

The Dispersion Bias

Lisa Goldberg^{*} Alex Papanicolaou[†] Alex Shkolnik[‡]

November 4, 2017

This Version: February 14, 2018[§]

Abstract

Estimation error has plagued quantitative finance since Harry Markowitz launched modern portfolio theory in 1952. Using random matrix theory, we characterize a source of bias in the sample eigenvectors of financial covariance matrices. Unchecked, the bias distorts weights of minimum variance portfolios and leads to risk forecasts that are severely biased downward. To address these issues, we develop an eigenvector bias correction. Our approach is distinct from the regularization and eigenvalue shrinkage methods found in the literature. We provide theoretical guarantees on the improvement our correction provides as well as estimation methods for computing the optimal correction from data.

^{*}Departments of Economics and Statistics and Consortium for Data Analytics in Risk, University of California, Berkeley, CA 94720 and Aperio Group, lrg@berkeley.edu.

[†]Consortium for Data Analytics in Risk, University of California, Berkeley, CA 94720, apapanicolaou@berkeley.edu.

[‡]Consortium for Data Analytics in Risk, University of California, Berkeley, CA 94720, ads2@berkeley.edu.

[§]We thank the Center for Risk Management Research, the Consortium for Data Analytics in Risk, and the Coleman Fung Chair for financial support. We thank Marco Avellaneda, Bob Anderson, Kay Giesecke, Nick Gunther, Guy Miller, George Papanicolaou, Yu-Ting Tai, participants at the the 3rd Annual CDAR Symposium in Berkeley, participants at the Swissquote Conference 2017 on FinTech, and participants at the UC Santa Barbara Seminar in Statistics and Applied Probability for discussion and comments. We are grateful to Stephen Bianchi, whose incisive experiment showing that it is errors in eigenvectors, and not in eigenvalues, that corrupt large minimum variance portfolios, pointed us in a good direction.

1 Introduction

Harry Markowitz transformed finance in 1952 by framing portfolio construction as a tradeoff between mean and variance of return. This application of mean-variance optimization is the basis of theoretical breakthroughs as fundamental as the Capital Asset Pricing Model (CAPM) and Arbitrage Pricing Theory (APT), as well as practical innovations as impactful as Exchange Traded Funds.¹ Still, all financial applications of mean-variance optimization suffer from estimation error in covariance matrices, and we highlight two difficulties. First, a portfolio that is optimized using an estimated covariance matrix is never the true Markowitz portfolio. Second, in current practice, the forecasted risk of the optimized portfolio is typically too low, sometimes by a wide margin. Thus, investors end up with the wrong portfolio, one that is riskier, perhaps a lot riskier, than anticipated.

In this article, we address these difficulties by correcting a systematic bias in the first eigenvector of a sample covariance matrix. Our setting is that of a typical factor model,² but our statistical setup differs from most recent literature. In the last two decades, theoretical and empirical emphasis has been on the case when the number of assets N and number of observations T are both large. In this regime, consistency of principal component analysis (PCA) estimates may be established (Bai & Ng 2008). Motivated by many applications, we consider the setting of relatively few observations (in asymptotic theory: N grows and T is fixed). Indeed, an investor often has a portfolio of thousands of securities but only hundreds of observations.³ PCA is applied in this environment in the early, pioneering work by Connor & Korajczyk (1986) and Connor & Korajczyk (1988), but also very recently (Wang & Fan 2017). In this high dimension, low sample-size regime, PCA factor estimates necessarily carry a finite-sample bias. This bias is further amplified by the optimization procedure that is required to compute a Markowitz portfolio.

An elementary simulation experiment reveals that in a large minimum variance portfolio, errors in portfolio weights are driven by the first principal component, not its variance.⁴ The fact that the eigenvalues of the sample covariance matrix are not important requires some nontrivial analysis, which we carry out. In particular, we show (in our asymptotic regime) that the bias in the dominant sample eigenvalue does not effect the performance of the estimated minimum variance portfolio. Only the bias in the dominant sample eigenvec-

¹ The seminal paper is Markowitz (1952). See Treynor (1962) and Sharpe (1964) for the Capital Asset Pricing Model and Ross (1976) for the Arbitrage Pricing Theory.

²More precisely, the eigenvalues of the covariance matrix corresponding to the factors grow linearly in the dimension. This is not the traditional random matrix theory setting in which all eigenvalues are bounded, nor that of “weak” factors, e.g., Onatski (2012).

³While high frequency data are available in some markets, many securities are observed only at a daily horizon or less frequently. Moreover, markets are non-stationary, so even when there is a long history of data available, its relevance to some problems is questionable.

⁴This experiment was first communicated to us by Stephen Bianchi.

tor needs to be addressed. We measure portfolio performance using two well-established metrics. Tracking error, the workhorse of financial practitioners, measures deviations in weights between the estimated (optimized) and optimal portfolios. We use the variance forecast ratio, familiar to both academics and practitioners, to measure the accuracy of the risk forecast of the portfolio, however right or wrong that portfolio may be.

To develop some intuition for the results to come, consider a simplistic world where all security exposures to the dominant (market) factor are identical. With probability one, a PCA estimate of our idealized, dominant factor will have higher dispersion (variation in its entries). Decreasing this dispersion, obviously, mitigates the estimation error. We denote our idealized, dominant factor by z . We prove that the same argument applies to any other dominant factor along the direction of z with high probability for N large. Thus moving our PCA estimate towards z , by some amount, is very likely to decrease estimation error. In the limit ($N \uparrow \infty$), the estimation error is reduced with probability one. The larger the component of the true dominant factor along z is, the more we can decrease the estimation error.

