Is Correlation Neglect Bad for Portfolio Diversification?

Yuan Chen *

ABSTRACT

While correlations between stocks play a central role in Markowitz's portfolio selection and diversification, empirical evidence shows that investors often neglect them, relying instead on simple heuristics rather than the Pearson correlation coefficient. Although standard theory suggests that incorporating correlations should improve portfolio performance, empirical studies reveal that ignoring correlations can sometimes yield better out-of-sample results. This raises two key questions: Is correlation neglect always harmful? And which aspects of correlation are truly essential for portfolio construction? In this paper, I propose a transformation that isolates the directional component of correlations and demonstrate that both fully ignoring correlations and fully relying on them are suboptimal. Empirically, I show that the directional component captures the most relevant information for diversification and plays a critical role in improving portfolio performance. By distinguishing between beneficial and irrelevant aspects of correlations, this paper provides a new framework for constructing more robust and effective investment portfolios.

Keywords: Correlation neglect; Portfolio optimization; Kendall's tau correlation; Portfolio diversitvation.

*affiliated with the Department of Statistics and Operations Research at the University of Vienna, as well as the Vienna Graduate School of Finance (VGSF) and the Oskar Morgenstern Doctoral School (OMDS). Email: yuan.chen@univie.ac.at

I. Introduction

Modern portfolio theory, pioneered by Markowitz (1952), is a cornerstone of financial economics. Central to its implementation is the covariance matrix of asset returns, which dictates the benefits of diversification. The correlation structure, in particular, is critical for constructing optimal portfolios. Accurately estimating and utilizing the correlation matrix is therefore considered a fundamental challenge for investors and asset managers.

Given its theoretical importance, a prominent literature in finance investigates "correlation neglect," a cognitive bias where investors tend to ignore or underweight correlation information when making portfolio decisions (Eyster and Weizsacker, 2010; Kallir and Sonsino, 2009; Enke and Zimmermann, 2019; Ungeheuer and Weber, 2021). This research stream views the failure to account for correlation as a heuristic error and seeks to "debias" investors, for instance by improving the experiential presentation of correlation data to encourage its use (Laudenbach et al., 2023). The implicit prescription from this literature is clear: investors should pay more attention to correlation.

In stark contrast, a vast literature in finance raises serious doubts about the practical utility of sample correlation matrices. It is well-documented that the sample covariance matrix, $\widehat{\mathbf{\Sigma}}$, suffers from severe estimation errors, leading to poor portfolio optimization performance based on these estimates (DeMiguel et al., 2009). In fact, many optimal portfolios constructed using sample covariance matrices can be easily outperformed by simple heuristic strategies. For example, the equal-weighted ($1/N$) portfolio, which implicitly ignores correlations, has been shown to consistently outperform optimized portfolios that rely on noisy sample estimates. Moreover, Kirby and Ostdiek (2012) demonstrate that the primary source of this instability arises not from the estimation of individual asset volatilities but rather from the noisy pairwise correlation estimates. When focusing solely on volatility information, it is possible to construct portfolios that outperform the equal-weighted benchmark while avoiding extreme portfolio weights. From this perspective, it may even be rational for investors to disregard sample correlations altogether.

This tension gives rise to a fundamental paradox. On the one hand, a large body of literature urges investors to incorporate correlation information that they are often prone to neglect. On the other hand, empirical evidence suggests that relying on sample correlations can be counterproductive. How can these two perspectives be reconciled? If the full correlation matrix is too noisy to be reliable, yet ignoring it entirely discards valuable information, an important question arises: which components of correlation provide stable diversification benefits, and which are dominated by estimation noise?

In this paper, we bridge this gap by proposing a simple yet powerful decomposition of the correlation coefficient into two distinct components: its magnitude (the strength of comovement) and its direction (the sign, indicating positive or negative comovement). We empirically demonstrate that the direction of correlation is remarkably stable and robust, even when using short time series or dealing with a large cross-section of assets. In contrast, the magnitude of correlation is highly unstable and constitutes the primary source of estimation error that undermines portfolio optimization, especially in low signal-to-noise ratio environments. By isolating the stable directional information, we are able to construct portfolios that retain the diversification benefits of correlations while discarding the noisy components.

To further explain these findings, we adopt both simulation-based and theoretical perspectives. We show that correlation neglect naturally arises in environments where relevant information must be estimated from noisy samples and where uncertainty about the future violates the assumptions of classical rational decision theory. Under such conditions, using the full correlation matrix often requires extremely large samples to achieve reliable asymptotic results. Counterintuitively, ignoring part of the correlation structure can sometimes improve out-of-sample performance, a phenomenon known as the less-is-more effect: there exists an inverse-U-shaped relationship between the amount of information used and the resulting accuracy of portfolio decisions.

Neglecting correlations introduces some specification error, but it can substantially

reduce the variance caused by sampling noise, thereby improving robustness. Our analysis provides guidance on how investors can selectively exploit correlation information without overfitting to unreliable estimates. In particular, we show that focusing on the directional component offers a tractable and effective solution: it captures the essential information needed for diversification, requires minimal additional computation, and can be seamlessly incorporated into investment applications. For instance, a simple "counting-based" approach, tracking the frequency of co-movement in stock returns, can be used to help investors make better portfolio decisions without relying on unstable correlation magnitudes.

We also analyze the statistical properties of our newly proposed correlation matrix through a spectral analysis of its eigenvalues. Our results show that the eigenvalues of the new matrix are more closely centered around their true values. This finding provides insight into how correlation neglect influences portfolio decisions: by effectively cleaning the extreme eigenvalues of the correlation matrix, the invertibility of the matrix becomes more stable, which in turn leads to less extreme portfolio weights and improves the robustness of diversification.

Our research contributes to the literature in three key ways. First, we introduce a new method for portfolio construction that relies solely on the directional component of the correlation matrix. We show that this approach can be interpreted as a theoretically grounded and intuitive form of elementwise shrinkage on the sample correlation matrix. Our empirical analysis demonstrates that this method delivers superior out-of-sample performance compared to the 1/N rule, full correlation neglect, and well-known shrinkage estimators such as Ledoit and Wolf (2004).

Second, we contribute to the literature on FinTech and financial education by providing a practical framework that helps investors incorporate robust correlation information into their investment decisions. Rather than simply encouraging reliance on the sample correlation estimator, our findings suggest that investors should be guided to focus on the

stable, directional nature of asset relationships. This helps explain why investors might intuitively distrust complex correlation measures and instead rely on simpler heuristics, while also providing a practical pathway to improve their decision-making.

Finally, we provide a clearer understanding of why and how shrinkage estimators work. By decomposing the correlation matrix into its magnitude and directional components, we identify that the magnitude benefits most from shrinkage, while the direction contains valuable information that should be preserved. This leads to a more targeted and effective approach to managing estimation error in portfolio choice. Our simulation and empirical results further reveal an important mechanism: overestimated correlations are the primary drivers of extreme negative weights and unstable portfolio allocations, which in turn hurt out-of-sample performance. Interestingly, correlation neglect benefits from intentionally sacrificing some model accuracy, as it is more robust to unpredictable future events and estimation errors.

Our study contributes to several strands of the literature. Since our work interprets correlation neglect as an elementwise shrinkage function for correlation matrix estimation, it naturally relates to the literature on shrinkage estimators. Ledoit and Wolf (2020) provide a comprehensive review of this literature over the past two decades, highlighting how shrinkage methods improve the statistical properties of estimators and enhance portfolio optimization. Our approach is also connected to the use of thresholding operators as regularization penalties (Rothman et al., 2009; Bickel and Levina, 2008; Karoui, 2008). Unlike shrinkage methods that primarily target eigenvalues, thresholding directly regularizes individual elements of the covariance matrix. A key advantage of thresholding is that it imposes essentially no computational burden, making it attractive for problems in very high dimensions and real-time applications. We show that our directional component shares similar properties, underscoring its relevance in high-dimensional portfolio problems.

Our paper is also closely connected to the literature on Kendall's tau correlation

matrices, as we exploit the relationship between Kendall's tau and Pearson correlation to isolate the directional component of correlation. Prior studies (Edirisinghe and Zhou, 2014; Espana et al., 2024) demonstrate that Kendall's tau improves portfolio optimization by enhancing the estimation of both eigenvalues and eigenvectors in data-poor regimes, thereby yielding better covariance matrix estimators. Related research on the empirical spectral distribution of Kendall rank correlation matrices (Bandeira et al., 2017; Bao, 2019) further supports our approach and helps explain why focusing solely on the directional component of Pearson correlation leads to favorable statistical properties.

Furthermore, our work connects to Gerber et al. (2022), who propose the Gerber statistic as a robust co-movement measure for covariance matrix estimation in portfolio construction. Their approach recognizes that very large or extremely small movements can distort correlation estimates and therefore relies on a more robust co-movement measure in a mean variance setting. In our paper, we formally establish the connection between rank-based correlation and traditional Pearson correlation, analyzing which component drives the improvement in estimation accuracy and portfolio performance.

Finally, our study relates to the empirical findings of Ungeheuer and Weber (2021), who show experimentally that investors understand dependence, but not necessarily in terms of Pearson correlation. Participants behave as if they apply a simple counting heuristic based on the frequency of comovement. In our framework, we formally define this heuristic using Kendall's tau and connect it to Pearson correlation. This connection provides an intuitive explanation of why correlation neglect can act as a robust method for addressing real-world diversification tasks, especially under high estimation uncertainty.

The remainder of the paper is organized as follows. Section II reviews the correlation neglect models considered in this study. Section III presents the data and reports the empirical results. Section IV describes the simulation studies. Section V concludes.

II. Description of the Correlation Neglect Models Considered

In this section, we discuss various models that capture how investors deal with correlation coefficients and provide a brief introduction to each. Our focus is on the diversification effects that arise from different ways of measuring correlations between stocks. We use R_t to denote the N-vector of returns on the N risky assets available for investment at date t. Let $\mathbf{\Sigma}_t$ denote the corresponding $N \times N$ variance–covariance matrix of returns, with its sample estimate given by $\hat{\mathbf{\Sigma}}_t$.

