-sample cross-covariance prediction errors. Our results demonstrate that combining random-matrix theory with machine learning makes asymptotic theories practically effective in realistic time-varying markets.

Keywords Spectral Tokenization · Random Matrix Theory · Equivariant Neural Networks · Nonlinear Shrinkage · Non-Stationary Dependence

1 Introduction

Empirical cross-covariance matrices provide a compact way to summarize how two groups of variables move together. Consider two zero-mean random vectors, $\mathbf{X} \in \mathbb{R}^{n_x}$ and $\mathbf{Y} \in \mathbb{R}^{n_y}$, observed over a time window of length Δt_{in} . Their cross-covariance matrix Σ_{XY} contains the expected products between components of \mathbf{X} and \mathbf{Y} , and can therefore be interpreted as a linear measure of co-movement. It indicates which fluctuations in one set of variables tend to be associated with fluctuations in the other. In empirical applications, Σ_{XY} is not known and must be estimated from the Δt_{in} paired observations available in the window.

A fundamental challenge arises when the number of variables is not small compared to the window length. In such high-dimensional regimes, the empirical estimator is heavily affected by sampling noise where apparent relationships may emerge purely by chance, and the resulting estimate can behave unreliably when evaluated Out-Of-Sample (OOS) [1]. For standard covariance matrices, this issue has been studied extensively, and a large body of work has developed principled denoising strategies, often based on spectral methods and Random Matrix Theory (RMT), to separate stable structure from noise [2]. The cross-covariance setting is less straightforward. When $n_x \neq n_y$, the matrix $\Sigma_{XY} \in \mathbb{R}^{n_x \times n_y}$ is rectangular, so the usual eigenvalue-based tools do not apply directly. Instead, one must work with a singular value decomposition. While singular-value filtering provides a natural analogue of covariance cleaning, the corresponding theoretical framework for choosing optimal shrinkage rules is comparatively less mature, and its effectiveness, in realistic time-varying systems, remains an open question.

Related denoising results exist for additive-noise models [3, 4], but sample (cross-)covariances arise from products of observations and therefore call for a multiplicative-noise cleaning framework. To the best of our knowledge, the only available analytical characterization of an asymptotically optimal singular-value shrinkage for empirical cross-covariance matrices under multiplicative noise is due to Benaych-Georges, Bouchaud, and Potters (BBP) [5]. In the high-dimensional asymptotic regime, they derive a Rotationally Invariant Estimator (RIE) that preserves the empirical singular vectors of $\widehat{\Sigma}_{XY}$ while shrinking its singular values so as to minimize the asymptotic Mean Squared Error (MSE) relative to the true cross-covariance Σ_{XY} . Their result is presented in two versions: a baseline setting with identity marginals ($\Sigma_{XX} = \mathbf{I}_{n_x}$ and $\Sigma_{YY} = \mathbf{I}_{n_y}$), which also covers the practically relevant case of whitened variables, and a more general setting that incorporates non-trivial marginal covariances into the optimal correction. The theory is validated through simulations and controlled benchmarks, but it does not consider real financial data and relies on strong stationarity assumptions that are rarely met in practice.

The impact of non-stationarity is already well documented in the simpler covariance case. Bongiorno, Challet and Loper [6] show that covariance-shrinkage estimators derived under stationarity assumptions can underperform naive asset-agnostic shrinkages when tested on real markets. In particular, they report that shrinkage-based portfolios may be dominated OOS by simple strategies such as the Average Oracle, which is calibrated on outdated returns and yet evaluated on more recent data, and that this behaviour is robust across different datasets and markets.

In the covariance setting, the mismatch between analytically optimal shrinkage rules and the empirical non-stationarity of financial data has been mitigated through physics-informed neural network architectures that embed the random-matrix solution as a limiting case [7, 8]. In these approaches, the RMT shrinkage is recast as a constrained network that remains expressive enough to recover the asymptotically optimal filter under stationarity, while gaining flexibility via additional trainable degrees of freedom to operate beyond that regime. A key element is the explicit enforcement of the problem’s symmetries, notably the rotationally invariant structure implied by RMT. Trained on rolling windows, the model learns to interpolate between a denoising behavior effective in quasi-stationary periods and a predictive correction that captures persistent time variation, yielding improved OOS performance relative to purely analytical shrinkage and naive baselines. Moreover, the neural formulation naturally accommodates decision-aware training. Rather than optimizing an MSE on covariance entries, one can directly minimize an objective aligned with the downstream portfolio loss [9, 10], thereby linking statistical estimation to the relevant economic criterion.

In this work, we introduce a physics-informed learning framework for cross-covariance estimation that preserves the rotational invariances of RMT while operating beyond the stationarity and boundedness regimes required by existing analytical shrinkage. The key idea is to parameterize the cleaned estimator directly in the empirical singular-vector basis and to learn a dimension-agnostic nonlinear map from empirical singular values and marginal projections to cleaned singular values, so that the BBP correction is recovered as a limiting case, but additional degrees of freedom can accommodate macroscopic modes and time variation. Empirically, we observe that the analytically optimal BBP shrinkage can lose robustness on equity returns as the universe size increases in the presence of a dominant market mode, whereas the proposed symmetry-preserving estimator remains stable and yields lower OOS reconstruction error.

2 Two-Stream Architecture for Cross-Covariance Cleaning

We start from the empirical cross-correlation $\widehat{\mathbf{C}}_{XY} \in \mathbb{R}^{n_x \times n_y}$ computed on a window of length Δt_{in} , together with the marginal empirical correlations $\widehat{\mathbf{C}}_{XX}$ and $\widehat{\mathbf{C}}_{YY}$. Let $r := \min(n_x, n_y)$ and write the Singular Value Decomposition (SVD)

$$\widehat{\mathbf{C}}_{XY} = \sum_{k=1}^r \widehat{s}_k \widehat{\mathbf{u}}_k \widehat{\mathbf{v}}_k^\top. \quad (1)$$

When needed, we complete $\{\widehat{\mathbf{u}}_k\}_{k=1}^r$ and $\{\widehat{\mathbf{v}}_k\}_{k=1}^r$ to orthonormal bases of \mathbb{R}^{n_x} and \mathbb{R}^{n_y} ; the extra directions correspond to zero singular values. A RIE keeps the empirical singular vectors $\{\widehat{\mathbf{u}}_k, \widehat{\mathbf{v}}_k\}$ and replaces each $\widehat{s}_k \mapsto \widetilde{s}_k$ by a cleaned value.

Motivated by the structure of BBP shrinkage [5], the correction should depend on the empirical singular spectrum and on how marginal energy projects onto the corresponding singular directions. We therefore define the marginal projections

$$\widehat{\gamma}_k^{(x)} := \widehat{\mathbf{u}}_k^\top \widehat{\mathbf{C}}_{XX} \widehat{\mathbf{u}}_k, \quad k = 1, \dots, n_x, \quad \widehat{\gamma}_k^{(y)} := \widehat{\mathbf{v}}_k^\top \widehat{\mathbf{C}}_{YY} \widehat{\mathbf{v}}_k, \quad k = 1, \dots, n_y, \quad (2)$$

and the aspect ratios $q_x := n_x / \Delta t_{\text{in}}$ and $q_y := n_y / \Delta t_{\text{in}}$. When $n_x \neq n_y$, only the first r singular values are nonzero, while the larger side retains additional marginal directions that enter BBP-type functionals.

To expose this information as a single aligned sequence, we set $p := \max(n_x, n_y)$ and define padded scalars for $k = 1, \dots, p$,

$$\bar{s}_k := \begin{cases} \widehat{s}_k, & k \leq r, \\ 0, & k > r, \end{cases} \quad \bar{\gamma}_k^{(x)} := \begin{cases} \widehat{\gamma}_k^{(x)}, & k \leq n_x, \\ 0, & k > n_x, \end{cases} \quad \bar{\gamma}_k^{(y)} := \begin{cases} \widehat{\gamma}_k^{(y)}, & k \leq n_y, \\ 0, & k > n_y. \end{cases} \quad (3)$$

In particular, for $k > r$ one has $\bar{s}_k = 0$ by construction, while $\bar{\gamma}_k^{(x)}$ (resp. $\bar{\gamma}_k^{(y)}$) can remain nonzero up to $k = n_x$ (resp. $k = n_y$), making the rectangular extra-marginal information explicit.

For each $k \in \{1, \dots, p\}$ we build two stream-specific tokens

$$\boldsymbol{\tau}_k^{(x)} := [\bar{\gamma}_k^{(x)}, \bar{s}_k, q_x]^\top, \quad \boldsymbol{\tau}_k^{(y)} := [\bar{\gamma}_k^{(y)}, \bar{s}_k, q_y]^\top. \quad (4)$$

A shared encoder E_θ , implemented as a token-wise two-layer Multi-Layer-Perceptron (MLP) with a 16-unit hidden layer and LeakyReLU activation, maps each token to a 2-dimensional embedding

$$\mathbf{e}_k^{(x)} = E_\theta(\boldsymbol{\tau}_k^{(x)}) \in \mathbb{R}^2, \quad \mathbf{e}_k^{(y)} = E_\theta(\boldsymbol{\tau}_k^{(y)}) \in \mathbb{R}^2. \quad (5)$$

Then, we fuse the two streams additively,

$$\mathbf{z}_k = \mathbf{e}_k^{(x)} + \mathbf{e}_k^{(y)}, \quad k = 1, \dots, p, \quad (6)$$

and pass the sequence $(\mathbf{z}_1, \dots, \mathbf{z}_p)$ through a two-layer bidirectional LSTM aggregator (hidden sizes 128 and 64). This yields context-aware states $(\mathbf{h}_1, \dots, \mathbf{h}_p)$ that provide each index with global spectral context [7].

A pointwise head g_θ , implemented as a token-wise two-layer MLP with a 252-unit hidden layer and LeakyReLU activation, maps each state to a scalar residual correction $\delta_k = g_\theta(\mathbf{h}_k)$. We keep only the first r corrections and define the cleaned singular values by the additive parametrization

$$\widetilde{s}_k = \widehat{s}_k + \delta_k, \quad k = 1, \dots, r. \quad (7)$$

See Appendix A for a step-by-step derivation showing that BBP is recovered as a special case.