While a careful proof of our result relies on some recent theory on sample eigenvector asymptotics, rule of thumb versions have been known to practitioners since the 1970s (see footnote 10 and the corresponding discussion). Indeed, the dominant risk factor consistently found in the US and many other developed public equity markets has most (if not all) individual equities positively exposed to it. In other words, empirically, the dominant risk factor has a significant component in z . Our characterization of the dispersion bias may then be viewed as a formalization of standard operation procedure.

The remainder of the introduction discusses our contributions and the related literature. Section 2 describes the problem and fundamental results around the sample covariance matrix and PCA. In Section 3, we present our main results on producing a bias corrected covariance estimate. Section 4 discusses the implementation of our correction for obtaining data-driven estimates. Finally, in Section 5 we present numerical results illustrating the performance of our method in improving the estimated portfolio and risk forecasts.

1.1 Our contributions

We contribute to the literature by providing a method that significantly improves the performance of PCA-estimated minimum-variance portfolios. Our approach and perspective appear to be new. We summarize some of the main points.

Several authors (see above) have noted that sample eigenvectors carry a bias in the statistical and model setting we adopt. We contribute in this direction by, first, recognizing that it is the bias in the first sample eigenvector that drives the performance of PCA-based, minimum-variance portfolios. Second, we show that this bias may in fact be corrected to some degree (cf. discussion below (3.7) in Wang & Fan (2017)). In our domain of application this degree

is material. We point out that eigenvalue bias, which has been the focus in most literature, does not have a material impact on minimum-variance portfolio performance. This motivates lines of research into more general Markowitz optimization problems. Finally, our correction can be framed geometrically in terms of the spherical law of cosines. This perspective illuminates possible extensions of our work. We discuss this further in our concluding remarks.

We also develop a bias correction and show that it outperforms standard PCA. Minimum variance portfolios constructed with our corrected covariance matrix are materially closer to optimal, and their risk forecasts are materially more accurate. In an idealized one-factor setting, we provide theoretical guarantees for the size of the improvement. Our theory also identifies some limitations. We demonstrate the efficacy of the method with an entirely data-driven correction. In an empirically calibrated simulation, its performance is far closer to the theoretically optimal than to standard PCA.

1.2 Related literature

The impact of estimation error on optimized portfolios has been investigated thoroughly in simulation and empirical settings. For example, see Jobson & Korkie (1980), Britten-Jones (1999), Bianchi, Goldberg & Rosenberg (2017) and the references therein. DeMiguel, Garlappi & Uppal (2007) compare a variety of methods for mitigating estimation error, benchmarking against the equally weighted portfolio in out-of-sample tests. They conclude that unreasonably long estimation windows are required for current methods to consistently outperform the benchmark. We review some methods most closely related to our approach.⁵

Early work on estimation error and the Markowitz problem was focused on Bayesian approaches. Vasicek (1973) and Frost & Savarino (1986) were perhaps the first to impose informative priors on the model parameters.⁶ More realistic priors incorporating multi-factor modeling are analyzed in Pástor (2000) (sample mean) and Gillen (2014) (sample covariance). Formulae for Bayes' estimates of the return mean and covariance matrix based on normal and inverted Wishart priors may be found in Lai & Xing (2008, Chapter 4, Section 4.4.1).

A related approach to the Bayesian framework is that of shrinkage or regularization of the sample covariance matrix.⁷ Shrinkage methods have been pro-

⁵The literature on this topic is extensive. We briefly mention a few important references that do not overlap at all with our work. Michaud & Michaud (2008) recommends the use of bootstrap resampling. Lai, Xing & Chen (2011) reformulate the Markowitz problem as one of stochastic optimization with unknown moments. Goldfarb & Iyengar (2003) develop a robust optimization procedure for the Markowitz problem by embedding a factor structure in the constraint set.

⁶Preceding work analyzed diffuse priors and was shown to be inefficient (Frost & Savarino 1986). The latter, instead, presumes all stocks are identical and have the same correlations. Vasicek (1973) specified a normal prior on the cross-sectional market betas (dominant factor).

⁷In the Bayesian setup, sample estimates are “shrunk” toward the prior (Lai & Xing 2008).

posed in contexts where little underlying structure is present (Bickel & Levina 2008) as well as those in which a factor or other correlation structure is presumed to exist (e.g. Ledoit & Wolf (2003), Ledoit & Wolf (2004), Fan, Liao & Mincheva (2013) and Bun, Bouchaud & Potters (2016)). Perhaps surprisingly, shrinkage methods turn out to be related to placing constraints on the portfolio weights in the Markowitz optimization. Jagannathan & Ma (2003) show that imposing a positivity constraint typically shrinks the large entries of the sample covariance downward.⁸

As already mentioned, factor analysis and PCA in particular play a prominent role in the literature. It appears that while eigenvector bias is acknowledged, direct⁹ bias corrections are made only to the eigenvalues corresponding to the principal components (e.g. Ledoit & P  ch   (2011) and Wang & Fan (2017)). Some work on characterizing the behavior of sample eigenvectors may be found in Paul (2007) and Shen, Shen, Zhu & Marron (2016). In the setting of Markowitz portfolios, the impact of eigenvalue bias and optimal corrections are investigated in El Karoui et al. (2010) and El Karoui (2013).

Our approach also builds upon several profound contributions in the literature on portfolio composition. In an influential paper, Green & Hollifield (1992) observe the importance of the structure of the dominant factor to the composition of minimum variance portfolios. In particular, the ‘‘dispersion’’ of the dominant factor exposures drives the extreme positions in the portfolio composition. This dispersion is further amplified by estimation error, as pointed out in earlier work by Blume (1975) (see also Vasicek (1973)). These early efforts have led to a number of heuristics¹⁰ to correct the sample bias of dominant factor estimates.