To facilitate comparison across different strategies, we consider an investor who chooses the portfolio of risky assets that minimizes the variance of returns:

$$\min_{\boldsymbol{w}_t} \ \boldsymbol{w}_t^\top \mathbf{\Sigma}_t \boldsymbol{w}_t \quad \text{s.t.} \quad \boldsymbol{w}_t^\top 1 = 1, \tag{1}$$

where $\boldsymbol{w}_t$ denotes the vector of portfolio weights at time t. The optimal portfolio weights are then given by:

$$\boldsymbol{w}_t = \frac{\mathbf{\Sigma}_t^{-1}\mathbf{1}}{\mathbf{1}^\top \mathbf{\Sigma}_t^{-1}\mathbf{1}}. \tag{2}$$

Since the covariance matrix can be decomposed into volatility and correlation components, we have:

$$\mathbf{\Sigma}_t = \begin{bmatrix} \sigma_{x_1,t} & & & 0 \\ & \sigma_{x_2,t} & & \\ & & \ddots & \\ 0 & & & \sigma_{x_n,t} \end{bmatrix} \begin{bmatrix} 1 & \rho_{x_1,x_2,t} & \cdots & \rho_{x_1,x_n,t} \\ \rho_{x_2,x_1,t} & 1 & \cdots & \rho_{x_2,x_n,t} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{x_n,x_1,t} & \rho_{x_n,x_2,t} & \cdots & 1 \end{bmatrix} \begin{bmatrix} \sigma_{x_1,t} & & & 0 \\ & \sigma_{x_2,t} & & \\ & & \ddots & \\ 0 & & & \sigma_{x_n,t} \end{bmatrix}. \tag{3}$$

Our analysis mainly focuses on how different treatments of the correlation component influence asset allocation and the resulting degree of diversification.

In what follows, we first present out-of-sample results that illustrate how correlation

neglect affects portfolio decisions and then provide detailed summary statistics on the distribution of portfolio weights across stocks. We show that the rank (directional) information contained in the correlation matrix is relatively stable, whereas the magnitude component is highly sensitive and often causes risk reduction to fail out-of-sample.

To quantify this effect, we proceed in several steps. First, we analyze what happens when we completely ignore correlation information and examine how this affects portfolio weights compared to an equal-weight portfolio. Next, we consider the case where only the directional (rank-based) information is incorporated and investigate its impact on portfolio weight adjustments. For this part, we propose a *parametric approach* to isolate directional information from correlations by exploiting the connection between Kendall's τ and Pearson correlation. This method is computationally efficient and avoids the long calculation times required by the standard nonparametric Kendall's τ estimator.

Finally, we compare our results with well-known shrinkage estimators to show how portfolio weights are adjusted under different regularization methods. Importantly, our goal is not to propose an optimal shrinkage estimator that minimizes estimation error. Instead, we aim to provide a method that achieves risk reduction comparable to an equal-weight portfolio while offering a more interpretable and practical framework. This allows us to offer clearer guidance on which aspects of correlation are most relevant when constructing portfolios designed to manage risk.

A. Naive Portfolio

The naive equal-weight ("EW" or "$1/N$") strategy allocates an equal portfolio weight of $\mathbf{w}_t = 1/N$ to each of the N risky assets. This approach requires neither estimation nor optimization and disregards all information about return correlations and volatilities. For comparison with the optimized weights in Equation (2), the $1/N$ strategy can be interpreted as implicitly assuming knowledge of the covariance matrix $\mathbf{\Sigma}_t$ but imposing the restriction that all assets exhibit identical volatility, i.e., $\mathbf{\Sigma}_t = \sigma_t^2\mathbf{I}$. Under this

Table I. List of various correlation neglect models considered

#	Model	Abbreviation
Naive		
1.	1/N portfolio with rebalancing (benchmark strategy)	EW or 1/N
Full Correlation Neglect		
2.	Volatility timing (Kirby and Ostdiek, 2012)	VT
Partial Correlation Neglect		
3.	Directional correlation via Kendall's τ (parametric): $\tau = \frac{2}{\pi}\arcsin(\rho)$	K's τ^2
Naive & **Partial Correlation Neglect**		
4.	Linear shrinkage estimator (Ledoit and Wolf, 2004)	LW
No Correlation Neglect		
5.	Sample Pearson Correlation	SampleC

formulation, the investor effectively ignores both the correlation structure and cross-sectional variation in volatility. DeMiguel et al. (2009) demonstrate that this simple rule can outperform a range of optimized portfolio strategies, particularly in the presence of estimation error.

B. Full Correlation Neglect

Kirby and Ostdiek (2012) identify the correlation estimator as a primary driver of high turnover and transaction costs in variance-based asset allocation strategies, frictions that can easily erode the theoretical benefits of portfolio optimization. They propose a simple yet powerful solution: setting all pairwise correlations between excess risky-asset returns to zero. This approach results in a diagonal covariance matrix, $\Sigma_{\text{total}} = \text{diag}(\sigma_1^2, \sigma_2^2, \ldots, \sigma_n^2)$, and a strategy they term volatility timing.

The rationale behind this method is to treat the zeroing out of the off-diagonal elements of the estimated covariance matrix, $\widehat{\mathbf{\Sigma}}$ as an aggressive form of shrinkage. While this intentionally discards correlation information, it drastically simplifies the estimation problem by reducing the number of parameters by $N(N-1)/2$. The central insight is that the significant reduction in estimation risk can outweigh the information loss, leading to more stable portfolio weights and reduced turnover. Consequently, this volatility

timing strategy can outperform naive diversification, even in the presence of significant transaction costs.

C. *Full Correlation Neglect*

Kirby and Ostdiek (2012) identify the correlation estimator as a primary driver of high turnover and transaction costs in variance-based asset allocation strategies, frictions that can erode the theoretical benefits of portfolio optimization. They propose a simple yet powerful solution: setting all pairwise correlations between risky-asset returns to zero. This approach results in a diagonal covariance matrix,

$$\widehat{\mathbf{\Sigma}}_t = \text{diag}(\hat{\sigma}^2_{1,t}, \hat{\sigma}^2_{2,t}, \ldots, \hat{\sigma}^2_{N,t}),$$

and a strategy termed *volatility timing*.

The rationale behind this method is to treat the zeroing out of the off-diagonal elements of $\widehat{\mathbf{\Sigma}}_t$ as an aggressive form of shrinkage. While this discards correlation information, it simplifies estimation by reducing the number of parameters from $N(N-1)/2$ to N. The key insight is that the significant reduction in estimation risk can outweigh the information loss, leading to more stable portfolio weights and reduced turnover.

D. *Partial Correlation Neglect*

In this section, we introduce a correlation measure that isolates the *directional co-movement* of stocks while deliberately ignoring the *magnitude* of their movements. We employ Kendall's τ, a nonparametric statistic that quantifies the ordinal association between two variables by comparing their pairwise rankings. Intuitively, Kendall's Tau captures how often two assets move in the same direction versus in opposite directions. This perspective allows us to view the standard Pearson correlation as comprising two components: a directional component (captured by Kendall's τ) and a magnitude com-

ponent ($\rho - \tau$). Our analysis focuses on how the directional component affects portfolio weights and risk reduction.

For two jointly distributed random variables (X, Y) with independent copies (X_1, Y_1) and (X_2, Y_2), the population Kendall's τ is

$$\tau = \mathbb{P}\big((X_1 - X_2)(Y_1 - Y_2) > 0\big) - \mathbb{P}\big((X_1 - X_2)(Y_1 - Y_2) < 0\big),$$

which measures the difference in probabilities that two pairs are concordant versus discordant.

Given a bivariate sample $(X_1, Y_1), \ldots, (X_n, Y_n)$, the sample Kendall's Tau is

$$\hat{\tau} = \frac{2}{n(n-1)} \sum_{i<j} \operatorname{sgn}(X_i - X_j) \operatorname{sgn}(Y_i - Y_j),$$

where

$$\operatorname{sgn}(x) = \begin{cases} +1 & \text{if } x > 0, \\ 0 & \text{if } x = 0, \\ -1 & \text{if } x < 0. \end{cases}$$

A pair of observations (X_i, Y_i) and (X_j, Y_j) is concordant if the assets move in the same direction (e.g., $X_i > X_j$ and $Y_i > Y_j$), and discordant if they move in opposite directions. The statistic $\hat{\tau}$ ranges in $[-1, 1]$, with $\hat{\tau} = 1$ indicating perfect concordance, $\hat{\tau} = -1$ perfect discordance, and $\hat{\tau} = 0$ no association.

The Kendall's Tau correlation matrix $\mathbf{T}$ generalizes this to multiple assets, with each entry $\mathbf{T}_{ij}$ representing the τ between the i-th and j-th variables.

Connection to Pearson Correlation

A natural question is how this directional correlation relates to the standard Pearson correlation. In financial applications, stock returns are often modeled as elliptically distributed (e.g., Normal or Student's t). In this setting, the two measures are linked by the

following proposition:

PROPOSITION 1 (Kendall's τ and Pearson Correlation): *Let* $\mathbf{X} = (X_1, X_2)$ *follow an elliptical distribution. Then Kendall's Tau and Pearson correlation* ρ *satisfy*

$$\tau = \frac{2}{\pi} \arcsin(\rho).$$

This proposition follows directly from Theorem 2 in Lindskog et al. (2003). This relation implies that focusing solely on the directional component corresponds to a *non-linear shrinkage of Pearson's correlation* toward zero. The transformation is mild for low correlations but becomes more pronounced as $|\rho|$ increases.

Figure 1 illustrates this effect, showing how extreme correlations are shrunk more strongly under Kendall's Tau. Accordingly, we define the estimator

$$\hat{\tau}_{ij} = \frac{2}{\pi} \arcsin(\hat{\rho}_{ij}),$$

which isolates the directional information in the Pearson correlation $\hat{\rho}_{ij}$ while suppressing its magnitude. The subsequent analysis investigates whether this transformation enhances portfolio diversification.