Alternatively, one may enforce a bounded multiplicative correction via $\widetilde{s}_k = \widehat{s}_k \sigma(\delta_k)$, where $\sigma(\cdot)$ denotes the logistic sigmoid, which guarantees $0 \leq \widetilde{s}_k \leq \widehat{s}_k$ (see Appendix A.3 and D.4 for more details). In our experiments, however, this constrained parametrization led to slightly more difficult optimization, and we therefore adopt the additive form. To provide a clearer conceptual overview of the integrated components and the associated data flow, we illustrate the architecture in Figure 1.

Finally, we reconstruct the cleaned cross-covariance by plugging Eq. (7) in Eq. (1) and rescaling with the marginal sample standard deviations,

$$\widetilde{\boldsymbol{\Sigma}}_{XY} = \widehat{\mathbf{D}}_X \widetilde{\mathbf{C}}_{XY} \widehat{\mathbf{D}}_Y. \quad (8)$$

In this work, we train in the cross-correlation domain to isolate multivariate filtering, but the same pipeline can be optimized end-to-end for cross-covariances by including the rescaling in the objective. For supervision, given an OOS

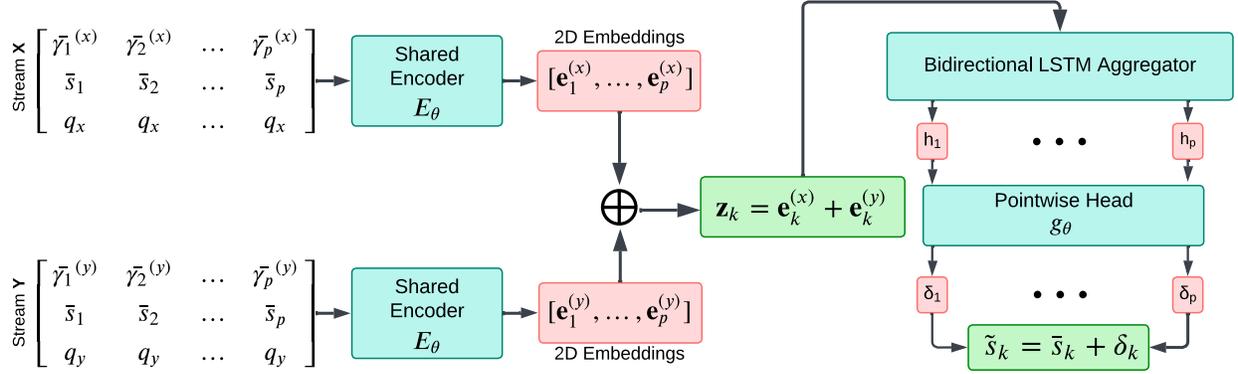


Figure 1: Flowchart of the neural singular value cleaning architecture. The diagram illustrates the construction of dual-stream tokens from marginal projections $\bar{\gamma}$ and singular values \bar{s} , their transformation through a shared encoder E_θ , and the global context aggregation via a bidirectional LSTM and pointwise head g_θ to produce the final additive corrections $\tilde{\delta}_k$.

target $\mathbf{C}_{XY}^{\text{OOS}}$ computed on the subsequent window of length Δt_{out} , we minimize the MSE, i.e., the squared Frobenius loss

$$\mathcal{L}(\theta) = \frac{1}{n_x n_y} \|\tilde{\mathbf{C}}_{XY} - \mathbf{C}_{XY}^{\text{OOS}}\|_F^2. \quad (9)$$

The network has 331,355 learnable parameters, independent of $(n_x, n_y, \Delta t_{\text{in}})$. It can be trained at one range of dimensions and deployed at other sizes without retraining. Moreover, it is equivariant to permutations of the X and Y coordinates. A Python implementation is available on GitHub [11]

3 Validation on Synthetic and Market Data

To assess our Neural Network (NN) based estimator, we first benchmark it on synthetic datasets that reproduce the spectral regimes of finite-rank signals and heavy-tailed noise under which the BBP shrinkage is derived [5]. These experiments therefore constitute a favorable setting for the analytical method, relying on stylized assumptions such as identity auto-covariance blocks that are mathematically convenient yet seldom satisfied in financial data. As reported in Table 1, our estimator matches the BBP performance across these validation tests, indicating that, in stationary regimes where the theoretical assumptions hold, the network can recover the optimal rotationally invariant filtration without introducing spurious structure. We further add a fourth benchmark in which a macroscopic one-factor mode is injected into the population correlation matrix \mathbf{C} , thereby violating the BBP boundedness assumption [5]; in this regime, BBP performance degrades. To confirm that this deterioration is indeed driven by the lack of boundedness, we also designed a Monte Carlo Cross-Validated (CV) BBP variant, which remains robust under the same perturbation. For details on the boundedness issue, the CV procedure, and the data-generating processes and validation protocol, we refer the reader to Appendix B and C.

Moving to real-world data, we train and test the estimators on a universe of the top 1000 U.S. equities by market capitalization (1995-2024), filtered for liquidity and stability (see Appendix D.1). To mitigate look-ahead bias, the top-1000 universe is reconstituted at each window start, so constituents vary over time. In this regime, the non-stationarity of financial time series becomes a predominant factor. According to standard practice, we train a separate model for each calendar year and then test it on the subsequent two years, using as a target a realized OOS correlation matrix computed over a hold-out window of length $\Delta t_{\text{out}} = 240$ trading days. Each yearly model is optimized for 50 epochs (500 gradient steps per epoch) with mini-batches of 32 samples. Each sample is generated by drawing a random date, the total number of assets $n := n_x + n_y \sim \mathcal{U}[50, 500]$, and the relative dimension $\nu := n_x / (n_x + n_y) \sim \mathcal{U}[0.05, 0.95]$, which enforces dimension-agnostic behavior and allows evaluation on larger shapes ($n > 500$) without retraining.

Figure 2 (upper left) reports the OOS MSE as a function of the total number of assets n for a fixed in-sample window ($\Delta t_{\text{in}} = 500$). The Maximum-Likelihood Estimator (MLE, yellow) provides the unfiltered baseline. Across both the n - and ν -sweeps, the proposed NN estimator (blue) achieves the lowest error. For small universes (e.g., $n = 250$), its performance is not statistically distinguishable from BBP (gray). However, as n increases, BBP deteriorates sharply and can even underperform MLE, indicating that the analytical shrinkage can become harmful in this regime. We attribute this failure to the presence of a macroscopic market mode in real returns, which violates the boundedness conditions

Condition	Empirical (MLE)	Analytical (BBP)	Numerical (CV)	Neural Network (NN)
<i>Panel A: Finite-Rank Spiked Models (varying spike % ξ)</i>				
$\xi = 40\%$	2.00×10^{-3}	2.46×10^{-4}	2.46×10^{-4}	2.53×10^{-4}
$\xi = 30\%$	2.00×10^{-3}	1.92×10^{-4}	1.92×10^{-4}	1.95×10^{-4}
$\xi = 20\%$	2.00×10^{-3}	1.34×10^{-4}	1.33×10^{-4}	1.34×10^{-4}
$\xi = 10\%$	2.00×10^{-3}	7.03×10^{-5}	7.00×10^{-5}	7.12×10^{-5}
$\xi = 0\%$	2.00×10^{-3}	3.73×10^{-6}	3.92×10^{-6}	9.02×10^{-6}
<i>Panel B: Heavy-Tailed Bulk Models (varying tail exponent α)</i>				
Gaussian	2.00×10^{-3}	6.16×10^{-4}	6.18×10^{-4}	6.16×10^{-4}
$\alpha = 5.0$	2.00×10^{-3}	6.16×10^{-4}	6.18×10^{-4}	6.16×10^{-4}
$\alpha = 2.5$	2.00×10^{-3}	6.16×10^{-4}	6.18×10^{-4}	6.16×10^{-4}
$\alpha = 1.5$	2.00×10^{-3}	6.19×10^{-4}	6.20×10^{-4}	6.18×10^{-4}
<i>Panel C: White Heavy-Tailed Models (varying tail exponent α)</i>				
Gaussian	2.00×10^{-3}	5.63×10^{-4}	5.64×10^{-4}	5.64×10^{-4}
$\alpha = 5.0$	2.00×10^{-3}	5.63×10^{-4}	5.65×10^{-4}	5.64×10^{-4}
$\alpha = 2.5$	2.00×10^{-3}	2.50×10^{-4}	2.50×10^{-4}	2.51×10^{-4}
$\alpha = 1.5$	2.00×10^{-3}	3.92×10^{-5}	3.90×10^{-5}	4.14×10^{-5}
<i>Panel D: Gaussian Bulk Models with Common Mode (varying mode m)</i>				
$m = 0.5$	1.09×10^{-3}	7.90×10^{-2}	7.45×10^{-4}	7.42×10^{-4}
$m = 0.4$	1.42×10^{-3}	3.29×10^{-2}	9.25×10^{-4}	9.18×10^{-4}
$m = 0.3$	1.66×10^{-3}	9.89×10^{-3}	9.83×10^{-4}	9.84×10^{-4}
$m = 0.2$	1.85×10^{-3}	2.23×10^{-3}	9.84×10^{-4}	9.79×10^{-4}
$m = 0.1$	1.96×10^{-3}	8.85×10^{-4}	8.52×10^{-4}	8.56×10^{-4}

Table 1: Summary of the average MSE for different estimators across three synthetic regimes over 1000 simulations. Bold marks the lowest MSE; ties are bold when a 95% percentile 5,000 copies bootstrap cannot distinguish them. Panel A varies the spike percentage ξ (Finite-Rank). Panel B and C vary the tail exponent α (Heavy-Tailed Bulk and White models), where a lower α implies heavier tails. Panel D is a variation of Panel B (Gaussian) with a one-factor mode injected. All simulations are performed with $n_x = 200$, $n_y = 350$, $\Delta t_{in} = 500$ [5].

underlying BBP. This interpretation is supported by two additional diagnostics. First, the numerical CV estimator is not affected by the same degradation and remains comparable to the NN. Second, in Appendix D.3 we show that pre-processing the data by removing the mode, operationally defined as the daily cross-sectional mean return, eliminates the gross of the BBP underperformance. Finally, when we shuffle the time series to break potential non-stationarity between in-sample and out-of-sample periods, BBP remains unstable whereas CV continues to track the NN, further indicating that the dominant driver is the mode rather than temporal non-stationarity, see Fig. 2 (right). Notably, the NN performance remains robust even for $n > 500$, which constitutes a fully out-of-distribution regime for the model. This region lies beyond the matrix sizes observed during training ($n \leq 500$), and it also corresponds to a different historical period than the data used to fit the network. To further probe this extrapolative regime, in the lower panel of Fig. 2 (lower) we unpack the $n = 1000$ case by varying the relative dimension ν , and we find that the NN gain remains stable across the full range of ν , indicating that its advantage is not concentrated in a specific configuration.