2 Problem formulation

We address the impact of estimation error on Markowitz portfolio optimization. To streamline the exposition, we restrict our focus to the minimum variance portfolio. In particular, given a covariance matrix $\widehat{\Sigma}$ estimated from T observations of returns to N securities, we consider the optimization problem,

$$\begin{aligned} \min_{w \in \mathbb{R}^N} w^\top \widehat{\Sigma} w \\ w^\top \mathbf{1}_N = 1. \end{aligned} \tag{1}$$

We denote by \widehat{w} the solution to (1), the estimated minimum variance portfolio. Throughout, $\mathbf{1}_N$ is the N -vector of all ones. It is well-known that the portfolio

⁸This is generalized and analyzed further in DeMiguel, Garlappi, Nogales & Uppal (2009).

⁹Several approaches to alter the sample eigenvectors indirectly (e.g. shrinking the sample towards some structured covariance) do exist. However, the analysis of these approaches is not focused on characterizing the bias inherent to the sample eigenvectors themselves.

¹⁰For example, the Blume and Vasicek (beta) adjustments. See the discussion of Exhibit 3 and footnote 7 in Clarke, De Silva & Thorley (2011).

weights, \hat{w} , are extremely sensitive to errors in the estimated model, and risk forecasts for the optimized portfolio tend to be too low.¹¹ We aim to address these issues in a high-dimension, low sample-size regime, i.e., $T \ll N$.

We are also interested in the equally weighted portfolio where $w_e = 1_N/N$, a very simple non-optimized portfolio frequently employed as a benchmark. We use this benchmark to test whether the corrections we make for improving the minimum variance portfolio are not offset by degraded performance elsewhere.

2.1 Model specification & assumptions

We consider a one-factor, linear model for returns to $N \in \mathbb{N}$ securities. Here, a generating process for the N -vector of excess returns R takes the form

$$R = \phi\beta + \epsilon \quad (2)$$

where ϕ is the return to the factor, $\beta = (\beta_1, \dots, \beta_N)$ is the vector of factor exposures, and $\epsilon = (\epsilon^1, \dots, \epsilon^N)$ is the vector of diversifiable specific returns. While the returns $(\phi, \epsilon) \in \mathbb{R} \times \mathbb{R}^N$ are random, we treat each exposure $\beta_n \in \mathbb{R}$ as a constant to be estimated. Assuming ϕ and the $\{\epsilon_n\}$ are mean zero and pairwise uncorrelated, the $N \times N$ covariance matrix of R can be expressed as

$$\Sigma = \sigma^2\beta\beta^\top + \Delta. \quad (3)$$

Here, σ^2 is the variance of ϕ and Δ a diagonal matrix with n th entry $\Delta_{nn} = \delta_n^2$, the variance of ϵ_n . Estimation of Σ is central to numerous applications.

We consider a setting in which T observations $\{R_t\}_{t=1}^T$ of the vector R are generated by a latent time-series $\{\phi_t, \epsilon_t\}_{t=1}^T$ of (ϕ, ϵ) . It is standard to assume the observations are i.i.d. and we do so throughout. Finite-sample error distorts measurement of the parameters $(\sigma^2, \beta, \delta^2)$ leading to the estimate,

$$\hat{\Sigma} = \hat{\sigma}^2\hat{\beta}\hat{\beta}^\top + \hat{\Delta}, \quad (4)$$

which approximates (3) by using the estimated model $(\hat{\sigma}^2, \hat{\beta}, \hat{\delta}^2)$.

Without loss of generality we assume the following condition throughout.

Condition 2.1. *Both β and any estimate $\hat{\beta}$ are normalized as $\|\beta\| = \|\hat{\beta}\| = 1$.*

We require further statistical and regularity assumptions on our model for our technical results. These conditions stem from our use of recent work on spiked covariance models (Wang & Fan 2017). Some may be relaxed in various ways. Our numerical results (Section 5) investigate a much more general setup.

¹¹Extreme sensitivity of portfolio weights to estimation error and the downward bias of risk forecasts are also found in the optimized portfolios constructed by asset managers. Portfolio specific corrections of the dispersion bias discussed in this article are useful in addressing these practical problems. The focus on the global minimum variance portfolio in this article highlights the essential logic of our analysis in the simplest possible setting.

Assumption 2.2. The factor variance $\sigma^2 = \sigma_N^2$ satisfies $\frac{N}{T\sigma_N^2} \rightarrow c_1$ as $N \uparrow \infty$ for fixed integer $T \geq 1$ and $c_1 \in (0, \infty)$. Also, $\mathbf{\Delta} = \delta^2 \mathbf{I}$ for fixed $\delta \in (0, \infty)$.

Assumption 2.3. The returns $\{R_t\}_{t=1}^T$ are i.i.d. with $R_1 \sim \mathcal{N}(0, \mathbf{\Sigma})$.

Assumption 2.4. For $z = \mathbf{1}_N / \sqrt{N}$ we have $\sup_N \gamma_{\beta, z} < 1$ and $\sup_N \gamma_{\hat{\beta}, z} < 1$. Also, β and $\hat{\beta}$ are oriented in such a way that $\beta^\top z$ and $\hat{\beta}^\top z$ are nonnegative.