E. Naive & Partial Correlation Neglect

Ledoit and Wolf (2004) and Ledoit and Wolf (2003) propose a widely used linear shrinkage estimator of the covariance matrix:

$$\widehat{\mathbf{\Sigma}}_{\mathrm{LW},t} = (1 - \lambda)\bar{\sigma}^2 \mathbf{I} + \lambda \widehat{\mathbf{\Sigma}}_t,$$

where $\bar{\sigma}^2$ denotes the cross-sectional average of the sample variances. This estimator shrinks the sample covariance matrix toward a scaled identity matrix, thereby reducing

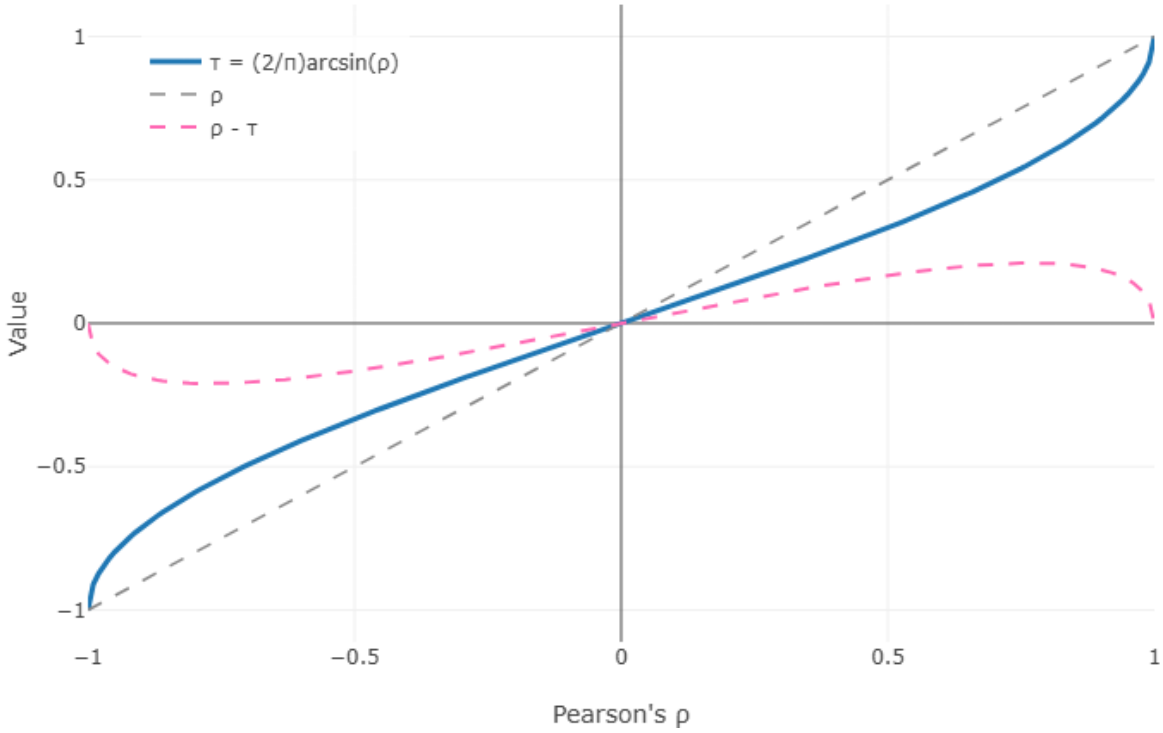

Figure 1. Relationship between Pearson's ρ and Kendall's τ

sensitivity to noisy correlation estimates. It can be interpreted as a linear combination of an equal-weighted portfolio (imposing identical volatilities and zero correlations) and the sample-based estimator. In this sense, it partially ignores correlation structure while simultaneously enforcing a homoskedasticity constraint, which improves estimation stability in high-dimensional settings.

F. No Correlation Neglect

To implement this model, we adopt the classic *plug-in* approach, whereby the optimization problem in Equation (2) is solved using the sample estimates of the covariance matrix, denoted by $\hat{\Sigma}$. We refer to this strategy as the *sample-based minimum variance portfolio*. While this approach is commonly used to capture the correlation structure among assets, our analysis highlights that the correlation component plays a dominant role in driving high portfolio turnover. These findings suggest that caution is warranted

when relying on sample correlations in diversification-based strategies, particularly in the presence of estimation error and transaction costs.

III. Empirical Results

A. Methodology for Evaluating Performance

Our goal is to study the influence of correlation neglect on real asset allocation performance across different setups, and to demonstrate how our proposed method can be applied to empirical data to derive insights from observed market behavior. The analysis focuses on two main objectives: (i) evaluating the performance of the Global Minimum Variance Portfolio (GMVP) under different types of correlation neglect and (ii) examining the impact ow correlation neglect affects portfolio weight dynamics.

We use the daily data obtained from the Center for Research in Security Prices (CRSP), covering the period from 01/01/1995 to 12/29/2023. The empirical analysis follows a *rolling-sample* approach. Consistent with common practice, 21 consecutive trading days are treated as one "month." The out-of-sample period spans from 01/14/2000 to 12/29/2023, resulting in a total of 287 "months" (6,027 trading days). Portfolios are rebalanced monthly, with the number of shares held remaining constant within each "month,â implying no transaction costs during this period. We denote the investment dates by $t = 1, ..., 287$.

For each combination (T, N), the investment universe is defined as the set of N stocks with complete return histories over both the most recent T trading days and the subsequent 21 trading days. These return series are used to estimate the covariance matrix, which then using as imput for portfolio optimization. At each investment date t, returns from the previous T days are used to estimate the covariance matrix $(\widehat{\mathbf{\Sigma}}_t)$ required to implement a given strategy. The estimated parameters determine the portfolio weights $(\widehat{\boldsymbol{w}}_t)$ for month $t+1$. This rolling window advances monthly, adding 21 new observations

and discarding the oldest 21. The outcome of this procedure is a time series of 6,027 out-of-sample daily returns generated by each portfolio strategy.

Given the out-of-sample returns, We evaluate portfolio performance using five metrics:

1. Out-of-sample annualized volatility is calculated as:

$$\hat{\sigma}_{\text{annual}} = \hat{\sigma}_{\text{daily}} \times \sqrt{252}, \tag{4}$$

   where the $\hat{\sigma}_{\text{daily}}$ is the daily sample standard deviation. To test whether the volatility of two strategies is statistically different, we also compute the p-value of the difference, using the approach suggested by Ledoit and Wolf (2008).

2. The out-of-sample Sharpe ratio of each portfolio is defined as the sample mean of excess returns, $\hat{\mu}_p$, divided by portfolio volatility, $\hat{\sigma}_p$:

$$\widehat{\text{SR}} = \frac{\hat{\mu}_p}{\hat{\sigma}_p}. \tag{5}$$

   To test whether the Sharpe ratios of two strategies are statistically different, we also compute the p-value of the difference, using the approach suggested by Ledoit and Wolf (2008).

3. Out-of-sample maximum drawdown (MDD), the largest peak-to-trough decline in portfolio value, is computed as:

$$\text{MDD} = \max_{\tau \in [0,T]} \left[\frac{\sup_{t \in [0,\tau]} P_t - P_\tau}{\sup_{t \in [0,\tau]} P_t} \right] \tag{6}$$

   where P_t denotes the portfolio value at time t.

4. The certainty-equivalent (CEQ) return represents the risk-free rate that an investor

would accept instead of adopting a given portfolio strategy:

$$\widehat{\text{CEQ}} = \widehat{\mu}_p - \frac{\gamma}{2}\widehat{\sigma}_p^2, \tag{7}$$

where γ is the coefficient of relative risk aversion. We report results for $\gamma = 4$; robustness checks using alternative values are provided in the appendix. To test whether the CEQ returns from two strategies are statistically different, we also compute the p-value of the difference, using the approach suggested by DeMiguel et al. (2009).

5. Portfolio turnover measures trading activity required to maintain a strategy:

$$\text{Turnover} = \sum_{j=1}^{N} |\hat{w}_{j,t+1} - \hat{w}_{j,t^+}| \tag{8}$$

where:

- $\hat{w}_{k,j,t}$ is the target weight in asset j at time t
- $\hat{w}_{k,j,t^+}$ denotes the pre-rebalancing weight at $t+1$ (before trading)
- $\hat{w}_{k,j,t+1}$ is the post-rebalancing weight at $t+1$ (after trading)

The average turnover is computed as the mean of turnover across all investment dates.

B. Empirical Results

Our analysis begins by investigating a challenging but realistic scenario for asset allocation: a "short sample" setting where the number of time-series observations is only slightly larger than the number of assets ($T/N = 1.05$). This environment is notorious for producing severe estimation error in the covariance matrix.

Table II. Performance Comparison of Different Portfolio Strategies

Method	Standard deviation	Sharpe ratio	Maximum drawdown	CE	Avg. Turnover
N=100, T=105					
K's τ^2	12.35	0.50	31.49	4.77	1.3655
EW	19.32	0.36	55.45	1.13	0.1036
	(0.00)	(0.39)		(0.11)	
VT	16.48	0.43	45.71	3.26	0.1672
	(0.00)	(0.59)		(0.25)	
LW	13.53	0.45	39.98	4.07	2.2976
	(0.00)	(0.44)		(0.22)	
SampleC	49.24	0.28	81.32	-32.92	29.1090
	(0.00)	(0.37)		(0.00)	
N=500, T=525					
K's τ^2	10.24	0.57	29.02	5.33	1.1235
EW	20.20	0.46	56.80	2.78	0.1005
	(0.00)	(0.57)		(0.22)	
VT	17.39	0.54	49.74	5.02	0.0897
	(0.00)	(0.88)		(0.45)	
LW	11.88	0.42	33.85	3.74	3.3802
	(0.00)	(0.14)		(0.10)	
SampleC	36.03	0.17	81.26	-18.08	36.1090
	(0.00)	(0.10)		(0.00)	
N=1000, T=1050					
K's τ^2	9.30	0.81	26.99	7.45	0.9819
EW	21.36	0.47	57.19	2.62	0.1036
	(0.00)	(0.09)		(0.09)	
VT	18.41	0.57	51.35	5.26	0.0822
	(0.00)	(0.17)		(0.23)	
LW	10.84	0.64	28.04	6.23	3.4244
	(0.00)	(0.12)		(0.14)	
SampleC	29.85	0.34	65.68	-6.00	33.6012
	(0.00)	(0.05)		(0.01)	

The portfolio constructed using the raw Sample Covariance Matrix (SampleC) performs disastrously. It exhibits extremely high volatility, large drawdowns, and deeply negative certainty-equivalent (CE) returns, implying that a risk-averse investor would pay to avoid this strategy. Moreover, its extremely high turnover indicates unstable and

erratic weight reallocations from month to month. This confirms that directly inverting the sample covariance matrix in high-dimensional settings is essentially a recipe for *"error maximization"*: the optimizer aggressively chases noise rather than meaningful signal.

Strategies that neglect correlation information perform far better. The Equally Weighted (EW) portfolio provides a reasonable baseline, but the Volatility-Targeting (VT) strategy — which uses only the diagonal elements of the covariance matrix — delivers a clear improvement. By simply weighting assets inversely to their variance, the VT portfolio achieves substantially lower risk and a higher Sharpe ratio.