Our working hypothesis is that this behavior is precisely the one targeted by the proposed architecture. In stationary settings, the network should primarily act as a noise-cleaning operator on the empirical spectrum. In contrast, when applied to non-stationary financial data, it can interpolate between denoising and a learned predictive correction that exploits persistent structure beyond sampling noise. To confirm that this gain is derived from capturing genuine temporal signals (non-stationarity), we replicate the analysis on temporally shuffled data. Shuffling destroys the serial correlations responsible for the predictive edge, creating a purely stationary environment. Figures 2 (right) illustrate that performance degradation of BBP persists even in the shuffled regime, confirming the underlying structural driver. As expected, the CV estimator does not degrade and instead matches the NN performance indicating that in a stationary regime the NN learn how to denoise only sample noise. This interpretation is further supported by the feasibility diagnostic reported in Appendix D.4.

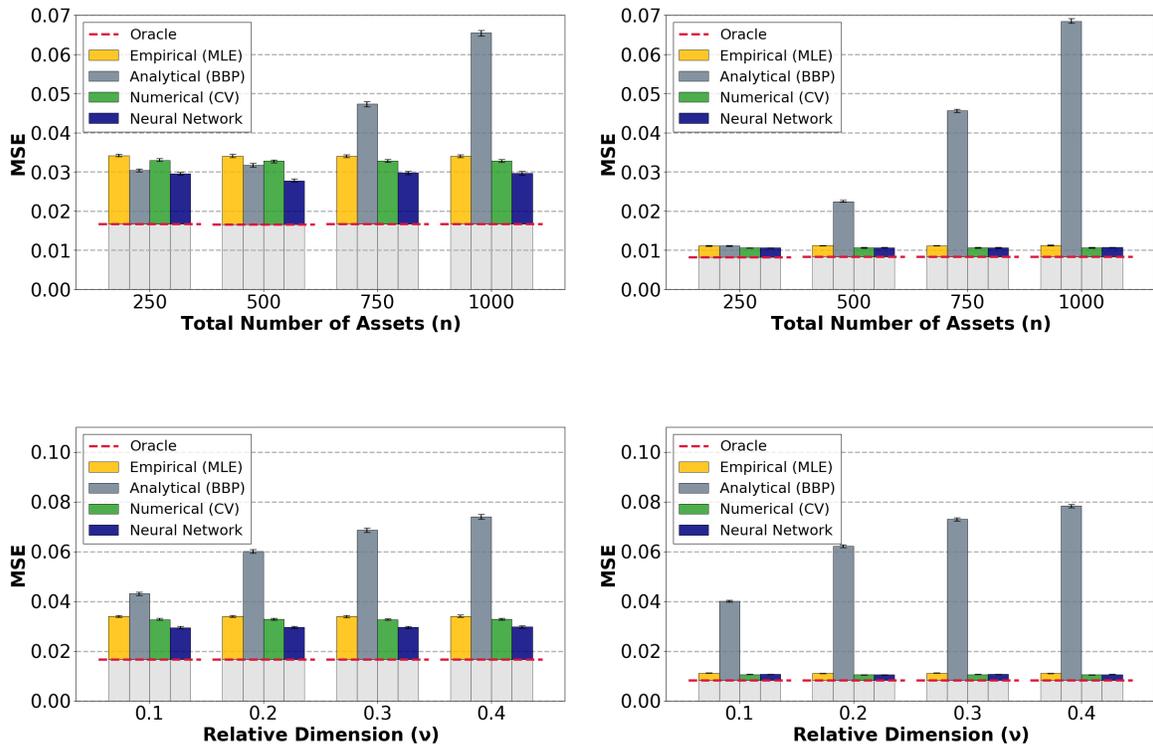


Figure 2: OOS MSE for cross-correlation estimators on financial data covering 2017–2024. For each test window, the estimator is trained on an expanding sample from 1995 up to the year immediately preceding the test period, and the reported value averages 1,000 independent runs for each out-of-sample year. The figures report also the Oracle (red dashed line), which is the lower achievable MSE by a RIE (see eq. (35)). The left panels preserve the chronological order between in-sample and out-of-sample windows, whereas the right panels shuffle dates. The upper panels varies the total number of assets for $\nu = 0.25$, the lower panels varies the relative dimension ν for $n = 1000$. The error bars indicate a 95% percentile 100,000 copies bootstrap confidence interval.

4 Discussion and Conclusion

Across controlled synthetic benchmarks and temporally shuffled equity returns, the proposed estimator matches the performance of the optimal rotationally invariant shrinkage, indicating that it has learned a stationary random-matrix-consistent spectral denoiser rather than introducing spurious structure. In the presence of a pronounced market mode, it also improves upon the corresponding analytic shrinkage method, which is consistent with a mode-aware generalization.

On original (unshuffled) market data, the two cleaners remain comparable for moderate universes, but their behavior diverges as n grows. The analytical BBP shrinkage deteriorates sharply with system size and can even underperform the unfiltered MLE, consistent with sensitivity to macroscopic components (notably a market mode) that violate the boundedness conditions underlying the asymptotics. In contrast, the neural estimator remains stable and achieves lower OOS reconstruction error, including for $n > 500$ (an out-of-distribution regime relative to training). Removing the market mode substantially reduces the BBP collapse, yet a residual gap to the neural estimator persists in the chronological evaluation and largely disappears under temporal shuffling; taken together, these controls support the view that the architecture preserves the BBP symmetry constraints while using OOS supervision to retain only cross-dependencies that persist across adjacent windows, i.e., a symmetry-preserving singular-value map that interpolates between denoising and a mild predictive correction.

A natural direction for future work is to extend the same symmetry-preserving parametrization to lead-lag cross-covariance matrices, where the relevant signal is intrinsically dynamical and cannot be reduced to a single synchronous estimator. In this setting, rather than training via the Frobenius loss, one can adopt an end-to-end decision-aware

objective consistent with a trading strategy so that the shrinkage rule is learned under the same decision loss used for evaluation.

Acknowledgments

This work was performed using HPC resources from the “Mésocentre” computing center of CentraleSupélec and École Normale Supérieure Paris-Saclay supported by CNRS and Région Île-de-France (mesocentre.centralesupelec.fr). E.M. acknowledges a fellowship funded by PNRR for the PhD DOT1608375 in Sistemi complessi per le scienze fisiche, socio-economiche e della vita of Catania University.

References

- [1] Vijay K Chopra, William T Ziemba, et al. “The effect of errors in means, variances, and covariances on optimal portfolio choice”. In: *Journal of Portfolio Management* 19.2 (1993), pp. 6–11.
- [2] Joël Bun, Jean-Philippe Bouchaud, and Marc Potters. “Cleaning large correlation matrices: tools from random matrix theory”. In: *Physics Reports* 666 (2017), pp. 1–109.
- [3] Emanuele Troiani et al. “Optimal denoising of rotationally invariant rectangular matrices”. In: *Mathematical and Scientific Machine Learning*. PMLR. 2022, pp. 97–112.
- [4] Matan Gavish and David L Donoho. “Optimal shrinkage of singular values”. In: *IEEE Transactions on Information Theory* 63.4 (2017), pp. 2137–2152.
- [5] Florent Benaych-Georges, Jean-Philippe Bouchaud, and Marc Potters. “Optimal cleaning for singular values of cross-covariance matrices”. In: *The Annals of Applied Probability* 33.2 (2023), pp. 1295–1326.
- [6] Christian Bongiorno, Damien Challet, and Grégoire Loeper. “Filtering time-dependent covariance matrices using time-independent eigenvalues”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2023.2 (2023), p. 023402.
- [7] Christian Bongiorno, Efstratios Manolakis, and Rosario Nunzio Mantegna. “End-to-End Large Portfolio Optimization for Variance Minimization with Neural Networks through Covariance Cleaning”. In: *arXiv preprint arXiv:2507.01918* (2025).
- [8] Christian Bongiorno, Efstratios Manolakis, and Rosario Nunzio Mantegna. “Neural Network-Driven Volatility Drag Mitigation under Aggressive Leverage”. In: *Proceedings of the 6th ACM International Conference on AI in Finance*. 2025, pp. 449–455.
- [9] Christian Bongiorno and Damien Challet. “Non-linear shrinkage of the price return covariance matrix is far from optimal for portfolio optimization”. In: *Finance Research Letters* 52 (2023), p. 103383.
- [10] Juchan Kim, Inwoo Tae, and Yongjae Lee. “Estimating Covariance for Global Minimum Variance Portfolio: A Decision-Focused Learning Approach”. In: New York, NY, USA: Association for Computing Machinery, 2025, pp. 105–113. ISBN: 9798400722202.
- [11] Manolakis Efstratios and Bongiorno Christian. *CrossRIE*. <https://github.com/Efstratios7/CrossRIE/>. 2026.

A BBP Analytical Cleaner as a Limiting Case of the Two-Stream NN

This section rewrites the analytical BBP singular-value cleaner [1] using the notation of Section 2, and clarifies why the proposed two-stream neural architecture contains that analytical mapping as a special case.