The requirement in Assumption 2.2 that the factor variance grow in dimension while the specific variance stays bounded is pervasive in the factor modeling literature (Bai & Ng 2008). The extra requirement that the specific risk is homogenous (i.e. $\mathbf{\Delta}$ is a scalar matrix) is restrictive but shortens the technical proofs significantly. It is also commonplace to the spiked covariance model literature.¹² We discuss the adjustments required for heterogenous specific risk (i.e. $\mathbf{\Delta}$ diagonal) in Section 4.2. The distributional Assumption 2.3 facilitates several steps in the proofs. In particular, it allows for an elegant characterization of the systematic bias in PCA, i.e., the bias in the first eigenvector of the sample covariance matrix of the returns $\{R_t\}$ (see Section 2.2). Assumption 2.4 is not much of a restriction. First, all results can be easily extended to the case $\beta = z$, which is simply a point of singularity and thus requires a separate treatment. The orientation requirement is essentially without loss of generality. We will see (in Section 3) that the vector z plays a special role in our bias correction procedure. And if $\beta^\top z < 0$, we would simply consider $-z$.

2.2 PCA bias characterization


We consider PCA as the starting point for our analysis as its use for risk factor identification is widespread in the literature. It is appropriate when σ^2 is much larger than $\|\mathbf{\Delta}\|$ (e.g., Assumption 2.2). Assembling a data matrix $\mathbf{R} = (R_1, \dots, R_T)$, we will denote the (data) sample covariance matrix by $\mathbf{S} = T^{-1} \mathbf{R} \mathbf{R}^\top$. PCA identifies $\hat{\sigma}^2$ with the largest eigenvalue of \mathbf{S} and $\hat{\beta}$ with the corresponding eigenvector (the first principal component). The diagonal matrix of specific risks $\mathbf{\Delta}$ is estimated as $\hat{\mathbf{\Delta}} = \mathbf{diag}(\mathbf{S} - \hat{\sigma}^2 \hat{\beta} \hat{\beta}^\top)$, which corresponds to least-squares regression of \mathbf{R} onto the estimated factors. Finally, the estimate $\hat{\mathbf{\Sigma}}$ of the covariance $\mathbf{\Sigma}$ is assembled as in (4).

Bias in the basic PCA estimator above arises from the use of the sample¹³ covariance matrix \mathbf{S} . We focus on the high-dimension, low sample size regime which is most appropriate for the practical applications we consider. Asymptotically, T is fixed and $N \uparrow \infty$, a regime in which the PCA estimates of σ^2

¹²In particular, our need for this condition stem from our use of results from Shen et al. (2016). The assumption may be relaxed by imposing regularity conditions on the entries of $\mathbf{\Delta}$ at the expense of a more cumbersome exposition. It may also be removed entirely if we consider the regime $N, T \uparrow \infty$ as is more common in the literature. We do not pursue this because many important practical applications are restricted to only a small number of observations.

¹³If $\mathbf{\Sigma}$ replaces \mathbf{S} , sample bias vanishes and the estimator is asymptotically exact as $N \uparrow \infty$.

and β are not consistent. We summarize some recent results from the literature below. We also state our characterization of PCA bias as it pertains to its principal components. Our result is novel in that it suggests a remedy.



(a) $\hat{\beta}$ lies near the cone defined by the Ψ in (5) w.h.p. for large N . There is no reference frame to detect the bias since β is unknown.

(b) $\theta_{\hat{\beta},z} > \theta_{\beta,z}$ w.h.p. for large N (their difference is shaded). PCA estimates exhibit a systematic (dispersion) bias relative to vector z .

Figure 1: A unit sphere with example vectors β , $\hat{\beta}$ and $z = 1_N/\sqrt{N}$. The vector z provides a reference frame in which PCA bias may be identified. Note, the angle $\theta_{x,y}$ is also the arc-length between points x and y on the unit sphere.

Let $\theta_{\hat{\beta},\beta}$ denote the angle between β and its PCA estimate $\hat{\beta}$. Shen et al. (2016) showed, under Assumption 2.2 and mild conditions on the moments of $\{R_t\}$, that there exist random variables $\Psi > 1$ and $\xi \geq 0$ such that¹⁴

$$\cos \theta_{\hat{\beta},\beta} \xrightarrow{a.s.} \Psi^{-1} \quad (5)$$

$$\hat{\sigma}^2/\sigma^2 \xrightarrow{a.s.} \Psi^2\xi \quad (6)$$

as $N \uparrow \infty$ where ξ and Ψ depend on T (fixed). The pair (ξ, Ψ) characterizes the error in the PCA estimated model asymptotically. As the sample size T increases, both ξ and Ψ approach one almost surely, i.e., PCA supplies consistent estimates. Since $\Psi > 1$, the estimate $\hat{\sigma}^2$ tends to be biased upward (whenever ξ fluctuates around one). Under Assumption 2.3, $\xi = \chi_T^2/T$ where χ_T^2 has the chi-square distribution with T degrees of freedom. Here, ξ is concentrated around one with high probability (w.h.p.) for even moderate values of T .

Identifying a (systematic) bias in PCA estimates of β requires a more subtle analysis. Observe from (5) that Ψ defines a cone near which the estimate $\hat{\beta}$ will lie with high probability for large N (see panel (a) of Figure 1). However, this

¹⁴More precisely, $\Psi^2 = 1 + \delta^2 c_1/\xi$ where c_1 is identified in Assumption 2.2.

does not point to a systematic error that can be corrected.¹⁵ Indeed, it is not a priori clear where on the cone the vector $\hat{\beta}$ resides as (5) provides information only about the angle $\theta_{\hat{\beta},\beta}$ away from the unknown vector β . We provide a result (see Theorem 2.5 below) that sheds light on this problem.

Recall the vector $z = 1_N/\sqrt{N}$. We consider, not the angle $\theta_{\hat{\beta},\beta}$ between β and $\hat{\beta}$, but the their respective angles $\theta_{\beta,z}$ and $\theta_{\hat{\beta},z}$ to this reference vector.