This demonstrates that even a naive form of "correlation neglect" — ignoring off-diagonal elements entirely — can be much more effective than relying on noisy correlation estimates. From statistical tests, we find that partially incorporating correlation information significantly improves portfolio risk control. For lower-dimensional settings, the improvement in the Sharpe ratio is modest and often not statistically significant, mainly because lower risk sometimes comes at the cost of reduced returns. However, as the stock universe expands, the benefit of correlation information becomes more pronounced, and the improvement in the Sharpe ratio becomes statistically significant.

The most sophisticated methods deliver the strongest results. Both the Ledoit-Wolf (LW) shrinkage estimator and our proposed rank-based Kendall's τ (K's τ) strategies significantly outperform simpler heuristics:

- LW achieves lower volatility than EW and VT but suffers from *very high turnover*, indicating frequent and costly portfolio adjustments.
- Our **K's** τ^2 strategy, which isolates the *most robust directional component* of correlation, consistently delivers one of the highest Sharpe ratios and CE returns while maintaining a much lower turnover than LW.

Statistical tests confirm that K's τ^2 portfolios extract stable and valuable information from the correlation structure. Among all methods, the closest competitor is VT, which

often shows larger p-values in tests of Sharpe ratio differences. This suggests that volatility alone carries strong predictive information for portfolio construction, and ignoring correlations entirely still leads to relatively good outcomes. However, volatility-based approaches sacrifice some diversification potential.

The strength of K's τ^2 lies in regularizing the correlation matrix, taming noise while preserving useful signals. By focusing on directional dependence, it balances signal extraction with portfolio stability. While correlation-based methods generally produce higher turnover than the $1/N$ strategy, they significantly improve risk-adjusted performance.

Next, we explore whether simply increasing the length of the estimation window can overcome estimation challenges. For a fixed portfolio size of $N = 500$, we expand the sample from $T = 525$ to $T = 1250$.

The findings in Table III reveal a nuanced picture:

- As T increases, the performance of the SampleC strategy *improves dramatically*. By $T = 1250$, its volatility and Sharpe ratio become respectable.
- However, consistent with the findings of DeMiguel et al. (2009), SampleC still fails to outperform simpler, more robust strategies. Achieving reliable results with SampleC requires unrealistically long time series.

The key insight from Table III is the remarkable stability of robust strategies. The performance of VT and especially K's τ^2 remains exceptionally consistent across different time-series lengths. While longer samples slightly reduce turnover for K's τ^2, the overall risk-return profile remains stable.

Interestingly, VT and K's τ^2 deliver statistically similar Sharpe ratios and CE returns, but K's τ^2 achieves lower volatility. This result arises because volatility alone provides strong signals about the risk-return trade-off, but correlation information further enhances volatility control. VT portfolios exhibit slightly higher volatility but achieve comparable returns, while K's τ^2 extracts additional diversification benefits.

Table III. Performance Comparison of Different Portfolio Strategies

Method	Standard deviation	Sharpe ratio	Maximum drawdown	CE	Avg. Turnover
N=500, T=525					
K's τ^2	10.24	0.57	29.02	5.33	1.1235
EW	20.20	0.46	56.80	2.78	0.1005
	(0.00)	(0.57)		(0.22)	
VT	17.39	0.54	49.74	5.02	0.0897
	(0.00)	(0.88)		(0.45)	
LW	11.88	0.42	33.85	3.74	3.3802
	(0.00)	(0.14)		(0.10)	
SampleC	36.03	0.17	81.26	-18.08	36.1090
	(0.00)	(0.10)		(0.00)	
N=500, T=600					
K's τ^2	10.26	0.62	28.54	5.90	1.0463
EW	20.20	0.46	56.80	2.78	0.1005
	(0.00)	(0.40)		(0.18)	
VT	17.42	0.55	49.82	5.06	0.0886
	(0.00)	(0.64)		(0.38)	
LW	11.72	0.50	31.29	4.78	3.2095
	(0.00)	(0.26)		(0.18)	
SampleC	19.58	0.23	62.67	-1.63	12.3525
	(0.00)	(0.05)		(0.01)	
N=500, T=1250					
K's τ^2	10.54	0.67	29.46	6.47	0.7117
EW	20.20	0.46	56.80	2.78	0.1005
	(0.00)	(0.26)		(0.13)	
VT	17.60	0.55	50.72	5.02	0.0846
	(0.00)	(0.43)		(0.29)	
LW	11.11	0.50	29.02	4.71	1.8859
	(0.00)	(0.04)		(0.03)	
SampleC	11.70	0.44	30.83	4.01	2.5060
	(0.00)	(0.03)		(0.02)	

Even though SampleC improves with longer data windows, it still tends to overestimate correlations and overweight certain stocks, which ultimately hurts returns. Counterintuitively, ignoring noisy correlation information often produces better portfolios than using SampleC directly.

This analysis highlights a crucial point:

- When sample sizes are small, naive inversion of the sample covariance matrix leads to unstable and poor-performing portfolios.
- Simple correlation-neglect strategies, such as VT, already deliver significant benefits.
- Among advanced methods, K's τ^2 stands out by isolating a stable directional signal, yielding strong performance with lower turnover and greater robustness.

For investors who are uncertain whether they have "enough" data, focusing on directional correlation provides a reliable and effective approach to portfolio construction.

C. The Effect of Correlation Neglect on Portfolio Weight

In this section, we conduct a detailed analysis of the summary statistics for the portfolio weights generated by various allocation strategies. Building on the preceding finding that neglecting the correlation component can mitigate estimation error and improve out-of-sample performance, we now seek to understand the mechanism through which this improvement is achieved. Specifically, we investigate how different treatments of the covariance matrix influence the resulting portfolio weights. The central hypothesis is that estimation errors, particularly in the off-diagonal elements (correlations) of the sample covariance matrix, induce extreme long and short positions, which in turn degrade portfolio performance. To characterize the portfolios, we compute five key summary statistics:

1. **Average Minimum Weight**: The mean of the lowest weight assigned to any single asset across all rebalancing periods.
2. **Average Maximum Weight**: The mean of the highest weight assigned to any single asset across all rebalancing periods.

3. **Average Percentage of Positive Weights**: The average proportion of assets held in long positions.

4. **Herfindahl Index**: A measure of weight concentration, calculated as the sum of squared portfolio weights, $H = \sum_{i=1}^{N} w_i^2$. A higher index indicates a more concentrated portfolio.

5. **Average Short Position Size**: The absolute sum of all negative weights, representing the total magnitude of short-selling.

These statistics are averaged across the 287 monthly rebalancing dates in our out-of-sample period. They collectively provide a comprehensive picture of the portfolio's sparsity, diversification, and reliance on leverage through short-selling.

Table IV. Summary Statistics of Weights of Different Portfolio Strategies

	Avg Min	Avg Max	Avg.. Pos. w	Avg. Conc.	Avg. Size of Short
N=100, T=105					
K's τ^2	-0.0562	0.1703	57.88	0.1339	0.6941
EW	0.0100	0.0100	100.00	0.0100	0.0000
VT	0.0016	0.0327	100.00	0.0136	0.0000
LW	-0.0970	0.1413	59.31	0.2085	1.2100
SampleC	-1.1021	1.1661	51.41	12.2742	10.2765
N=500, T=525					
K's τ^2	-0.0283	0.0729	52.64	0.0662	1.4705
EW	0.0020	0.0020	100.00	0.0020	0.0000
VT	0.0002	0.0073	100.00	0.0027	0.0000
LW	-0.0567	0.0748	53.37	0.1923	3.2240
SampleC	-0.4595	0.4724	50.79	4.5141	15.9578
N=1000, T=1050					
K's τ^2	-0.0193	0.0559	51.24	0.0462	1.7550
EW	0.0010	0.0010	100.00	0.0010	0.0000
VT	0.0001	0.0042	100.00	0.0014	0.0000
LW	-0.0378	0.0557	52.58	0.1367	3.8644
SampleC	-0.2571	0.2619	50.65	2.3843	16.4928

Table IV reports these results for the short-sample case where the time-series length

T and the cross-sectional dimension N satisfy $T/N = 1.05$. This data-limited scenario is common in practice and magnifies the impact of estimation error. The results, presented in Table IV, are striking. The findings in Table IV reveal stark differences in the characteristics of the portfolios constructed using varying levels of correlation neglect.

The portfolio based on the Sample Covariance Matrix (SampleC), which directly inverts the sample covariance matrix, exhibits severe instability. The average minimum and maximum weights are extreme, often exceeding 100% in magnitude. This indicates that the portfolio takes massive, highly leveraged positions in a few assets. This is confirmed by the astronomical Herfindahl Index (12.27), signifying extreme concentration, and a total short position that is over 10 times the portfolio's net value. Concurrently, the average short position size is substantial, highlighting an aggressive and risky allocation strategy driven by spurious correlations in the data. These extreme weights are a classic symptom of estimation error in an ill-conditioned covariance matrix, where small changes in input parameters lead to dramatic shifts in the optimal portfolio. This behavior is a classic symptom of "error maximization," where the optimizer aggressively exploits spurious correlations found in noisy data.

In stark contrast, the Volatility-Targeting (VT) strategy, which completely ignores correlation information and weights assets based solely on their volatility, produces highly stable weights. The weights are constrained, with a complete absence of short positions. This approach is analogous to the equally weighted (EW) portfolio in its simplicity and robustness. By ignoring the noisy correlation structure, the VT strategy avoids the optimization error that plagues the sample-based approach, leading to the superior out-of-sample performance documented in Table II.

The strategy based on Kendall's τ^2, which utilizes only the directional component of correlation, offers a compelling intermediate solution. This method implicitly restricts the magnitude of short positions while still permitting significant long positions. The asymmetric effect on long and short weights is evident from the summary statistics.

The concentration ratio, as measured by the Herfindahl index, is significantly reduced compared to the sample covariance strategy. This behavior can be interpreted as a form of implicit regularization or shrinkage. By focusing on the more robust, rank-based measure of co-movement, the model filters out much of the estimation noise associated with linear correlation, thereby preventing the optimization from assigning extreme negative weights.

This shrinkage-like effect is also observed in the portfolio constructed using the Ledoit-Wolf (LW) shrinkage estimator. The LW method explicitly regularizes the covariance matrix by shrinking the sample estimates towards a more structured target. This process effectively dampens extreme weights, reduces short-selling intensity, and increases portfolio diversification, as evidenced by the lower Herfindahl index. Both the Kendall's τ^2 and LW approaches demonstrate that moderating the influence of estimated correlations is critical for achieving well-diversified and robust portfolios.