A.1 BBP Analytical Shrinkage: Algorithm 1

Let $\{\mathbf{x}(t) \in \mathbb{R}^{n_x}, \mathbf{y}(t) \in \mathbb{R}^{n_y}\}_{t=1}^{\Delta t_{\text{in}}}$ denote standardized observations on a window of length Δt_{in} , and let $\widehat{\mathbf{C}}_{XY}$, $\widehat{\mathbf{C}}_{XX}$, and $\widehat{\mathbf{C}}_{YY}$ be the corresponding empirical cross-correlation and marginal correlation matrices. Write the SVD

$$\widehat{\mathbf{C}}_{XY} = \sum_{k=1}^r \widehat{s}_k \widehat{\mathbf{u}}_k \widehat{\mathbf{v}}_k^\top, \quad r := \min(n_x, n_y), \quad (10)$$

and define the marginal projections (Section 2)

$$\widehat{\gamma}_k^{(x)} := \widehat{\mathbf{u}}_k^\top \widehat{\mathbf{C}}_{XX} \widehat{\mathbf{u}}_k, \quad \widehat{\gamma}_k^{(y)} := \widehat{\mathbf{v}}_k^\top \widehat{\mathbf{C}}_{YY} \widehat{\mathbf{v}}_k. \quad (11)$$

When $n_x \neq n_y$, BBP requires also the aggregate marginal energy in the ‘‘extra’’ directions of the larger side. Without loss of generality assume $n_x \leq n_y$ (otherwise swap the roles of X and Y); then $r = n_x$ and we define

$$\widehat{\gamma}_{r+}^{(y)} := \sum_{\ell=r+1}^{n_y} \widehat{\mathbf{v}}_\ell^\top \widehat{\mathbf{C}}_{YY} \widehat{\mathbf{v}}_\ell, \quad (12)$$

where $\{\widehat{\mathbf{v}}_\ell\}_{\ell=r+1}^{n_y}$ completes $\{\widehat{\mathbf{v}}_k\}_{k=1}^r$ to an orthonormal basis of \mathbb{R}^{n_y} .

The authors then defines, for each index $k \in \{1, \dots, r\}$, the complex spectral point

$$\zeta_k := \widehat{s}_k + i\eta, \quad \eta := (r n_y \Delta t_{\text{in}})^{-1/12}, \quad (13)$$

and the scalar functionals

$$H(\zeta_k) := \frac{1}{\Delta t_{\text{in}}} \sum_{\ell=1}^r \frac{\widehat{s}_\ell^2}{\zeta_k^2 - \widehat{s}_\ell^2}, \quad (14)$$

$$A(\zeta_k) := \frac{1}{\Delta t_{\text{in}}} \sum_{\ell=1}^r \frac{\widehat{\gamma}_\ell^{(x)}}{\zeta_k^2 - \widehat{s}_\ell^2}, \quad (15)$$

$$B(\zeta_k) := \frac{1}{\Delta t_{\text{in}}} \left(\sum_{\ell=1}^r \frac{\widehat{\gamma}_\ell^{(y)}}{\zeta_k^2 - \widehat{s}_\ell^2} + \zeta_k^{-2} \widehat{\gamma}_{r+}^{(y)} \right). \quad (16)$$

With

$$\Theta(\zeta_k) := \frac{\zeta_k^2 A(\zeta_k) B(\zeta_k)}{1 + H(\zeta_k)}, \quad L(\zeta_k) := H(\zeta_k) - \frac{\Theta(\zeta_k)}{1 + H(\zeta_k) - \Theta(\zeta_k)}, \quad (17)$$

the analytical cleaned singular values are

$$\widehat{s}_k^{\text{BBP}} := \widehat{s}_k \left(\frac{\Im L(\zeta_k)}{\Im H(\zeta_k)} \right)_+, \quad (x)_+ := \max\{x, 0\}. \quad (18)$$

Optionally, Ref. [1] suggests enforcing monotonicity of $k \mapsto \widehat{s}_k^{\text{BBP}}$ via isotonic regression on the pairs $(\widehat{s}_k, \widehat{s}_k^{\text{BBP}})$.

Finally, the BBP rotationally invariant estimator is reconstructed as

$$\widehat{\mathbf{C}}_{XY}^{\text{BBP}} := \sum_{k=1}^r \widehat{s}_k^{\text{BBP}} \widehat{\mathbf{u}}_k \widehat{\mathbf{v}}_k^\top. \quad (19)$$

A.2 Relationship Between the Two-Stream NN and the Analytical Shrinkage

The key observation is that the BBP map in Eq. (18) depends on the data only through (i) the empirical singular spectrum $\{\widehat{s}_\ell\}_{\ell=1}^r$, (ii) the marginal projections $\{\widehat{\gamma}_\ell^{(x)}\}_{\ell=1}^r$ and $\{\widehat{\gamma}_\ell^{(y)}\}_{\ell=1}^{n_y}$ (through $\widehat{\gamma}_{r+}^{(y)}$), and (iii) the aspect ratios $q_x = n_x / \Delta t_{\text{in}}$

and $q_y = n_y / \Delta t_{\text{in}}$ (which enter through η and, more generally, through the high-dimensional scaling). In particular, once the empirical singular vectors are fixed, BBP is a deterministic permutation-equivariant mapping

$$(\widehat{\mathbf{s}}, \widehat{\boldsymbol{\gamma}}^{(x)}, \widehat{\boldsymbol{\gamma}}^{(y)}, q_x, q_y) \mapsto (\widehat{s}_1^{\text{BBP}}, \dots, \widehat{s}_r^{\text{BBP}}), \quad (20)$$

with $\widehat{\mathbf{s}} = (\widehat{s}_1, \dots, \widehat{s}_r)$ and similarly for the projection vectors.

Our NN is constructed precisely to take as input these BBP-sufficient statistics. Using the padded sequences of Eq. (3) in Section 2, we build the stream-specific tokens

$$\boldsymbol{\tau}_k^{(x)} := (\gamma_k^{(x)}, s_k, q_x)^\top, \quad \boldsymbol{\tau}_k^{(y)} := (\gamma_k^{(y)}, s_k, q_y)^\top, \quad k = 1, \dots, p, \quad (21)$$

and process them via a shared encoder, additive fusion, a bidirectional LSTM aggregator, and a pointwise head to output δ_k and hence $\widetilde{s}_k = \widehat{s}_k + \delta_k$ (Eq. (7) in Section 2). Because BBP is rotationally invariant, recovering it amounts to producing the particular residual

$$\delta_k^{\text{BBP}} := \widehat{s}_k^{\text{BBP}} - \widehat{s}_k = \widehat{s}_k \left[\left(\frac{\Im L(\zeta_k)}{\Im H(\zeta_k)} \right)_+ - 1 \right], \quad k = 1, \dots, r, \quad (22)$$

and leaving the empirical singular vectors unchanged, which matches our reconstruction by design.

To see that the NN modules can represent this mapping, it suffices to track how each BBP ingredient can be implemented within the architecture. First, the shared token-wise encoder E_θ can be set to pass through (possibly after a linear rescaling) the components $(s_k, \gamma_k^{(x)}, \gamma_k^{(y)}, q_x, q_y)$ into the latent sequence representation; in particular, the additive fusion $\mathbf{z}_k = \mathbf{e}_k^{(x)} + \mathbf{e}_k^{(y)}$ provides a symmetric mechanism to combine the two marginals while preserving equivariance under permutations of X and Y coordinates.

Second, the bidirectional LSTM aggregator is precisely a mechanism to attach global spectral context to each index k . In BBP, this global context appears through the sums defining $H(\zeta_k)$, $A(\zeta_k)$, and $B(\zeta_k)$, which couple the local point ζ_k (hence \widehat{s}_k) to the entire set $\{(\widehat{s}_\ell, \widehat{\gamma}_\ell^{(x)}, \widehat{\gamma}_\ell^{(y)})\}_\ell$. Since $\zeta_k = \widehat{s}_k + i\eta$ is fixed once k is fixed, each of these quantities is a continuous function of the finite sequence of tokens, obtained by composing elementary operations (squaring, addition, reciprocal) with sequence-wise summation. The LSTM hidden state provides exactly the capacity required to encode such sequence-level aggregates (including the extra-marginal scalar $\widehat{\gamma}_{r+}^{(y)}$ when $n_y > n_x$) and to make them available at each position k .

Third, given access to approximations of $H(\zeta_k)$, $A(\zeta_k)$, and $B(\zeta_k)$ in its context-aware state \mathbf{h}_k , the pointwise head g_θ can implement the remaining BBP algebra: it can form $\Theta(\zeta_k)$, then $L(\zeta_k)$, and finally the ratio $\Im L(\zeta_k) / \Im H(\zeta_k)$ and the residual δ_k^{BBP} . Concretely, this is a scalar continuous map from the state variables to δ_k , and thus falls within the representational capacity of a two-layer MLP head. Therefore, there exists a parameter configuration θ_{BBP} for which the overall NN mapping coincides with BBP on the domain of interest:

$$\widetilde{s}_k(\theta_{\text{BBP}}) = \widehat{s}_k^{\text{BBP}}, \quad k = 1, \dots, r. \quad (23)$$

In summary, the architecture is a symmetry-preserving parametrization of rotationally invariant shrinkage rules: BBP corresponds to a particular choice of parameters that reproduces its analytical nonlinear map from empirical singular values and marginal projections to cleaned singular values, while training allows the same constrained hypothesis class to depart from that stationary prescription when the data are non-stationary or violate BBP regularity conditions.

A.3 Admissible Singular Values Under Fixed Marginals

A cross-correlation block cannot be specified independently of its marginals: given marginal correlation matrices $\mathbf{C}_{XX} \in \mathbb{R}^{n_x \times n_x}$ and $\mathbf{C}_{YY} \in \mathbb{R}^{n_y \times n_y}$, a candidate cross-correlation $\mathbf{C}_{XY} \in \mathbb{R}^{n_x \times n_y}$ is feasible if and only if the full block correlation matrix

$$\mathbf{C} := \begin{pmatrix} \mathbf{C}_{XX} & \mathbf{C}_{XY} \\ \mathbf{C}_{YX} & \mathbf{C}_{YY} \end{pmatrix} \quad (24)$$

is positive semidefinite[2]. We now express this constraint in terms of the singular values after whitening [3, 4], and show how the general (non-whitened) case reduces to the identity-marginal case.