Theorem 2.5 (PCA bias). *Suppose that Assumptions 2.2, 2.3 and 2.4 hold and let $\hat{\beta}$ be a PCA estimate of β . Then, $\cos \theta_{\beta,z} \stackrel{a.s.}{\sim} \Psi \cos \theta_{\hat{\beta},z}$ as $N \uparrow \infty$ and, in particular, we have that $\theta_{\hat{\beta},z}$ exceeds $\theta_{\beta,z}$ with high probability for large N .*

The proof of this result is deferred to Appendix C. It applies the spiked covariance model results of Shen et al. (2016) and Paul (2007) on sample eigenvectors using a decomposition in terms of the reference vector z .

2.3 Errors in optimized portfolios

Estimation error causes two types of difficulties in optimized portfolios. It distorts portfolio weights, and it biases the risk of optimized portfolios downward. Both effects are present for the minimum variance portfolio \hat{w} , constructed as the solution to (1) using some estimate $\hat{\Sigma}$ of the returns covariance Σ . We now define the metrics for assessing the magnitude of these two errors.

We denote by w_* the optimal portfolio, i.e., the solution of (1) with $\hat{\Sigma}$ replaced by Σ . Since the latter is positive definite, the optimal portfolio weights w_* may be given explicitly.¹⁶ We define,

$$\mathcal{T}_w^2 = (w_* - \hat{w})^\top \Sigma (w_* - \hat{w}), \quad (8)$$

the (squared) tracking error of \hat{w} . Here, \mathcal{T}_w^2 measures the distance between the optimal and estimated portfolios, w_* and \hat{w} . Specifically, it is the square of the width of the distribution of return differences $w_* - \hat{w}$.

The variance of portfolio \hat{w} is given by $\hat{w}^\top \hat{\Sigma} \hat{w}$ and its true variance is $\hat{w}^\top \Sigma \hat{w}$. We define,

$$\mathcal{R}_{\hat{w}} = \frac{\hat{w}^\top \hat{\Sigma} \hat{w}}{\hat{w}^\top \Sigma \hat{w}}, \quad (9)$$

¹⁵Recently, Wang & Fan (2017) provided CLT-type results for the asymptotic distribution of sample eigenvectors (Theorem 3.2). They remark on the near ‘‘impossibility’’ of correcting sample eigenvector bias. This follows from their choice of coordinate system (cf. Figure 1).

¹⁶Indeed, the Sherman-Morrison-Woodbury formula yields an explicit solution even for a multi-factor model and with a guaranteed mean return (El Karoui 2008). In our setting,

$$w_* \propto \Delta^{-1} (1_N \beta_{MV} - \beta), \quad \beta_{MV} = \frac{1 + \sigma^2 \sum_{k=1}^N \beta_k^2 / \delta_k^2}{\sigma^2 \sum_{k=1}^N \beta_k / \delta_k^2}. \quad (7)$$

the variance forecast ratio. Ratio (9) is less than one when the risk of the portfolio \hat{w} is underforecast.¹⁷

Metrics (8) and (9) quantify the errors in portfolio weights and risk forecasts induced by estimation error.¹⁸ We analyze \mathcal{T}_w and \mathcal{R}_w asymptotically. Again recall $z = 1_N/\sqrt{N}$ and let $\gamma_{x,y} = x^\top y$. Note, $\gamma_{x,y} = \cos \theta_{x,y}$ whenever x and y lie on the surface of a unit sphere as in Figure 1. Define,

$$\mathcal{E} = \frac{\gamma_{\beta,z} - \gamma_{\hat{\beta},z}}{\sin^2 \theta_{\hat{\beta},z}}. \quad (10)$$

The variable \mathcal{E} drives the asymptotics of our error metrics. Note that $\mathcal{E} = 0$ when $\hat{\beta} = \beta$. Indeed, it is not difficult to show that for any estimates $(\hat{\sigma}, \hat{\beta}, \hat{\delta})$ satisfying Assumption 2.4, if \mathcal{E} is bounded away from zero (i.e. $\inf_N \mathcal{E} > 0$),

$$\begin{aligned} \mathcal{T}_w^2 &\sim \mu^2 \mathcal{E}^2 \\ \mathcal{R}_w &\sim \frac{N^{-1} \hat{\delta}^2}{\mu^2 \mathcal{E}^2 \sin^2 \theta_{\hat{\beta},z}} \quad (N \uparrow \infty). \end{aligned} \quad (11)$$

where $\mu^2 = \sigma^2/N$ which is bounded in N . This result (with our all assumptions above relaxed) is given by Proposition D.1 of Appendix D.

Corollary 2.6 (PCA performance). *Suppose Assumptions 2.2, 2.3 and 2.4 hold. For the PCA-estimator of the minimum variance portfolio \hat{w} , the tracking error squared \mathcal{T}_w^2 is bounded away from zero and variance forecast ratio \mathcal{R}_w tends to zero as $N \uparrow \infty$. In particular, for the PCA estimator, the error \mathcal{E} in (10) satisfies*

$$\mathcal{E} \stackrel{a.s.}{\sim} \gamma_{\beta,z} \left(\frac{1 - \Psi^{-2}}{1 - \Psi^{-2} \gamma_{\hat{\beta},z}^2} \right) \quad (12)$$

as $N \uparrow \infty$ where $\Psi > 1$ is the random variable in (5).