Furthermore, our analysis shows that as the number of assets in the portfolio increases, the limitations of the sample covariance matrix become more pronounced. Although larger portfolios based on the sample matrix exhibit slightly less extreme maximum weights, they tend to concentrate leverage into a single substantial short position. This indicates that as portfolio dimensionality grows, the risk of estimation errors leading to dangerously concentrated short positions also increases.

This analysis reveals a key channel through which estimation error harms portfolio performance: the overestimation of correlation coefficients in finite samples induces extreme negative weights, leading to poorly diversified and unstable portfolios. Strategies that neglect correlation (like VT) or employ robust estimators that implicitly or explicitly shrink the correlation structure (like Kendall's τ^2 and Ledoit-Wolf) effectively mitigate this problem. This finding is consistent with the work of Jagannathan and Ma (2003), who demonstrate that imposing no-short-sale constraints can be an effective remedy for estimation error in practice. Our results suggest that correlation neglect acts similarly to such a constraint, primarily by preventing the large, erroneous short positions that arise

from noisy correlation estimates. Ultimately, we find that the directional component of correlation is a more stable and reliable input for portfolio decisions, especially when data is limited. By focusing on this robust feature, investors can construct better-diversified portfolios and enhance out-of-sample performance.

Table V. Summary Statistics of Weights of Different Portfolio Strategies

	Avg. Min	Avg. Max	Avg. Pos. w	Avg. Conc.	Avg. Size of Short
N=500, T=525					
K's τ^2	-0.0283	0.0729	52.64	0.0662	1.4705
EW	0.0020	0.0020	100.00	0.0020	0.0000
VT	0.0002	0.0073	100.00	0.0027	0.0000
LW	-0.0567	0.0748	53.37	0.1923	3.2240
SampleC	-0.4595	0.4724	50.79	4.5141	15.9578
N=500, T=600					
K's τ^2	-0.0287	0.0738	52.74	0.0659	1.4610
EW	0.0020	0.0020	100.00	0.0020	0.0000
VT	0.0002	0.0072	100.00	0.0027	0.0000
LW	-0.0586	0.0792	53.33	0.1978	3.2592
SampleC	-0.2398	0.2528	51.21	1.1949	8.0010
N=500, T=1250					
K's τ^2	-0.0304	0.0739	53.20	0.0621	1.3651
EW	0.0020	0.0020	100.00	0.0020	0.0000
VT	0.0002	0.0068	100.00	0.0027	0.0000
LW	-0.0699	0.0923	53.19	0.1761	2.8480
SampleC	-0.1946	0.2043	52.57	0.3576	3.5903

We next investigate whether simply having more data can solve the problem. In Table V, we fix the number of assets at N=500 and increase the time-series length from T=525 ($T/N = 1$) to T=1250 ($T/N = 2.5$).

As expected, increasing the amount of data helps the SampleC strategy. The average extreme weights decrease, and the concentration index drops from 4.51 to a more moderate 0.36. However, even with 2.5 times more data points than assets, the portfolio remains significantly more concentrated and leveraged compared to those constructed using alternative correlation neglect methods. This highlights that simply relying on "brute-force"

longer samples is often insufficient to resolve the underlying estimation issues.

The most important result from this table is the remarkable stability of the K's τ^2 strategies. Across all data lengths ($T = 525, 600, 1250$), the summary statistics for the K's τ^2 portfolios remain almost unchanged. This provides strong evidence supporting our central thesis: the directional component of correlation is a stable signal that is largely insensitive to the noise induced by limited data. By isolating this signal, we develop a portfolio construction method that remains robust in both high-dimensional and short-sample environments.

Our analysis of portfolio weights reveals the precise channel through which estimation error harms performance: noisy, overestimated sample correlations lead to extreme and heavily leveraged negative weights. This finding is consistent with Jagannathan and Ma (2003), who show that no-short-sale constraints can improve performance by mitigating the impact of estimation error.

Strategies that neglect correlation act as a blunt but effective tool, similar to a no-short-sale constraint. However, our proposed method, based on Kendall's τ, offers a more nuanced solution. By retaining the stable directional information in correlations while discarding the noisy magnitude component, it effectively regularizes the portfolio, preventing extreme weights and enhancing diversification. This approach provides a robust and reliable method for portfolio construction, especially when data is limited.

IV. Simulation Studies

The results in Section III demonstrate that correlation neglect influences diversification in real data. To further investigate this mechanism, we employ simulated data to analyze how the performance of the strategies considered in our empirical analysis varies with the number of assets (N) and the length of the estimation window (T). The main advantage of using simulated data is that their economic and statistical properties are fully specified. Specifically, we generate returns from a simple single-factor model, assum-

ing that they are independently and identically distributed over time and follow a normal distribution. Since the normal distribution is elliptical, this framework provides a clearer understanding of how correlation neglect affects portfolio diversification. Moreover, it ensures that the results are not confounded by empirical irregularities, such as small-firm effects, calendar effects, momentum, mean reversion, or the fat tails typically observed in real financial data.

Our approach for simulating returns, as well as our choice of parameter values, follows (Craig MacKinlay and Pástor, 2000; DeMiguel et al., 2009). We assume that the market consists of a risk-free asset and N risky assets, which are driven by K underlying factors. The excess returns of the remaining $N - K$ risky assets are generated according to the factor model:$R_{a,t} = \alpha + BR_{b,t} + \epsilon_t$, where $R_{a,t}$ is the vector of excess asset returns, α is the vector of mispricing coefficients, B is the factor loadings matrix, $R_{b,t}$ is the vector of excess returns on the factor portfolios, $R_b \sim N(\mu_b, \Sigma_b)$, and ϵ_t is the vector of idiosyncratic noise, $\epsilon \sim N(0, \Sigma_\epsilon)$, which is assumed to be independent of the factor portfolios.

Following the setup of DeMiguel et al. (2009), we assume that the risk-free rate follows a normal distribution with an annual mean of 2% and a standard deviation of 2%. We further assume that there is only one factor ($K = 1$), whose annual excess return has a mean of 8% and a standard deviation of 16%. The mispricing term α is set to zero, and the factor loadings B are evenly distributed between 0.5 and 1.5. The variance-covariance matrix of the idiosyncratic noise, Σ_ϵ, is assumed to be diagonal, with elements drawn from a uniform distribution on $[0.10, 0.30]$, resulting in a cross-sectional average annual idiosyncratic volatility of 20%. We also report results for different volatility levels to examine their impact on portfolio weights.

We consider scenarios with the number of assets set to $N \in \{100, 500, 1000\}$ and estimation window lengths $T \in \{505, 1250, 2500\}$, corresponding to short, medium, and long data setups, respectively. For each configuration, Monte Carlo sampling is used to generate daily return data matching the out-of-sample length in the empirical section.

Table VI. Standard deviation for simulated data

	$N = 50$			$N = 100$			$N = 500$		
Strategy	$T = 505$	$T = 1250$	$T = 2500$	$T = 505$	$T = 1250$	$T = 2500$	$T = 505$	$T = 1250$	$T = 2500$
Noise Level					20%				
K's τ^2	8.1157	8.0386	8.0157	6.5545	6.4171	6.3712	3.4955	3.2918	3.1945
True	7.0302	7.0302	7.0302	5.6165	5.6165	5.6165	2.7232	2.7232	2.7232
LW	7.6104	7.2821	7.1539	6.5176	6.0047	5.8010	5.2705	3.6121	3.1337
SampleC	7.3929	7.1947	7.1213	6.2989	5.8620	5.7270	29.9687	3.5149	3.0346
Noise Level					50%				
K's τ^2	14.7048	14.5270	14.4541	13.2714	12.9257	12.7864	8.7571	8.0828	7.7909
True	13.3294	13.3294	13.3294	11.7035	11.7035	11.7035	7.0109	7.0109	7.0109
LW	15.4019	14.7944	14.3222	13.7530	13.1116	12.7011	9.7165	8.9103	8.1472
SampleC	14.0484	13.6540	13.5118	13.1537	12.2475	11.9414	87.6054	9.0725	7.8357
Noise Level					80%				
K's τ^2	16.2519	16.0396	15.9552	15.7439	15.3083	15.1374	12.6712	11.6559	11.2194
True	14.9434	14.9434	14.9434	13.9758	13.9758	13.9758	10.0217	10.0217	10.0217
LW	19.0862	18.6778	18.2467	17.0460	16.6083	16.2984	12.6521	11.8875	11.4796
SampleC	15.7809	15.3057	15.1426	15.7181	14.6241	14.2653	360.4017	12.9470	11.2015

From the simulation results, we observe that the sample correlation coefficient produces estimates close to the true diversification outcome only when idiosyncratic volatility is low and sufficiently long time-series samples are available-roughly ten times the cross-sectional dimension. However, once idiosyncratic volatility increases relative to factor variance (as is typically the case in financial markets with a low signal-to-noise ratio), the performance of the sample correlation coefficient deteriorates significantly. In such settings, much longer data histories are required to obtain reliable results. When the available sample is short or the cross-sectional dimension is large, the estimates become extremely noisy, leading to unstable and highly concentrated portfolio allocations.

By contrast, the directional component of correlation remains remarkably stable across different sample lengths and cross-sectional dimensions. Although it necessarily discards part of the correlation structure and thereby introduces specification error, it consistently yields results close to the true benchmark. As illustrated in Figures VII and VIII (provided in the appendix for space considerations), this stability translates into more diversified portfolio weights and substantially fewer short positions. Conceptually, this reflects a trade-off between specification error and sampling error: by ignoring part of the correlation information, we incur specification error but benefit from more stable

portfolios that are robust to sample length and dimensionality.

A comparison with the Ledoit-Wolf (LW) shrinkage estimator reveals an interesting distinction. Shrinkage methods stabilize the covariance estimator but shrink both variances and correlations, which leads to lower performance relative to the directional component. This finding reinforces our central result: the correlation matrix is the primary source of sampling error and extreme weights, more so than variance estimates. Correlation neglect therefore acts as an implicit shrinkage mechanism. By down-weighting correlations, it regularizes extreme portfolio weights and reduces reliance on short positions, ultimately producing lower-risk portfolios, consistent with the findings of Jagannathan and Ma (2003).