Assume $\mathbf{C}_{XX} \succ \mathbf{0}$ and $\mathbf{C}_{YY} \succ \mathbf{0}$, so that the symmetric square roots $\mathbf{C}_{XX}^{\pm 1/2}$ and $\mathbf{C}_{YY}^{\pm 1/2}$ are well-defined. Define the whitened cross-correlation

$$\mathbf{C}_{XY}^{(w)} := \mathbf{C}_{XX}^{-1/2} \mathbf{C}_{XY} \mathbf{C}_{YY}^{-1/2} \in \mathbb{R}^{n_x \times n_y}. \quad (25)$$

Consider the block-diagonal congruence transform with $\mathbf{W} := \text{diag}(\mathbf{C}_{XX}^{-1/2}, \mathbf{C}_{YY}^{-1/2})$, which yields

$$\mathbf{W} \mathbf{C} \mathbf{W}^\top = \begin{pmatrix} \mathbf{I}_{n_x} & \mathbf{C}_{XY}^{(w)} \\ \mathbf{C}_{XY}^{(w)\top} & \mathbf{I}_{n_y} \end{pmatrix} =: \mathbf{C}^{(w)}. \quad (26)$$

Since \mathbf{W} is invertible, congruence preserves definiteness [2], hence

$$\mathbf{C} \succeq \mathbf{0} \iff \mathbf{C}^{(w)} \succeq \mathbf{0}, \quad \mathbf{C} \succ \mathbf{0} \iff \mathbf{C}^{(w)} \succ \mathbf{0}. \quad (27)$$

We can therefore analyze feasibility in the identity-marginal regime represented by $\mathbf{C}^{(w)}$. Because the top-left block of $\mathbf{C}^{(w)}$ is \mathbf{I}_{n_x} , the Schur complement [5] gives

$$\mathbf{C}^{(w)} \succeq \mathbf{0} \iff \mathbf{I}_{n_y} - \mathbf{C}_{XY}^{(w)\top} \mathbf{C}_{XY}^{(w)} \succeq \mathbf{0}, \quad (28)$$

and similarly $\mathbf{I}_{n_x} - \mathbf{C}_{XY}^{(w)} \mathbf{C}_{XY}^{(w)\top} \succeq \mathbf{0}$. Let $r := \min(n_x, n_y)$ and write the SVD $\mathbf{C}_{XY}^{(w)} = \sum_{k=1}^r s_k^{(w)} \mathbf{u}_k \mathbf{v}_k^\top$, with singular values $s_1^{(w)} \geq \dots \geq s_r^{(w)} \geq 0$. The nonzero eigenvalues of $\mathbf{C}_{XY}^{(w)\top} \mathbf{C}_{XY}^{(w)}$ are $\{(s_k^{(w)})^2\}_{k=1}^r$, so Eq. (28) is equivalent to

$$0 \leq s_k^{(w)} \leq 1, \quad k = 1, \dots, r. \quad (29)$$

Moreover, the strict constraint

$$0 \leq s_k^{(w)} < 1, \quad k = 1, \dots, r, \quad (30)$$

is sufficient to ensure positive definiteness: indeed, Eq. (30) implies $\mathbf{I}_{n_y} - \mathbf{C}_{XY}^{(w)\top} \mathbf{C}_{XY}^{(w)} \succ \mathbf{0}$, hence $\mathbf{C}^{(w)} \succ \mathbf{0}$ by the Schur complement, and therefore $\mathbf{C} \succ \mathbf{0}$ by Eq. (27). Consequently, in the general-marginal setting, testing feasibility reduces to computing $\mathbf{C}_{XY}^{(w)}$ via Eq. (25) and verifying that its singular values satisfy $0 \leq s_k^{(w)} \leq 1$ (or $s_k^{(w)} < 1$ for strict positive definiteness).

Finally, the whitening relation can be inverted as

$$\mathbf{C}_{XY} = \mathbf{C}_{XX}^{1/2} \mathbf{C}_{XY}^{(w)} \mathbf{C}_{YY}^{1/2}, \quad (31)$$

so any procedure that produces a cleaned $\mathbf{C}_{XY}^{(w)}$ with singular values in $[0, 1]$ automatically yields a block matrix \mathbf{C} that is positive semidefinite, and enforcing $s_k^{(w)} < 1$ guarantees that \mathbf{C} is positive definite when \mathbf{C}_{XX} and \mathbf{C}_{YY} are positive definite.

A final remark concerns the meaning of “shrinkage” versus feasibility. Conditions such as $\tilde{s}_k \leq \hat{s}_k$ (cleaned singular values smaller than empirical ones) [1] are typical of pure denoising: under stationarity, the empirical spectrum is interpreted as signal contaminated by sampling noise, and optimal corrections shrink modes toward zero. In our setting, however, the objective is not only denoising but also OOS forecasting. Enforcing a universal shrinkage inequality $\tilde{s}_k \leq \hat{s}_k$ would be unnecessarily stringent, because improving predictive performance may require amplifying specific modes that are weak in-sample yet systematically predictive OOS. The relevant requirement is therefore feasibility in the sense of producing a valid block correlation matrix under fixed marginals, while allowing \tilde{s}_k to exceed \hat{s}_k in selected directions when this improves forecasting. This distinction is supported by the analysis in Appendix D.4.

B Cross-validated Singular-Value Shrinkage

B.1 Boundness in the Presence of a Mode

The analytical BBP shrinkage is derived under high-dimensional assumptions that include boundedness of the operator norm of the population covariance of the concatenated vector. In practice this condition can fail in the presence of a strong common mode. A standard illustration is the equicorrelation matrix

$$\mathbf{A}_n = (1 - m)\mathbf{I}_n + m \mathbf{1}\mathbf{1}^\top, \quad 0 < m < 1, \quad (32)$$

where $\mathbf{1}\mathbf{1}^\top$ has eigenvalue n along $\mathbf{1}$ and 0 on the $(n - 1)$ -dimensional orthogonal subspace. Hence the eigenvalues of \mathbf{A}_n are

$$\lambda_1 = (1 - m) + mn = 1 + m(n - 1), \quad \lambda_2 = \dots = \lambda_n = 1 - m, \quad (33)$$

so that $\|\mathbf{A}_n\|_{\text{op}} = \lambda_1 \sim mn \rightarrow \infty$ as $n \rightarrow \infty$. If such a block is embedded as a principal submatrix of the full population covariance $\mathbf{\Sigma}$, then $\|\mathbf{\Sigma}\|_{\text{op}} \geq \|\mathbf{A}_n\|_{\text{op}}$, and $\mathbf{\Sigma}$ is not bounded in the sense required by the analytical derivation. We therefore introduce a Cross-validation (CV) singular-value shrinkage that provides a stationary BBP-like construction without relying on boundedness of $\|\mathbf{\Sigma}\|_{\text{op}}$.

B.2 Cross-Validation Algorithm

CV offers a direct, model-free proxy for cleaning singular values. The key idea is to estimate the singular directions on a training subsample and to evaluate their associated mode strengths on an independent test subsample; under approximate stationarity within the window, this replaces an optimistically biased in-sample assessment of each mode by an OOS estimate, and fold-averaging further reduces the variance of these mode-wise proxies. This logic is not ad hoc: in the covariance setting, closely related holdout/CV eigenvalue-regularization schemes (NERCOME and variants) come with high-dimensional convergence guarantees, showing that the holdout/CV error approaches the nonlinear-shrinkage/oracle error under standard assumptions [6, 7].

Let $\{\mathbf{X}, \mathbf{Y}\} = \{\mathbf{x}_t \in \mathbb{R}^{n_x}, \mathbf{y}_t \in \mathbb{R}^{n_y}\}_{t=1}^{\Delta t_{\text{in}}}$ denote standardized observations on an in-sample window of length Δt_{in} . We form the empirical cross-correlation

$$\widehat{\mathbf{C}}_{XY} = \frac{1}{\Delta t_{\text{in}}} \sum_{t=1}^{\Delta t_{\text{in}}} \mathbf{x}_t \mathbf{y}_t^\top = \widehat{\mathbf{U}} \text{diag}(\widehat{\mathbf{s}}) \widehat{\mathbf{V}}^\top, \quad (34)$$

where $\widehat{\mathbf{s}} = (\widehat{s}_1, \dots, \widehat{s}_r)$ and $r := \min(n_x, n_y)$ are the empirical singular values, and $(\widehat{\mathbf{U}}, \widehat{\mathbf{V}})$ are empirical singular-vector matrices. The oracle within the rotationally invariant class keeps the empirical singular vectors fixed and replaces \widehat{s}_k by

$$s_k^* = \widehat{\mathbf{u}}_k^\top \mathbf{C}_{XY} \widehat{\mathbf{v}}_k, \quad (35)$$

where \mathbf{C}_{XY} denotes the population cross-correlation and $(\widehat{\mathbf{u}}_k, \widehat{\mathbf{v}}_k)$ are the empirical singular vectors¹. We estimate these oracle diagonals by CV.