PROOF. Expression (12) follows from (5), Theorem 2.5 (i.e. $\gamma_{\beta,z} \stackrel{a.s.}{\sim} \Psi \gamma_{\hat{\beta},z}$) and the identity $1 - \gamma_{\hat{\beta},z}^2 = \sin^2 \theta_{\hat{\beta},z}$. The remaining claims follow from (11). ■

The result states that the variance forecast ratio for the minimum variance portfolio \hat{w} that uses the PCA estimator of Section 2.2 is asymptotically 0. The estimated risk will be increasingly optimistic as $N \rightarrow \infty$. This is entirely due to estimation error between the sample eigenvector and the population eigenvector. As N grows, the forecast risk becomes negligible relative to the true risk, rendering the PCA estimated minimum variance portfolio worse and worse. The tracking error is also driven by the error of the sample eigenvector and for increasing dimension, its proximity to the true minimum variance portfolio as measured by tracking error is asymptotically bounded below (away from zero).

¹⁷With respect to the equally weighted portfolio, w_e , (the tracking error of which is zero) we only consider the variance forecast ratio.

¹⁸For a relationship to more standard error norms see Wang & Fan (2017).

3 Bias correction

Our Theorem 2.5 characterizes the bias of the PCA estimator in terms of the vector z . This is the unique (up to negation) dispersionless vector on the unit sphere, i.e., its entries do not vary. Of course when $z = \beta$, then the its PCA estimate $\hat{\beta}$ will have higher dispersion with probability one. The argument works along the projection of any β along z , given by $\gamma_{\beta,z}$. Our PCA bias characterization implies that $\gamma_{\beta,z} > \gamma_{\hat{\beta},z}$, or equivalently, $\theta_{\hat{\beta},z} > \theta_{\beta,z}$ with high probability (for large N). Figure 1b illustrates this systematic PCA bias and clearly suggests a correction: a shrinkage of the PCA estimate $\hat{\beta}$ towards z .

3.1 Intuition for the correction

Given an estimate $\hat{\beta}$, we analyze a parametrized correction of the form

$$\hat{\beta}(\rho) = \frac{\hat{\beta} + \rho z}{\|\hat{\beta} + \rho z\|}. \quad (13)$$

for $\rho \in \mathbb{R}$. The curve $\hat{\beta}(\rho)$ represents a geodesic between $-\hat{\beta}$ and z that passes through $\hat{\beta}$ as ρ passes the origin. We select the optimal “shrinkage” parameter ρ^* as the ρ that minimizes the error in our metrics \mathcal{T}_w^2 and \mathcal{R}_w asymptotically.

With \mathbb{S}^{N-1} denoting the unit N -sphere, we define the space

$$\mathbb{S}_\beta = \{x \in \mathbb{S}^{N-1} : \gamma_{\beta,z} - \gamma_{x,z}\gamma_{x,\beta} = 0\}. \quad (14)$$

Lemma 3.1. *Let $\beta \neq z$. There is a ρ^* such that $\hat{\beta}(\rho^*) \in \mathbb{S}_\beta$ and $\hat{\beta}(\rho^*) \neq z$.*


PROOF. By the spherical law of cosines, we obtain

$$\gamma_{\beta,z} - \gamma_{x,z}\gamma_{x,\beta} = \sin \theta_{x,z} \sin \theta_{x,\beta} \cos \kappa \quad (15)$$

where κ is the angle (in \mathbb{S}^{N-1}) between the geodesic from $\hat{\beta}$ to z and the one from $\hat{\beta}$ to β (see Figure 2a). Write $x = \hat{\beta}(\rho)$ and $\kappa = \kappa_\rho$. Then, $\cos \kappa_\rho = 0$ when $\rho = \rho^*$ for which the geodesic between $\hat{\beta}(\rho^*)$ and β is perpendicular to that between $\hat{\beta}$ and z .¹⁹ By construction, $\hat{\beta}(\rho^*)$ is not z and is in \mathbb{S}_β . ■

We observe that $\hat{\beta}(\rho^*)$ of Lemma 3.1 minimizes the asymptotic error of our metrics. In particular, replacing $\hat{\beta}$ with $\hat{\beta}(\rho^*)$ ensures that \mathcal{E} is zero. Note that $\hat{\beta}(\rho^*) \neq z$ (unless $\beta = z$), i.e., we do not shrink $\hat{\beta}$ all the way to z . While $z \in \mathbb{S}_\beta$, when $\beta \neq z$, the asymptotic error \mathcal{E} explodes as $\hat{\beta}(\rho) \rightarrow z$.

¹⁹If $\hat{\beta}$ and β do not lie on the same side of the sphere we must amend (14) by replacing $\hat{\beta}$ with $-\hat{\beta}$. Note, $\gamma_{\beta,z} - \gamma_{x,z}\gamma_{x,\beta}$ has the same value for x and $-x$.



(a) The spherical law of cosines states: $\gamma_{\beta,z} - \gamma_{\hat{\beta},z} \gamma_{\hat{\beta},\beta} = \sin \theta_{\hat{\beta},z} \sin \theta_{\hat{\beta},\beta} \cos \kappa$ where κ is the angle of the corner opposite $\theta_{\beta,z}$, arc-length between β and z .

(b) Setting the angle $\kappa = \pi/2$ corresponds to setting the error driving term $\mathcal{E} = 0$. Thus, $\hat{\beta}(\rho^*) \in \mathbb{S}_\beta$.

Figure 2: Illustration of the spherical law of cosines and the geodesic $\hat{\beta}(\rho)$ between the points $\hat{\beta}$ and z along which we shrink the PCA estimate $\hat{\beta}$.