Finally, to further illustrate the role of the estimation window size, we provide additional plots in the appendix showing the portfolio risk of different methods as T increases. The results reveal that both K's τ^2 and LW remain relatively stable across different sample sizes, whereas SampleC performs poorly in small samples and only begins to outperform K's τ^2 when the sample size reaches approximately five to ten times the cross-sectional dimension. This finding confirms that extremely long histories are required for SampleC to achieve its true performance potential, consistent with insights from previous literature (DeMiguel et al., 2009).

Taken together, the simulation study offers a deeper understanding of correlation neglect. It can be interpreted as a form of shrinkage that, while introducing specification error, enhances portfolio performance by mitigating extreme weights and reducing short exposures. The exceptional stability of the directional component provides a practical guideline for practitioners seeking to incorporate correlation information under limited data availability or in noisy market environments.

A. *Empirical Spectral Distribution*

Modern Portfolio Theory (MPT), pioneered by Markowitz, provides a theoretically elegant framework for constructing optimal portfolios based on estimated means, variances, and correlations of asset returns. However, its application in practice has been limited by a fundamental challenge: estimating the correlation matrix accurately in high-dimensional, data-poor environments.

As our previous simulation results demonstrate, obtaining reliable correlation estimates typically requires the sample size T to be at least five to ten times larger than the number of assets N. In real-world applications, where the cross-sectional dimension is large and time-series data are limited, such conditions rarely hold. Consequently, the sample correlation matrix often suffers from severe estimation noise, leading to unstable portfolio weights and poor out-of-sample performance.

To better understand the effects of data limitations on correlation estimation, researchers have increasingly relied on Random Matrix Theory (RMT; Marčenko and Pastur (1967)), which provides powerful insights into the spectral properties of sample covariance and correlation matrices. RMT characterizes the limiting spectral distribution of eigenvalues under the null hypothesis of independent and identically distributed (i.i.d.) returns. This theoretical benchmark allows us to distinguish between eigenvalue components dominated by signal and those dominated by noise, thereby guiding the development of methods for regularizing or denoising empirical correlation matrices.

Building on these insights, a rich body of literature has emerged to improve correlation estimation in high-dimensional settings. Techniques such as eigenvalue cleaning, shrinkage, and factor-based approaches have been proposed and extensively studied. For comprehensive reviews, see Bun et al. (2016) and Bouchaud and Potters (2009). In particular, many shrinkage estimators, including the widely used Ledoit-Wolf methods, are grounded in RMT principles and provide theoretically justified improvements over the naive sample estimator (Ledoit and Wolf (2020)).

In this section, we investigate how our proposed method affects the spectral distribution of eigenvalues of the correlation matrix. By comparing our approach with standard shrinkage and eigenvalue-cleaning techniques, we highlight its effectiveness in reducing estimation noise while preserving essential dependency structures among assets. This analysis provides a deeper understanding of the mechanism through which our method improves portfolio diversification and risk management in high-dimensional environments.

A.1. The Marčenko–Pastur Law and Eigenvalue Distributions

The resulting cross-sectional distribution of the sample eigenvalues is known as the Marčenko–Pastur (MP) law (Marčenko and Pastur, 1967), which provides a theoretical benchmark for understanding the behavior of sample eigenvalues in high-dimensional settings. Consider a data matrix $\mathbf{X} \in \mathbb{R}^{T \times N}$, where returns are i.i.d. with zero mean and variance 1. Define the sample correlation matrix as:

$$\mathbf{S} = \frac{1}{T}\mathbf{X}^\top \mathbf{X}.$$

When both N and T grow to infinity with their ratio converging to $\gamma = N/T$, the empirical spectral distribution of $\mathbf{S}$ converges almost surely to the MP distribution with density:

$$f_{\text{MP}}(\lambda) = \begin{cases} \dfrac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{2\pi\gamma\lambda}, & \lambda \in [\lambda_-, \lambda_+], \\ 0, & \text{otherwise}, \end{cases}$$

where the lower and upper bounds of the support are:

$$\lambda_\pm = \left(1 \pm \sqrt{\gamma}\right)^2.$$

The MP law therefore characterizes the eigenvalue spectrum of the sample correlation matrix under the null hypothesis of no correlation. When γ is large (i.e., T is small

relative to N), the eigenvalues become highly dispersed, and many accumulate near zero, leading to numerical instability in matrix inversion. Conversely, when $\gamma \to 0$ (data-rich environments), the eigenvalues concentrate tightly around 1, recovering the population correlation structure.

This result provides the theoretical foundation for modern techniques in correlation estimation and eigenvalue cleaning. In our context, we leverage the MP framework to understand how focusing on the directional component of correlations affects the eigenvalue distribution and, consequently, the stability of portfolio weights.

PROPOSITION 2 (Marčenko–Pastur Law for the Directional Component of Pearson Correlation): *Let* $\mathbf{X} \in \mathbb{R}^{T \times N}$ *be a data matrix with i.i.d. entries of zero mean and variance* 1*, and assume* $\mathbf{X}$ *follows an elliptical distribution. Define the sample Pearson correlation matrix* $\hat{\rho}_{ij}$ *and construct the directional correlation matrix* $\mathbf{D}$ *as:*

$$D_{ij} = \frac{2}{\pi} \arcsin(\hat{\rho}_{ij}).$$

Then, as $N, T \to \infty$ *with* $N/T \to \gamma > 0$*, the empirical spectral distribution of* $\mathbf{D}$ *converges in probability to:*

$$\frac{2}{3} Y_\gamma + \frac{1}{3},$$

where Y_γ *follows the standard Marčenko–Pastur distribution with parameter* γ *as defined above.*

Proof. As $T \to \infty$, we have $\hat{\rho}_{ij} \to \rho_{ij}$. From Proposition 1, under the elliptical distribution assumption, we have the relationship between Kendall's τ and Pearson's ρ:

$$\tau_{ij} = \frac{2}{\pi} \arcsin(\rho_{ij}).$$

Since $\frac{2}{\pi} \arcsin(1) = 1$, the diagonal elements of $\mathbf{D}$ equal one, and the off-diagonal elements correspond to the Kendall's τ correlation matrix. By Bandeira et al. (2017), the empirical

spectral distribution of the Kendall's τ matrix converges to:

$$\frac{2}{3}Y_\gamma + \frac{1}{3},$$

where Y_γ follows the standard Marčenko–Pastur distribution for the Pearson correlation matrix $\mathbf{X}$. □

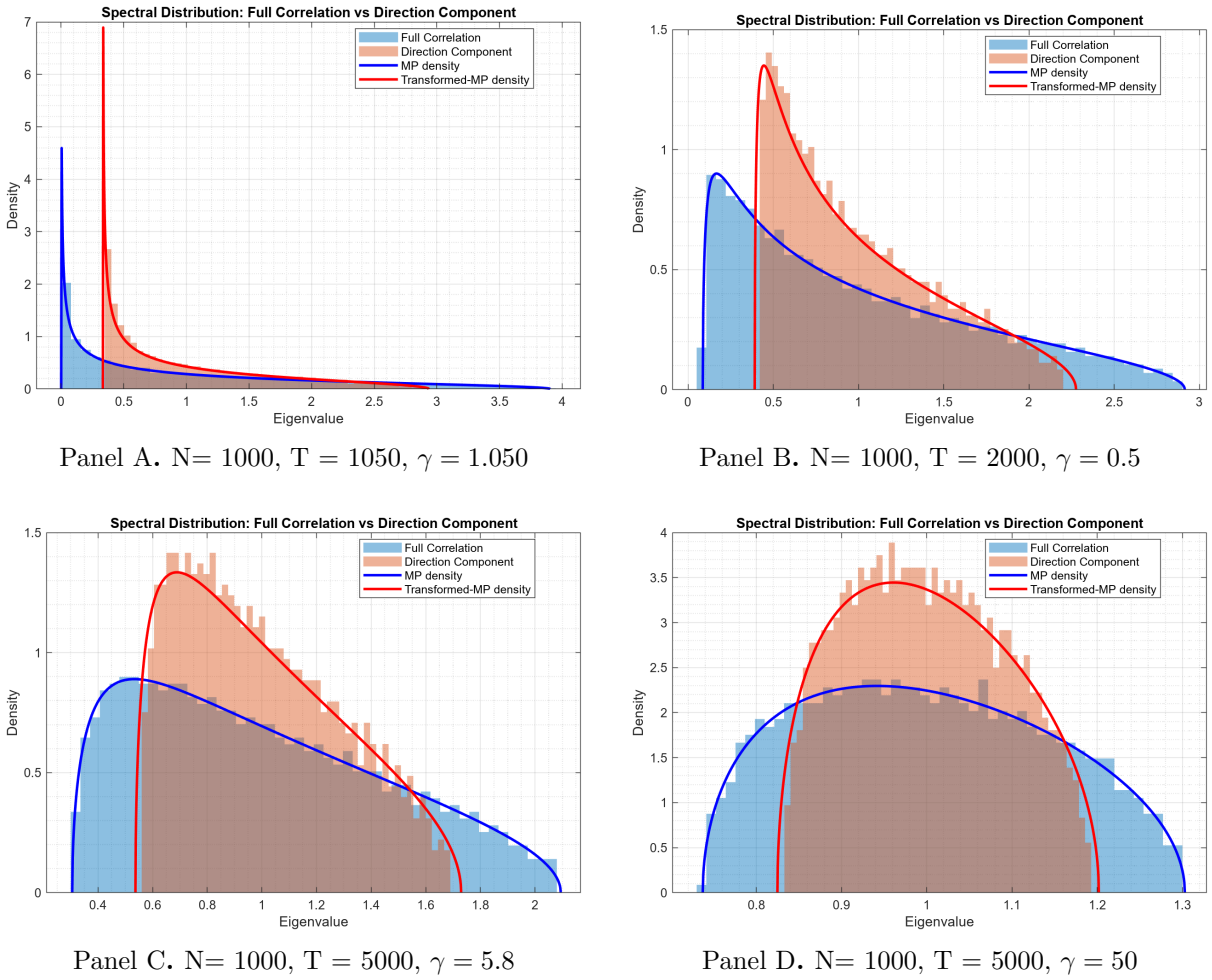

Panel A. N= 1000, T = 1050, $\gamma = 1.050$

Panel B. N= 1000, T = 2000, $\gamma = 0.5$

Panel C. N= 1000, T = 5000, $\gamma = 5.8$

Panel D. N= 1000, T = 5000, $\gamma = 50$

Figure 2. Comparison of empirical eigenvalue distributions and theoretical MP densities under different values of q. The black solid line is the original MP density, and the red dashed line is the affine-transformed MP density.