We split the time index set $\{1, \dots, \Delta t_{\text{in}}\}$ into κ disjoint folds. For each fold f , we define training and test sets $\mathcal{I}_{\text{train}}^{(f)}$ and $\mathcal{I}_{\text{test}}^{(f)}$ and compute

$$\widehat{\mathbf{C}}_{XY, \text{train}}^{(f)} = \frac{1}{|\mathcal{I}_{\text{train}}^{(f)}|} \sum_{t \in \mathcal{I}_{\text{train}}^{(f)}} \mathbf{x}_t \mathbf{y}_t^\top, \quad \widehat{\mathbf{C}}_{XY, \text{test}}^{(f)} = \frac{1}{|\mathcal{I}_{\text{test}}^{(f)}|} \sum_{t \in \mathcal{I}_{\text{test}}^{(f)}} \mathbf{x}_t \mathbf{y}_t^\top. \quad (36)$$

Let the SVD of the training matrix be

$$\widehat{\mathbf{C}}_{XY, \text{train}}^{(f)} = \widehat{\mathbf{U}}_{\text{train}}^{(f)} \text{diag}(\widehat{\mathbf{s}}_{\text{train}}^{(f)}) \widehat{\mathbf{V}}_{\text{train}}^{(f)\top}. \quad (37)$$

We then define fold-wise CV targets by projecting the held-out matrix onto the training singular directions,

$$\tilde{s}_k^{(f)} := \widehat{\mathbf{u}}_{\text{train}, k}^{(f)\top} \widehat{\mathbf{C}}_{XY, \text{test}}^{(f)} \widehat{\mathbf{v}}_{\text{train}, k}^{(f)}, \quad k = 1, \dots, r, \quad (38)$$

and average over folds,

$$\tilde{s}_k^{\text{CV}} := \frac{1}{\kappa} \sum_{f=1}^{\kappa} \tilde{s}_k^{(f)}. \quad (39)$$

Optionally, we enforce monotonicity of the shrinkage map by isotonic regression, fitting a non-decreasing sequence \tilde{s}_k^{iso} to the pairs $(\tilde{s}_k, \tilde{s}_k^{\text{CV}})$. The final CV-cleaned estimator is reconstructed in the empirical singular-vector basis of the full-sample matrix,

$$\tilde{\mathbf{C}}_{XY}^{\text{CV}} := \widehat{\mathbf{U}} \text{diag}(\tilde{\mathbf{s}}^{\text{iso}}) \widehat{\mathbf{V}}^\top. \quad (40)$$

This provides a stationary BBP-like baseline that is directly computable from data splits and projections and does not require boundedness of $\|\Sigma\|_{\text{op}}$.

C Synthetic Data Validation

Most of the benchmark models presented here are based on Ref. [1]. These experiments are designed to verify the performance of the proposed NN shrinkage estimator within specific spectral regimes, ranging from finite-rank signals to heavy-tailed noise. For completeness, we included a fourth benchmark that includes a positive mode.

¹When the Oracle is reported on real data, it is used only as an ex post lower bound for the achievable OOS MSE within the rotationally invariant class, and not to construct the cross-validated estimator. Accordingly, the unobservable population cross-correlation \mathbf{C}_{XY} in Eq. (35) is replaced by the realized OOS cross-covariance/cross-correlation computed on the OOS window.

C.1 Synthetic Benchmarks

Benchmark I: Finite-Rank Spiked Models

The first set of benchmarks are reported as Models 1-5 in Ref. [1]. We consider a Gaussian block model for (\mathbf{X}, \mathbf{Y}) with population covariance

$$\Sigma := \begin{pmatrix} \mathbf{I}_{n_x} & \Sigma_{XY} \\ \Sigma_{XY}^\top & \Sigma_{XY}^\top \Sigma_{XY} + \sigma^2 \mathbf{I}_{n_y} \end{pmatrix}, \quad (41)$$

with $\sigma^2 = 0.5$ as in Ref. [1]. The cross-block Σ_{XY} is constructed as a finite-rank matrix with Haar-distributed singular vectors and a fraction ξ of non-zero singular values (“spikes”): letting $\rho := \lfloor \xi \min(n_x, n_y) \rfloor$,

$$\Sigma_{XY} = \mathbf{U} \text{diag}(s_1, \dots, s_\rho, 0, \dots) \mathbf{V}^\top, \quad s_i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}[0.2, 0.5], \quad (42)$$

where $\mathbf{U} \in \mathbb{R}^{n_x \times n_x}$ and $\mathbf{V} \in \mathbb{R}^{n_y \times n_y}$ are independent Haar orthogonal matrices.

This benchmark therefore fixes an isotropic and perfectly conditioned X -marginal, while inducing a structured Y -marginal aligned with the cross-covariance directions plus an additional isotropic noise component $\sigma^2 \mathbf{I}_{n_y}$; Panel A of Table 1 reports performance as ξ varies from 0% to 40%.

Benchmark II: Heavy-Tailed Bulk Models

The second family (Models 6–10 in Ref. [1]) assesses performance under heavy-tailed regimes where the spectrum is broad and no explicit finite-rank spike component is planted. We define a rectangular matrix \mathbf{W} of size $n \times 2n$ (with $n := n_x + n_y$) with entries drawn from a distribution which is either standard Gaussian (Model 6) or a symmetric heavy-tailed law with tail exponent α (Models 7–10). The population covariance matrix of the concatenated vector $(\mathbf{X}^\top, \mathbf{Y}^\top)^\top$ is defined as:

$$\Sigma := \frac{1}{2n} \mathbf{W} \mathbf{W}^\top, \quad (43)$$

and $\Sigma_{XX}, \Sigma_{XY}, \Sigma_{YY}$ denote its blocks under the partition $n := n_x + n_y$. Lower values of α represent heavier tails. In this construction, any cross-dependence between \mathbf{X} and \mathbf{Y} is not driven by isolated spikes but instead emerges from the random anisotropy of Σ , whose eigenvalue distribution becomes increasingly dispersed as α decreases. Panel B of Table 1 compares the estimators across varying tail exponents. In our simulations, the $\alpha = 5$ instance is statistically indistinguishable from the Gaussian case, which is expected since $\alpha = 5$ already has sufficiently light tails (finite fourth moment) and convergence to the Gaussian/universal regime is rapid at the considered sample sizes; we nevertheless report both cases to match the benchmark specification of Ref. [1].

Benchmark III: White Heavy-Tailed Cross-Correlation Models

The third family (Models 11–15 in Ref. [1]) isolates the cross-block estimation problem by enforcing spherical auto-covariance blocks while retaining the same cross-block induced by the heavy-tailed construction of Benchmark II. Starting from the population covariance Σ in Benchmark II, we extract the corresponding cross-block Σ_{XY} . We then define a “white-marginals” population covariance by replacing the marginal blocks with scaled identities while keeping the cross-block fixed:

$$\Sigma := \begin{pmatrix} s_{\max} \mathbf{I}_{n_x} & \Sigma_{XY} \\ \Sigma_{XY}^\top & s_{\max} \mathbf{I}_{n_y} \end{pmatrix}, \quad s_{\max} := \|\Sigma_{XY}\|_2. \quad (44)$$

This construction removes any marginal anisotropy present in Σ_{XX} and Σ_{YY} . Panel C of Table 1 reports the results as the tail exponent α (inherited from Benchmark II) varies across the values considered in Ref. [1]. As in the previous benchmark, results for $\alpha = 5$ are statistically indistinguishable from those for Gaussian data. This is expected because $\alpha = 5$ corresponds to relatively light tails (in particular, a finite fourth moment), so finite-sample behavior quickly approaches the Gaussian/universal regime at the considered dimensions. We nevertheless report both settings to adhere to the benchmark definition in Ref. [1].

Benchmark IV: Gaussian Bulk Models with Common Mode

This benchmark starts from the Gaussian instance of Benchmark II (Model 6), i.e., the bulk-only construction with \mathbf{W} i.i.d. standard normal and population covariance

$$\Sigma^{(G)} := \frac{1}{2n} \mathbf{W} \mathbf{W}^\top, \quad n := n_x + n_y. \quad (45)$$

To emulate empirical “market-mode” behavior (a positive common component across all variables), we inject an equicorrelated mode through a multiplicative (congruence) transformation. For a mode strength $m \in [0, 1]$, we define

$$\mathbf{M} := (1 - m)\mathbf{I}_n + m\mathbf{1}\mathbf{1}^\top, \quad (46)$$

and set

$$\boldsymbol{\Sigma}^{(M)} = \mathbf{M}^{1/2} \boldsymbol{\Sigma}^{(G)} \mathbf{M}^{1/2}, \quad \mathbf{D} := \text{diag}\left(\boldsymbol{\Sigma}^{(M)}\right), \quad \boldsymbol{\Sigma} = \mathbf{D}^{-1/2} \boldsymbol{\Sigma}^{(M)} \mathbf{D}^{-1/2}. \quad (47)$$

Here $\mathbf{M}^{1/2}$ denotes the unique positive definite square root of \mathbf{M} , and the diagonal normalization $\mathbf{D}^{-1/2}(\cdot)\mathbf{D}^{-1/2}$ rescales the result to unit diagonal (correlation form). Finally, we extract by $\boldsymbol{\Sigma}_{XX}$, $\boldsymbol{\Sigma}_{XY}$, and $\boldsymbol{\Sigma}_{YY}$ the blocks of $\boldsymbol{\Sigma}$ under the partition $n := n_x + n_y$.

C.2 Training on Synthetic Data

We train an NN model for each benchmark described above. The training uses batches that are generated on the fly, so the network is effectively exposed to a virtually unbounded stream of non-repeated observations. This design mitigates memorization of a finite training set and helps the learned shrinkage map generalize across the benchmark’s parameter range.

The total system dimension is sampled uniformly from $n \sim \mathcal{U}[50, 500]$, and the sample length is sampled from $\Delta t_{\text{in}} \sim \mathcal{U}[200, 1200]$ every batch. The relative dimension is drawn as $\nu \sim \mathcal{U}[0.05, 0.95]$, and to ensure broad coverage, the spike percentage ξ is sampled uniformly on every instance from $\mathcal{U}[0.2, 0.35]$ for model I, and the tail exponent α is uniformly sampled from $\mathcal{U}[1.5, 2.5]$ for models II and II models; finally, the mode m of model IV is sampled uniformly in $\mathcal{U}[0.2, 0.3]$,

The neural network’s weights are optimized using the Adam optimizer with an initial learning rate of 0.001 and an exponential decay rate of 0.99 per epoch over 50 epochs. We employ gradient accumulation over 4 steps with a base batch size of 32 (resulting in an effective batch size of 128) to allow for multiple shapes $(n, \Delta t_{\text{in}})$ within the same effective batch.

For testing (Table 1), we set $n_x = 200$, $n_y = 350$, and $\Delta t_{\text{in}} = 500$ as in Ref. [1]. While the dimensions n and Δt_{in} are set roughly around their training ranges, the benchmark parameters (ξ, α, m) are tested out of the training distribution.

D Market Data Validation

D.1 Universe Selection

The data filtering procedure largely follows Ref. [8]. Our objective is to construct a time-varying investable universe that enforces (i) stable short-term liquidity and (ii) sufficiently long and homogeneous trading histories, while aiming to mitigate lookahead biases in the selection process.