In a setting where $\hat{\beta}$ is the PCA estimator, observe that our selection of $\hat{\beta}(\rho^*)$ does more than just correct PCA bias. Theorem 2.5, under the proper assumptions, states that $\gamma_{\beta,z} \stackrel{a.s.}{\sim} \Psi \gamma_{\hat{\beta},z}$ for a random variable $\Psi > 1$. This suggests taking ρ such that $\gamma_{\beta,z}$ equals $\gamma_{\hat{\beta}(\rho),z}$. This choice is not optimal however as it lies on the contour $\{x \in \mathbb{S}^{N-1} : \gamma_{x,z} = \gamma_{\beta,z}\}$ (see Figure 2b). It is not in \mathbb{S}_β unless $\beta = z$, and its asymptotic error \mathcal{E} is bounded away from zero.

3.2 Statement of the main theorem

As noted above, we can improve tracking error and variance forecast ratio by reducing the angle between the estimated eigenvector and the true underlying eigenvector, β , equivalently, by replacing $\hat{\beta}$ with an appropriate choice of $\hat{\beta}_\rho$. We find an optimal value ρ_N^* for a particular N . We present our method and its impact on tracking error and variance forecast ratio for a minimum variance portfolio below. We restrict to the case of homogeneous specific risk where $\Delta = \delta^2 \mathbf{I}$ for expositional purposes but consider the full-fledged case in empirical results in Section 5.

In the following theorem, we also provide a correction to the sample eigenvalue. Our bias correction for the sample eigenvector introduces a bias in the variance forecast ratio for the equally weighted portfolio. We shrink the sample eigenvalue, treated as the variance of the estimated factor, to debias the variance forecast ratio for the equally weighted portfolio.

Theorem 3.2. Suppose Assumptions 2.2, 2.3, 2.4 hold and denote by $\hat{\delta}$ any estimator of the specific risk δ . Define the oracle corrected estimate by $\hat{\beta}_{\rho^*} := \hat{\beta}_{\rho_N^*}$ where the finite sample (fixed N and T) optimal value ρ_N^* solves the equation $0 = r_{\hat{\beta}_\rho} - \gamma_{\beta, \hat{\beta}_\rho}$ (or equivalently, maximizes $\gamma_{\beta, \hat{\beta}_\rho}$).

i. The finite sample optimal oracle ρ_N^* is given by,

$$\rho_N^* = \frac{\gamma_{\beta, z} - \gamma_{\beta, \hat{\beta}} \gamma_{\hat{\beta}, z}}{\gamma_{\beta, \hat{\beta}} - \gamma_{\beta, z} \gamma_{\hat{\beta}, z}}. \quad (16)$$

ii. For the oracle value ρ_N^* , the tracking error and forecast variance ratio for the minimum variance portfolio for $\hat{\Sigma}_{\rho_N^*} = \hat{\sigma}^2 \hat{\beta}_{\rho^*} \hat{\beta}_{\rho^*}^T + \hat{\delta}^2 \mathbf{I}$ satisfy,

$$\mathcal{T}_w^2 \stackrel{a.s.}{\sim} \frac{1}{N} \delta^2 \frac{\gamma_{\hat{\beta}_{\rho^*}, z}^2 - \gamma_{\beta, z}^2}{(1 - \gamma_{\hat{\beta}_{\rho^*}, z}^2)(1 - \gamma_{\beta, z}^2)} \quad (17)$$

$$\mathcal{R}_w \stackrel{a.s.}{\sim} \hat{\delta}^2 / \delta^2. \quad (18)$$

That is, after the optimal correction, the forecast variance ratio for the minimum variance portfolio no longer converges to 0 while the tracking error to the true minimum variance portfolio does.

iii. For T fixed, $N \rightarrow \infty$, we have, $\rho_N^* \stackrel{a.s.}{\sim} \bar{\rho}_N$ where

$$\bar{\rho}_N = \frac{\gamma_{\beta, z}}{1 - \gamma_{\beta, z}^2} (\Psi - \Psi^{-1}). \quad (19)$$

where $\Psi^2 = 1 + \delta^2 c_1 / \xi$, $\xi = \chi_T^2 / T$, and χ_T^2 is the chi-squared distribution with T degrees of freedom. Also, $\bar{\rho}_N > 0$ almost surely if $\gamma_{\beta, z} > 0$. And the asymptotic improvement of the optimal angles $\theta_{\beta, \hat{\beta}_{\rho^*}}$ and $\theta_{\beta, \hat{\beta}_{\bar{\rho}_N}}$ over the original angle $\theta_{\beta, \hat{\beta}}$ as $N \rightarrow \infty$ is,

$$\frac{\sin^2 \theta_{\beta, \hat{\beta}_{\rho_N^*}}}{\sin^2 \theta_{\beta, \hat{\beta}}} \stackrel{a.s.}{\sim} \frac{1 - \gamma_{\beta, z}^2}{1 - (\frac{\gamma_{\beta, z}}{\Psi})^2} = \frac{\sin^2 \theta_{\beta, \hat{\beta}_{\bar{\rho}_N}}}{\sin^2 \theta_{\beta, \hat{\beta}}}.$$

Remark 3.3. Geometrically, there are two views of $\hat{\beta}_{\rho^*}$. One is that $\hat{\beta}_{\rho^*}$ is the projection of $\hat{\beta}$ onto \mathbb{S}_β . The other is that $\hat{\beta}_{\rho^*}$ is the projection of β onto the geodesic defined by $\hat{\beta}_\rho$. In either case, our goal is to find the intersection of the geodesic and the space \mathbb{S}_β .

Remark 3.4. While we consider a specific target, z , in principle the target does not matter. It is possible that these kind of factor corrections can be applied beyond the first factor, given enough information to create a reasonable prior.