From Fig. 6, we observe that in data-poor settings (small T relative to N), the eigenvalues of the sample correlation matrix are highly concentrated near zero. Since portfolio weights are proportional to the inverse of the covariance (or correlation) matrix, this

concentration induces extremely unstable and extreme weights.

However, when we focus only on the directional component of the correlation—obtained via the arcsin transformation of Kendall's τ—the eigenvalue distribution becomes significantly shrunk toward 1. This stabilization of the spectrum leads to much more stable portfolio weights, even when data are limited.

Our simulation further illustrates this mechanism by fixing $N = 100$ and increasing T. When T is extremely large, the eigenvalues of the Pearson correlation matrix naturally become well distributed around 1, and the directional component yields a spectrum even closer to 1. This explains the earlier finding: in data-rich environments, Kendall's τ may introduce a small model misspecification error relative to Pearson correlation, but its impact on portfolio performance is negligible because the estimation error disappears when sufficient data are available.

In all other cases, however, focusing only on the directional component offers a practical and robust solution. This section identifies the channel through which extreme portfolio weights arise—namely, the distortion of the eigenvalue distribution in high-dimensional, data-poor setups. From this perspective, shrinking correlations is effectively equivalent to shrinking the eigenvalue distribution, which aligns with other eigenvalue-cleaning techniques in the literature. Therefore, applying the arcsin transformation to isolate the directional component of correlation provides a theoretically justified and numerically stable approach to covariance estimation, avoiding the need for extremely long time-series data.

V. Conclusion

The financial literature presents conflicting views on the role of correlation in portfolio choice. On one hand, behavioral research identifies "correlation neglect" as a cognitive bias, suggesting that investors should be encouraged to incorporate correlation information into their decisions. On the other hand, a large body of empirical evidence demonstrates that portfolios constructed using sample correlation estimates often perform poorly out of sample, implying that ignoring these correlations-as simple heuristics do-can be a rational response to estimation error.

This paper reconciles this tension by showing that the sample correlation coefficient is not a monolithic signal. We decompose correlation into two components: its *magnitude* (the strength of comovement) and its *direction* (the sign of comovement). Our central finding is that the directional component is remarkably stable and robust, even in high-dimensional settings with limited data. In contrast, the magnitude is the primary source of estimation error, which destabilizes portfolio weights and deteriorates out-of-sample performance.

We systematically analyze how correlation neglect influences portfolio diversification. Using simulation studies, we demonstrate that ignoring the magnitude of correlations introduces some specification error but substantially reduces the variance caused by sampling noise, thereby improving robustness. A spectral analysis of the correlation matrix further reveals the mechanism: by effectively shrinking extreme eigenvalues, correlation neglect stabilizes the matrix inversion required for portfolio optimization, which in turn reduces extreme portfolio weights. Our empirical analysis confirms that this mechanism also operates in real financial data, where our proposed method consistently delivers improved diversification and risk-adjusted performance.

The implications of our findings are threefold. First, for financial econometrics, our approach offers a new and intuitive form of elementwise shrinkage on the correlation matrix, providing a theoretical foundation for discarding noisy components while retaining

stable, valuable information. Second, for behavioral finance, our work refines the debate on correlation neglect: investors' skepticism toward sample correlation estimates may not reflect a cognitive bias but rather a rational response to excessive noise. The prescriptive takeaway is not to indiscriminately "use correlation" but instead to focus on its stable directional component. Third, for practitioners, our proposed method is simple to implement and provides a practical pathway to harness the benefits of diversification without being misled by unstable correlation estimates.

In conclusion, the question is not whether to use correlation, but how. By disentangling the stable signal from the volatile noise within the correlation matrix, this paper provides both a theoretically grounded and practically effective framework for building more robust, well-diversified, and stable portfolios.

REFERENCES

Bandeira, Afonso S., Asad Lodhia, and Philippe Rigollet, 2017, Marčenko-Pastur law for Kendall's tau, *Electronic Communications in Probability* 22, 1 – 7.

Bao, Zhigang, 2019, TracyâWidom limit for Kendall's tau, *The Annals of Statistics* 47, 3504 – 3532.

Bickel, Peter J., and Elizaveta Levina, 2008, Covariance regularization by thresholding, *The Annals of Statistics* 36, 2577–2604.

Bouchaud, Jean-Philippe, and Marc Potters, 2009, Financial applications of random matrix theory: a short review, *arXiv preprint arXiv:0910.1205* .

Bun, Joël, Jean-Philippe Bouchaud, and Marc Potters, 2016, My beautiful laundrette: Cleaning correlation matrices for portfolio optimization.

Craig MacKinlay, A, and L'uboš Pástor, 2000, Asset pricing models: Implications for expected returns and portfolio selection, *The Review of Financial Studies* 13, 883–916.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal, 2009, Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?, *The Review of Financial Studies* 22, 1915–1953.

Edirisinghe, Chanaka, and Wenjun Zhou, 2014, Portfolio optimization using rank correlation, in *Encyclopedia of Business Analytics and Optimization*, 1866–1879 (IGI Global Scientific Publishing).

Enke, Benjamin, and Florian Zimmermann, 2019, Correlation neglect in belief formation, *The Review of Economic Studies* 86, 313–332.

Espana, Tomas, Victor Le Coz, and Matteo Smerlak, 2024, Kendall correlation coefficients for portfolio optimization, *arXiv preprint arXiv:2410.17366* .

Eyster, Erik, and Georg Weizsacker, 2010, Correlation neglect in financial decision-making .

Gerber, Sander, Harry Markowitz, Philip Ernst, Yinsen Miao, Paul Sargen, et al., 2022, The gerber statistic: A robust co-movement measure for portfolio optimization, *The Journal of Portfolio Management* 48, 87–102.

Jagannathan, Ravi, and Tongshu Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps, *Journal of Finance* 58, 1651–1683.

Kallir, Ido, and Doron Sonsino, 2009, The neglect of correlation in allocation decisions, *Southern Economic Journal* 75, 1045–1066.

Karoui, Noureddine El, 2008, Operator norm consistent estimation of large-dimensional sparse covariance matrices, *The Annals of Statistics* 36, 2717–2756.

Kirby, Chris, and Barbara Ostdiek, 2012, It's all in the timing: simple active portfolio strategies that outperform naive diversification, *Journal of Financial and Quantitative Analysis* 47, 437–467.

Laudenbach, Christine, Michael Ungeheuer, and Martin Weber, 2023, How to alleviate correlation neglect in investment decisions, *Management Science* 69, 3400–3414.

Ledoit, Oliver, and Michael Wolf, 2008, Robust performance hypothesis testing with the sharpe ratio, *Journal of Empirical Finance* 15, 850–859.

Ledoit, Olivier, and Michael Wolf, 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, *Journal of empirical finance* 10, 603–621.

Ledoit, Olivier, and Michael Wolf, 2004, A well-conditioned estimator for large-dimensional covariance matrices, *Journal of multivariate analysis* 88, 365–411.

Ledoit, Olivier, and Michael Wolf, 2020, The power of (non-)linear shrinking: A review and guide to covariance matrix estimation, *Journal of Financial Econometrics* 20, 187–218.

Lindskog, Filip, Alexander McNeil, and Uwe Schmock, 2003, Kendallâs tau for elliptical distributions, in *Credit risk: Measurement, evaluation and management*, 149–156 (Springer).

Marčenko, Vladimir A, and Leonid Andreevich Pastur, 1967, Distribution of eigenvalues for some sets of random matrices, *Mathematics of the USSR-Sbornik* 1, 457.

Markowitz, Harry, 1952, Portfolio selection, *Journal of Finance* 7, 77–91.

Rothman, Adam J, Elizaveta Levina, and Ji Zhu, 2009, Generalized thresholding of large covariance matrices, *Journal of the American Statistical Association* 104, 177–186.

Ungeheuer, Michael, and Martin Weber, 2021, The perception of dependence, investment decisions, and stock prices, *Journal of Finance* 76, 797–844.

VI. Appendix

Appendix A. More Stuff

Table VII. Herfindahl Index for simulated data

	$N = 50$			$N = 100$			$N = 500$		
Strategy	$T = 505$	$T = 1250$	$T = 2500$	$T = 505$	$T = 1250$	$T = 2500$	$T = 505$	$T = 1250$	$T = 2500$
Noise Level					20%				
K's τ^2	0.1471	0.1462	0.1465	0.0741	0.0735	0.0736	0.0252	0.0236	0.0228
True	1.7570	1.7570	1.7570	1.5329	1.5329	1.5329	1.1679	1.1679	1.1679
LW	0.4491	0.7850	1.0945	0.2068	0.3682	0.6159	0.0902	0.0600	0.0854
SampleC	1.7259	1.7127	1.7187	1.5932	1.5524	1.5517	7.4667	1.2617	1.2432
Noise Level					50%				
K's τ^2	0.1172	0.1013	0.0979	0.2633	0.2415	0.2343	0.8094	0.7581	0.7285
True	1.0913	1.0913	1.0913	1.1000	1.1000	1.1000	1.0496	1.0496	1.0496
LW	0.0542	0.0876	0.1640	0.0429	0.0508	0.0750	0.0303	0.0283	0.0250
SampleC	1.1266	1.1081	1.1027	1.1635	1.1361	1.1383	6.1870	1.1003	1.0978
Noise Level					80%				
K's τ^2	0.3377	0.3392	0.3403	0.1694	0.1742	0.1765	0.0217	0.0226	0.0233
True	1.0234	1.0234	1.0234	1.0306	1.0306	1.0306	1.0249	1.0249	1.0249
LW	0.0238	0.0275	0.0364	0.0144	0.0169	0.0196	0.0086	0.0118	0.0127
SampleC	1.0601	1.0469	1.0419	1.0873	1.0686	1.0719	4.9758	1.0507	1.0600