The dataset comprises all NYSE- and NASDAQ-listed equities from January 1, 1990, to December 31, 2024, including ADRs, and excludes ETFs, funds, and non-operating vehicles. Only common shares are retained. To avoid the main source of look-ahead bias, the eligible universe on each day t is determined using information available strictly prior to t , and any stock that will be delisted within the subsequent $\Delta t_{\text{out}} = 240$ trading days is excluded. While this screening necessarily introduces a mild survivorship bias, our objective is to assess cross-covariance estimation and forecasting accuracy rather than to backtest a tradable investment strategy; accordingly, we expect its impact on our conclusions to be limited. See Appendix D.1.1 for a discussion of the alternatives.

Each candidate stock must exhibit a stable trading history over the full five-year lookback window. In particular, within every rolling one-year subwindow, the stock must participate in the closing auction on at least 95% of trading days. This removes securities with fragmented price histories arising from recent IPOs, prolonged suspensions, ticker changes, or other corporate events. Additional liquidity requirements are imposed on the most recent 5–220 days: the closing auction must execute each day, daily traded volume must exceed 1% of both shares outstanding and market capitalization, the prior-day share price must satisfy $(10 \leq p_{t-1} \leq 2000)$ USD, and the number of shares outstanding must exceed 5 million. These constraints prevent the inclusion of illiquid or distressed securities that would lead to unstable cross-covariance estimates.

To avoid degeneracies in the covariance matrices, we remove stocks with almost zero short-term variance. Such cases, typically acquisition targets trading near their offer price, flatten the volatility structure and collapse multivariate optimization problems into trivial univariate ones. We eliminate all securities whose log-standard-deviation falls below 1.5 IQRs of the cross-sectional distribution, computed over both 5- and 20-day windows. We also retain only one share

class per issuer (selecting the class with the highest market capitalization) and remove the lower-cap member of any pair with in-sample correlation greater than 0.95.

After applying all filters, the remaining stocks are ranked by market capitalization, and the top 1000 are selected each trading day. The universe size is therefore fixed at 1000, though membership naturally changes over time, particularly near the lower tail of the size distribution. The final panel spans 1995 – 2024, comprising roughly 2.7 million stock–day observations and approximately 4,000 unique securities.

D.1.1 Delisting Risk and Alternative Remedies

In the stock-selection step, we restrict the OOS universe to names for which returns are observed throughout the full one-year OOS window, thereby avoiding mechanically missing observations induced by delistings. This choice introduces a mild look-ahead (survivorship) bias, since survival over the OOS horizon is implicitly conditioned on future information. There are, however, three practical ways to remove this bias.

One option is to shorten the OOS horizon to a window that does not exceed the public-notice period preceding an exchange delisting; in particular, delisting via Form 25 becomes effective 10 days after filing and exchanges are required to provide public notice no fewer than 10 days in advance [9], so an OOS window on the order of 10 days limits the amount of forward-looking conditioning, at the cost of higher variance in the learning signal due to fewer OOS observations.

A second option is to estimate correlations using pairwise-overlapping observations (i.e., computing each product only over dates on which both series are observed), which eliminates the need to pre-filter for full-horizon availability but can yield an indefinite “correlation” estimate with negative eigenvalues; since our objective is an MSE criterion against a target and feasibility is enforced downstream, this alternative target need not be detrimental in principle, although it changes the estimation problem.

A third option is to keep the present protocol for training (to maintain strict comparability across methods) and, at the application stage, quantify sensitivity by relaxing the survivorship restriction in the OOS evaluation and reporting how performance varies with the degree of testing-time rigor.

At the current stage, where we do not yet map forecasts into a trading strategy with a well-defined economic objective, this last dimension is secondary, because we lack an external criterion that would justify selecting one debiasing protocol over another.

D.2 Training on Real Market Data

The model is trained under a walk-forward optimization protocol on a universe of U.S. equities, with an OOS evaluation spanning 2017–2024. We fit a sequence of 7 distinct models using an expanding training window. Specifically, at each iteration, the network is trained from scratch on data from 1995 up to the year immediately preceding the corresponding test segment; the training set is then expanded, and the procedure is repeated for the next iteration.

To ensure that the model learns to operate across varying problem dimensions, we adopt a stochastic dynamic sampling strategy. In each training batch, the total number of assets is drawn uniformly at random as $n \sim \mathcal{U}[50, 500]$. This set is partitioned into two disjoint groups, \mathbf{X} and \mathbf{Y} , according to a random relative dimension $\nu \sim \mathcal{U}[0.05, 0.95]$. The lookback window length is sampled every batch in the range $\Delta t_{\text{in}} \sim \mathcal{U}[200, 1200]$ days.

The inputs are the empirical correlation matrices computed from daily adjusted-close returns. Specifically, we estimate the in-sample correlation matrices $\hat{\mathbf{C}}_{xx}$ and $\hat{\mathbf{C}}_{yy}$, together with the in-sample cross-correlation matrix $\hat{\mathbf{C}}_{xy}$, all computed over a window of length Δt_{in} . These three matrices are provided as input to the network. The model outputs a cleaned estimate of the cross-correlation, which is trained by minimizing the MSE with respect to the realized OOS target $\mathbf{C}_{xy}^{\text{OOS}}$, computed over a future hold-out window of $\Delta t_{\text{out}} = 240$ trading days.

When explicitly testing the role of non-stationarity, the pipeline above is modified through a shuffling procedure designed to destroy temporal structure while preserving the daily cross-sectional distribution of returns. Concretely, we first merge the IS and OOS segments into a single pool of trading days (each day represented by the full vector of cross-sectional returns), and then apply a random permutation to these days. We then re-split the permuted sequence into in-sample and OOS windows with the same lengths as in the original protocol. From these reshuffled returns, we recompute the corresponding empirical input matrices ($\hat{\mathbf{C}}_{xx}$, $\hat{\mathbf{C}}_{yy}$, $\hat{\mathbf{C}}_{xy}$) and the realized target $\mathbf{C}_{xy}^{\text{OOS}}$ using the same estimators and window lengths as above. By construction, the resulting IS and OOS sets are drawn from the same distribution and therefore differ only through finite-sample fluctuations (i.e., sampling noise induced by finite Δt_{in} and Δt_{out}).

Finally, the neural network’s weights are optimized using the Adam optimizer with an initial learning rate of 0.001 and an exponential decay rate of 0.99 per epoch over 50 epochs. We employ gradient accumulation over 4 steps with a base batch size of 32 (resulting in an effective batch size of 128) to allow for multiple shapes $(n, \Delta t_{\text{in}})$ within the same effective batch.

D.3 Market-Mode Removal as a Boundedness Diagnostic

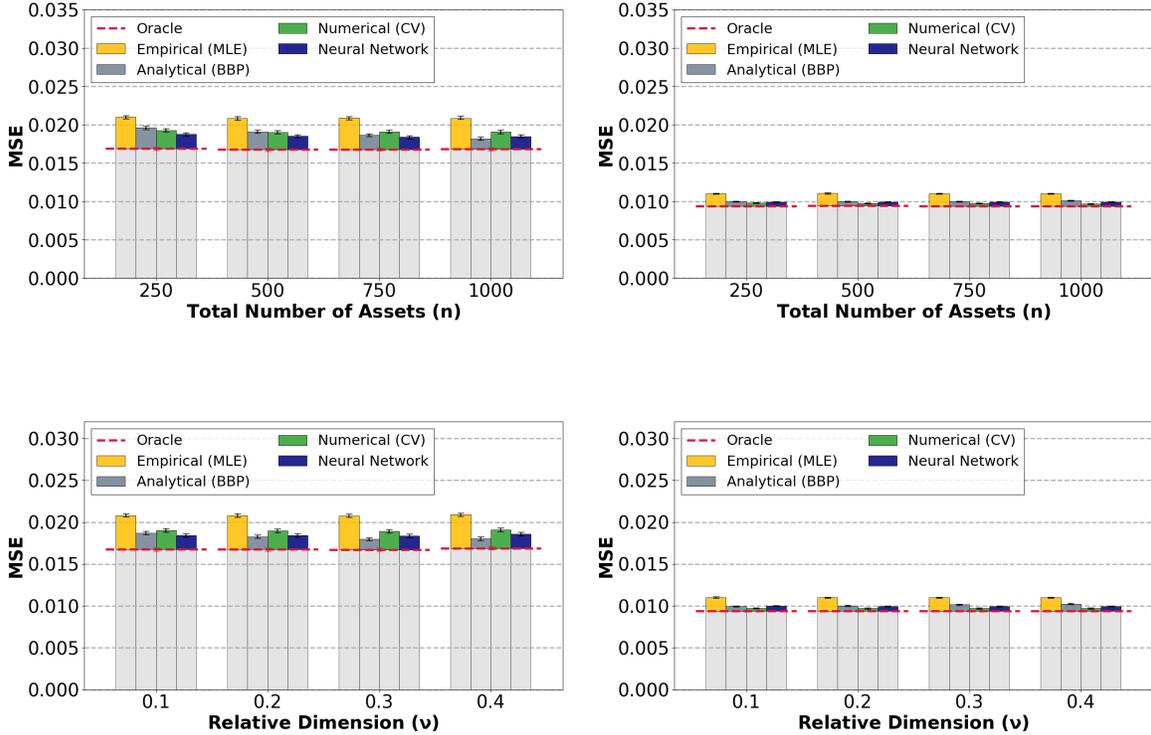


Figure 3: OOS MSE for cross-correlation estimators on financial data covering 2017–2024 with the market mode removed. For each test window, the estimator is trained on an expanding sample from 1995 up to the year immediately preceding the test period, and the reported value averages 1,000 independent runs for each OOS year. The left panels preserve the chronological order between in-sample and OOS windows, whereas the right panels shuffle dates. The upper panels vary the total number of assets for $\nu = 0.25$, while the lower panels vary the relative dimension ν for $n = 1000$. The error bars indicate a 95% bootstrap confidence interval.

In the market experiment of Section 3, the analytical BBP shrinkage deteriorates sharply as the system size increases, and can even underperform the unfiltered MLE. Our working hypothesis is that this is not primarily a consequence of temporal non-stationarity, but rather a structural failure induced by the presence of a macroscopic market mode. In particular, a strong common component can violate the boundedness assumption required by the BBP asymptotics (cf. Appendix B.1), making the corresponding shrinkage map unstable in finite samples. To test this mechanism directly, we repeat the full real-data pipeline after explicitly removing the market mode from returns, and we re-train the neural estimator on the transformed data to keep the comparison protocol unchanged.

Let $R_{i,t}$ denote the raw daily return of asset $i \in \{1, \dots, n\}$ on day $t \in \{1, \dots, \Delta t_{\text{in}}\}$ within a given estimation window. We define the market mode operationally as the cross-sectional mean return on each day and subtract it from every asset,

$$\hat{\mu}_t := \frac{1}{n} \sum_{i=1}^n R_{i,t}, \quad R_{i,t}^{(0)} := R_{i,t} - \hat{\mu}_t. \quad (48)$$

This transformation is applied independently within each window, both in-sample and OOS. From the mode-removed returns $\{R_{i,t}^{(0)}\}$ we then recompute the same empirical input matrices as in the baseline protocol, namely $(\hat{C}_{xx}, \hat{C}_{yy}, \hat{C}_{xy})$

on an in-sample window of length Δt_{in} , as well as the realized target $\mathbf{C}_{xy}^{\text{OOS}}$ on a subsequent hold-out window of length $\Delta t_{\text{out}} = 240$ trading days, leaving all estimators unchanged. The NN is trained again under the identical walk-forward expanding-window scheme described in Appendix C.2 (including the same dynamic batching over $(n, \nu, \Delta t_{\text{in}})$), so that any performance change can be attributed to the preprocessing rather than to a modified learning protocol.

Figure 3 reports the resulting OOS MSE for 2017–2024, averaging 1,000 independent runs per OOS year. In this mode-removed regime, the gross BBP underperformance is substantially reduced: the sharp deterioration observed in the raw-return experiment is no longer present, and BBP becomes again competitive with the other cleaners across both the n -sweep (at fixed ν) and the ν -sweep (at fixed $n = 1000$).

Importantly, this improvement does not eliminate the systematic gap to the learned estimator when the data are kept in the original chronological order. Across most problem sizes, both BBP and the CV-based shrinkage remain consistently worse than the NN in OOS MSE. The only notable exception occurs at $n = 1000$, where BBP attains NN-level performance and can even marginally outperform it. This regime, however, is decisively out-of-distribution for the network, since training batches were restricted to $n \leq 500$; as a consequence, the $n = 1000$ comparison probes a different spectral scale than the one observed during training and should be interpreted as a robustness stress test rather than an in-distribution performance claim.

D.4 Feasibility Diagnostic via Canonical-Correlation Boundedness

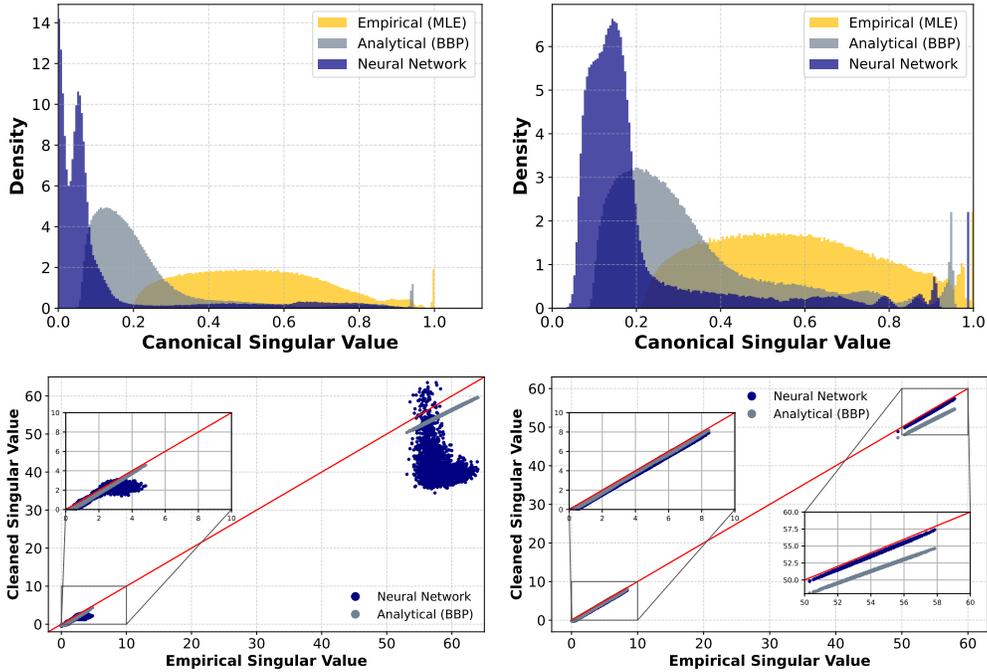


Figure 4: Feasibility diagnostic for the reconstructed cross-correlation block. Top row: empirical distributions of canonical singular values (singular values of the whitened block $\tilde{\mathbf{C}}_{XY}^{(w)}$ obtained via Eq. (25)) for the original chronological pipeline (left) and the shuffled control (right). Bottom row: cleaned singular values in the unwhitened domain plotted against the empirical singular values, again for the original pipeline (left) and the shuffled control (right).

Appendix A.3 characterizes when a candidate cross-correlation block is compatible with fixed marginal correlation matrices. Specifically, for given marginals $\hat{\mathbf{C}}_{XX} \succ \mathbf{0}$ and $\hat{\mathbf{C}}_{YY} \succ \mathbf{0}$, feasibility of a cleaned cross-block $\tilde{\mathbf{C}}_{XY}$ is equivalent to positive semidefiniteness of the associated block correlation matrix in Eq. (24). By the whitening reduction in Eq. (25), this constraint is in turn equivalent to boundedness of the singular values of the whitened block, i.e., the canonical correlations must satisfy $0 \leq s_k^{(w)} \leq 1$ as stated in Eq. (29) (and the strict inequality in Eq. (30) suffices for positive definiteness).

We use this result as an ex-post diagnostic for the neural estimator in the market experiment. We consider 3,200 OOS realizations produced by the final walk-forward model (trained on the longest available history), fixing $n = 300$, $\nu = 0.3$, and $\Delta t_{\text{in}} = 1000$. For each realization we form $(\hat{\mathbf{C}}_{XX}, \hat{\mathbf{C}}_{YY}, \hat{\mathbf{C}}_{XY})$, apply the network to obtain $\tilde{\mathbf{C}}_{XY}$,

whiten it by applying Eq. (25) with $(\widehat{\mathbf{C}}_{XX}, \widehat{\mathbf{C}}_{YY}, \widetilde{\mathbf{C}}_{XY})$, and record the singular values of the resulting $\widetilde{\mathbf{C}}_{XY}^{(w)}$. Figure 4 (top left) shows that the canonical singular values are overwhelmingly feasible: 99.96% of them lie in $[0, 1]$, only 0.04% exceed 1, and none fall below 0. Hence, feasibility violations are extremely rare, and the reconstructed block matrix is almost always positive semidefinite in the sense of Eq. (24).

Importantly, this diagnostic is not equivalent to requiring a shrinkage condition ($\tilde{s}_k \leq \widehat{s}_k$) in the unwhitened singular spectrum. Indeed, Figure 4 (bottom left) shows that the learned map can produce $\tilde{s}_k > \widehat{s}_k$ for a non-negligible set of directions, which would be ruled out by a multiplicative sigmoid parametrization $\tilde{s}_k = \widehat{s}_k \sigma(\cdot)$. This behavior is nonetheless compatible with feasibility because positive semidefiniteness is controlled by the canonical bound of Eq. (29), not by an ordering between unwhitened singular values. The latter type of inequality is specific to purely denoising interpretations under stationarity assumptions, whereas here the network is trained against an OOS target and may optimally amplify some directions while still satisfying the feasibility constraint implied by Appendix A.3.

Finally, in the shuffled control (where the chronological structure is destroyed), the feasibility diagnostic becomes exact: no canonical singular values fall outside $[0, 1]$ (top right). In addition, the predicted singular values become automatically monotone (bottom right), without requiring isotonic post-processing, and closely resemble the stationary BBP-like shrinkage. The remaining discrepancy concentrates in the leading modes, which are associated with the market component and are precisely the directions where BBP is most fragile in the real-data benchmarks.

Supplementary References

- [1] Florent Benaych-Georges, Jean-Philippe Bouchaud, and Marc Potters. “Optimal cleaning for singular values of cross-covariance matrices”. In: *The Annals of Applied Probability* 33.2 (2023), pp. 1295–1326.
- [2] Robb J Muirhead. *Aspects of multivariate statistical theory*. John Wiley & Sons, 2009, pp. 239–246, 585–586.
- [3] Anna Bykhovskaya and Vadim Gorin. “Canonical Correlation Analysis”. In: *arXiv preprint arXiv:2411.15625* (2024).
- [4] Takoua Jendoubi and Korbinian Strimmer. “A whitening approach to probabilistic canonical correlation analysis for omics data integration”. In: *BMC bioinformatics* 20.1 (2019), p. 15.
- [5] Diane Valerie Ouellette. “Schur complements and statistics”. In: *Linear Algebra and its Applications* 36 (1981), pp. 187–295.
- [6] Clifford Lam. “Nonparametric Eigenvalue-Regularized Precision or Covariance Matrix Estimator”. In: *The Annals of Statistics* 44.3 (2016), pp. 928–953.
- [7] Karim M. Abadir, Walter Distaso, and Filip Žikeš. “Design-free estimation of variance matrices”. In: *Journal of Econometrics* 181.2 (2014), pp. 165–180.
- [8] Christian Bongiorno, Efstratios Manolakis, and Rosario Nunzio Mantegna. “End-to-End Large Portfolio Optimization for Variance Minimization with Neural Networks through Covariance Cleaning”. In: *arXiv preprint arXiv:2507.01918* (2025).
- [9] U.S. Securities and Exchange Commission. *17 CFR §240.12d2-2 – Removal from Listing and Registration*. <https://www.law.cornell.edu/cfr/text/17/240.12d2-2>. 2024.