Table VIII. Average Short Position Size for simulated data

	$N = 50$			$N = 100$			$N = 500$		
Strategy	$T = 505$	$T = 1250$	$T = 2500$	$T = 505$	$T = 1250$	$T = 2500$	$T = 505$	$T = 1250$	$T = 2500$
Noise Level					20%				
K's τ^2	0.5814	0.5761	0.5730	0.7484	0.7355	0.7274	1.1115	1.0488	1.0184
True	1.2696	1.2696	1.2696	1.3761	1.3761	1.3761	1.5144	1.5144	1.5144
LW	0.9578	1.0622	1.1389	1.0510	1.0902	1.1623	2.0853	1.3986	1.2504
SampleC	1.3111	1.2867	1.2788	1.4874	1.4140	1.3928	14.2269	1.8698	1.6586
Noise Level					50%				
K's τ^2	0.1850	0.1832	0.1822	0.1246	0.1226	0.1216	0.0463	0.0449	0.0440
True	0.4906	0.4906	0.4906	0.7184	0.7184	0.7184	1.2369	1.2369	1.2369
LW	0.1375	0.1657	0.2085	0.3337	0.3514	0.3763	1.0713	1.0221	0.9192
SampleC	0.5479	0.5172	0.5112	0.8576	0.7761	0.7551	13.6635	1.5936	1.3839
Noise Level					80%				
K's τ^2	0.0372	0.0238	0.0212	0.1034	0.0789	0.0691	0.5400	0.4800	0.4466
True	0.2277	0.2277	0.2277	0.3844	0.3844	0.3844	0.9486	0.9486	0.9486
LW	0.0004	0.0014	0.0029	0.0161	0.0371	0.0496	0.3414	0.4911	0.5341
SampleC	0.2849	0.2507	0.2456	0.5193	0.4397	0.4210	12.1656	1.2949	1.0917

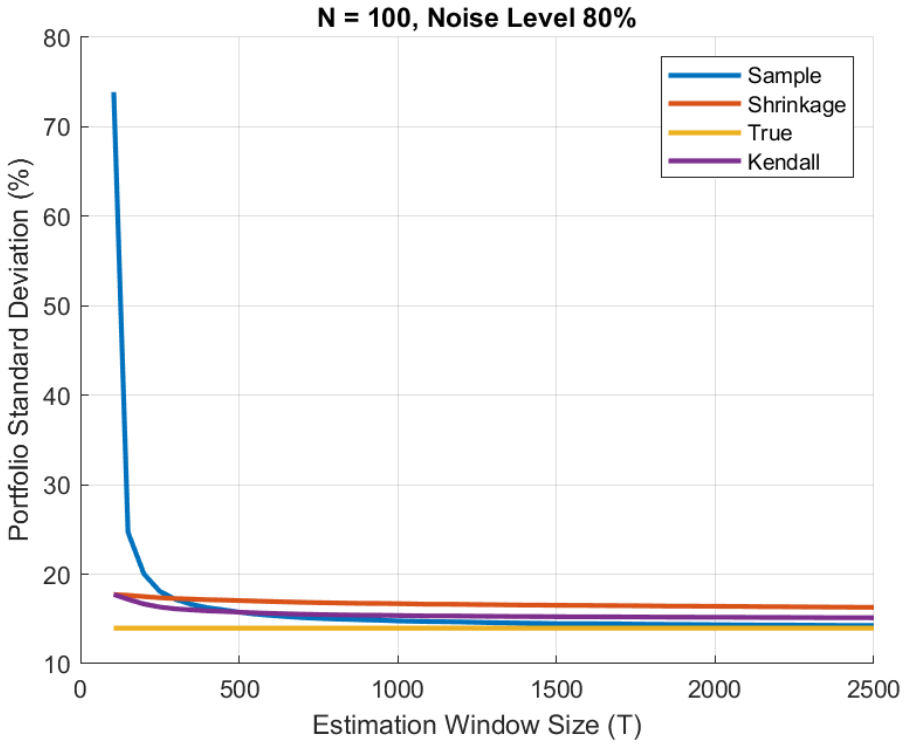

Figure 3. Comparison of portfolio standard deviation for simulated data as the estimation window size increases.

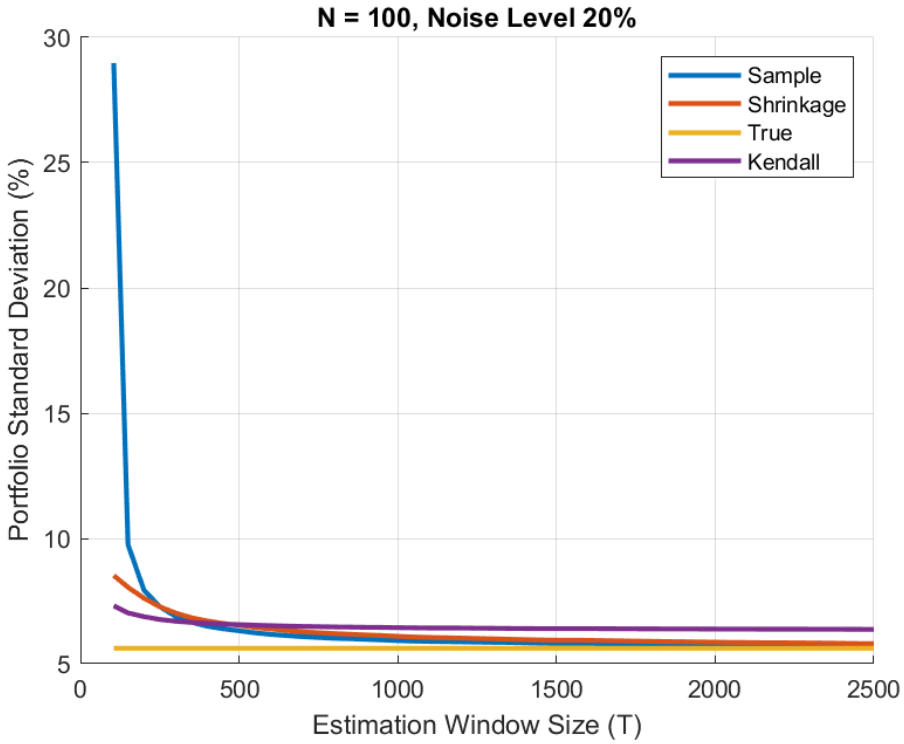

Figure 4. Comparison of portfolio standard deviation for simulated data as the estimation window size increases.

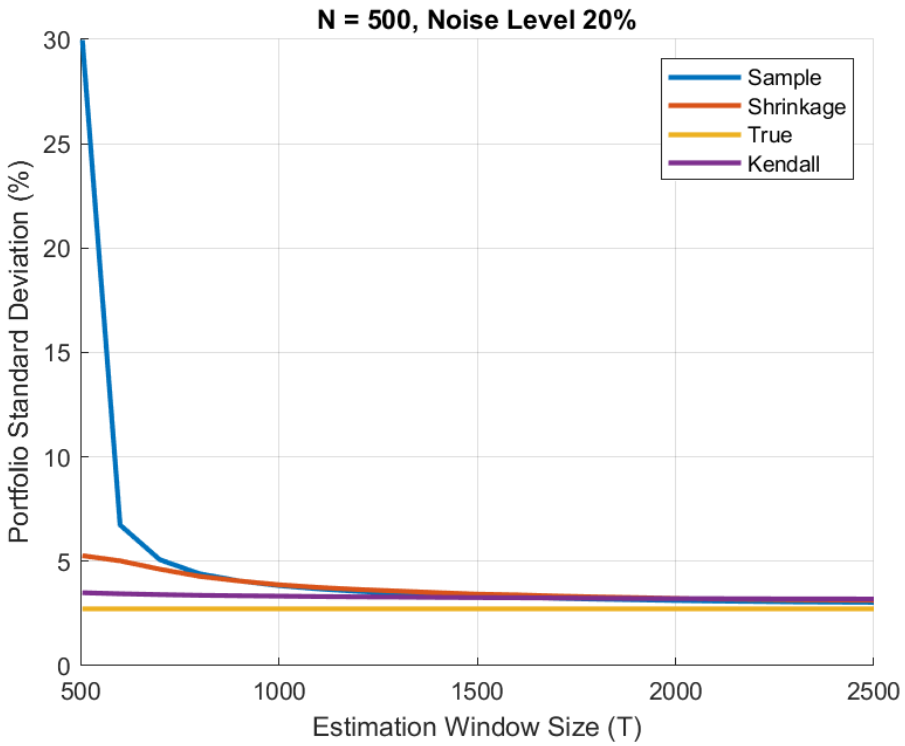

Figure 5. Comparison of portfolio standard deviation for simulated data as the estimation window size increases.

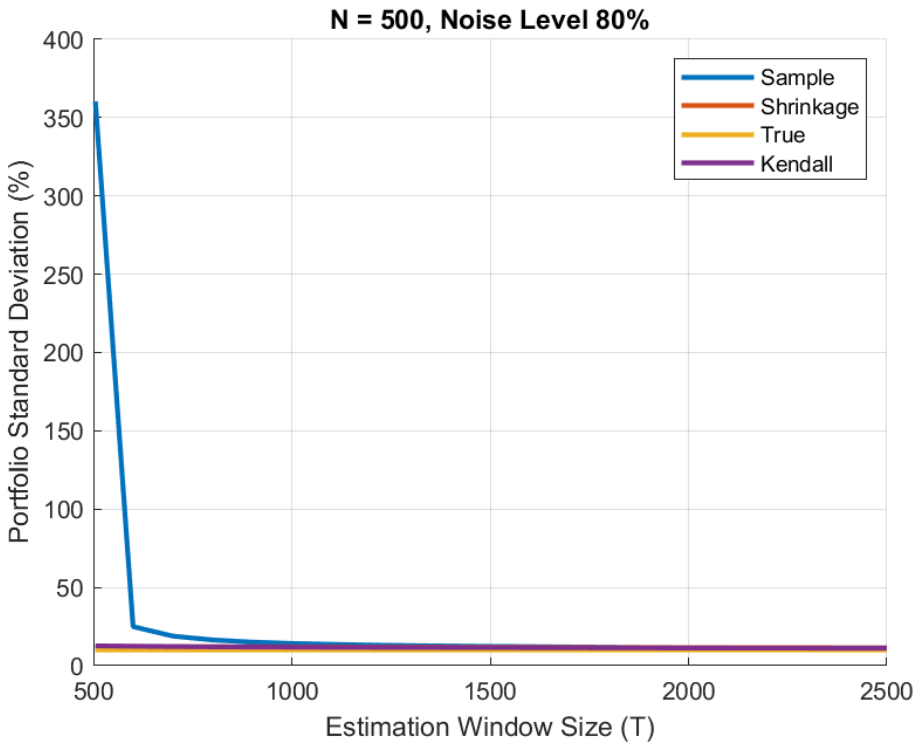

Figure 6. Comparison of portfolio standard deviation for simulated data as the estimation window size increases.