The first takeaway from this result is that in the high dimensional limit, it is always possible to improve on the PCA estimate by moving along the geodesic between $\hat{\beta}$ and z . As $\gamma_{\beta,z}$ approaches 0 or for a larger T , the optimal correction approaches 0. Conversely, as $\gamma_{\beta,z}$ approaches 1 or for smaller T , the magnitude of the correction is larger. For $\gamma_{\beta,z} = 1$, the proper choice is naturally to choose z since β and z are aligned in that case.

The improvement in the angle as measured by the ratio of squared sines is bounded in the interval $(1 - \gamma_{\beta,z}^2, 1)$. As $\gamma_{\beta,z}$ approaches 0 or for larger T , the improvement diminishes and the ratio approaches 1. Conversely, for large values of c_1 , the improvement approaches $1 - \gamma_{\beta,z}^2$, indicating that improvement is naturally constrained by how close β is to z in the first place.

In the application to the minimum variance portfolio, the initial idea is to correct the sample eigenvector so that we reduce the angle to the population eigenvector. However, it is not immediately clear that this should have a dramatic effect. Even more surprising is that underestimation of risk has a large component due to sample eigenvector bias and not any sample eigenvalue bias. While an improved estimate $\hat{\beta}_\rho$ has the potential to greatly improve forecast risk, this represents only a single dimension on which to evaluate the performance of a portfolio. We could be sacrificing small tracking error to the true long-short minimum variance portfolio in exchange for better forecasting. That however is not the case here.

Since ξ and Ψ are unobservable non-degenerate random variables, determining their realized values, even with asymptotic estimates, is an impossible task. Hence perfect corrections to kill off the driving term of underestimation of risk are not possible. However, it is possible to make corrections that materially improve risk forecasts.

3.3 Eigenvalue corrections

Our bias correction, based on formula (14), adjusts the dominant eigenvector of the sample covariance matrix \mathbf{S} . It does not involve standard corrections to the eigenvalues of \mathbf{S} , which are well known to be biased. This distinguishes our results from the existing literature (see Section 1.2).

For the purpose of improving accuracy of minimum variance portfolios, there is no need to adjust the dominant eigenvalue. As shown in formulas (11) of Section 2.3, the main drivers of \mathcal{T}^2 and \mathcal{R} for large minimum variance portfolios do not depend on the dominant eigenvalue of \mathbf{S} when returns follow a one-factor model. This is a particular feature of minimum variance, where the dominant factor is effectively hedged.

Since our correction removes a systematic form of bias, it can be used to improve accuracy of other portfolios. If these portfolios have substantial factor exposure, however, a *compatibility adjustment* to the dominant eigenvalue may be required. As our eigenvector adjustment, the compatibility adjustment to the eigenvalue is distinct from the eigenvalue corrections in the literature.

Here, we provide some discussion of compatible eigenvalue corrections for “simple” portfolios, i.e., those that do not depend on the realized matrix of returns \mathbf{R} . Note, that for these weights, the tracking error \mathcal{T} is zero, so we treat the risk forecast ratio \mathcal{R} only.

Under assumptions on our one-factor model that hold for most cases of interest, one can write, for a simple portfolio w that

$$\mathcal{R}_w \sim \frac{\hat{\sigma}_N^2}{\sigma_N^2} C(w, \beta, \hat{\beta})^2 \quad (20)$$

where $C(w, \beta, \hat{\beta}) = (w^\top \beta) / (w^\top \hat{\beta})$. Our correction addresses only the quantity C , but asymptotic formula (20) reveals that for simple portfolios, sample eigenvalues play a material role. Another difference between simple portfolios and minimum variance is that the estimate $\hat{\delta}^2$ of δ^2 does not play any role; the factor risk is all that matters. From the discussion in Section 2.2 we know that $\hat{\sigma}_N^2$ tends to be larger than σ_N^2 for large N , but it is not apriori clear that an eigenvalue correction should aim to lower $\hat{\sigma}_N^2$. This would depend on the behavior of the coefficient $C(w, \beta, \hat{\beta})$ given the estimate $\hat{\beta}$ and the simple portfolio w at hand. Moreover, a correction that decreases $\hat{\sigma}_N^2$ will adjust risk forecast ratios downward, potentially leading to unintended underforecasts. Thus, sample eigenvalue corrections should be coupled with those for the sample eigenvectors to balance their respective terms in (20).

We state a sharp result in this direction for the equally weighted portfolio, a widely used simple portfolio.

Proposition 3.5. *Suppose Assumptions 2.2, 2.3, 2.4 hold. Let $w_e = 1_N/N$ and define the corrected eigenvalue via*

$$\hat{\sigma}_{\rho_N^*}^2 = \left(\frac{\gamma_{\hat{\beta}, z}}{\gamma_{\hat{\beta}_{\rho^*}, z}} \right)^2 \hat{\sigma}_N^2$$

where $\hat{\beta}_{\rho^*}$ is our corrected PCA estimate $\hat{\beta}$ of Theorem 3.2. Then, the forecast variance ratio satisfies $\mathcal{R}_w \stackrel{a.s.}{\sim} \chi_T^2 / T$ as $N \rightarrow \infty$ where χ_T^2 has a chi-squared distribution with T degrees of freedom.

Note that we adjust $\hat{\sigma}_N^2$ downward since by design since $\gamma_{\hat{\beta}, z} \leq \gamma_{\hat{\beta}_{\rho^*}, z}$.

4 Algorithm and extensions

For the precise statement of our algorithm see Appendix A. In what follows we address data-driven corrections for the case $\mathbf{\Delta} = \delta^2 \mathbf{I}$ of Theorem 3.2 as well as the extension to heterogeneous specific risk.

4.1 Data-driven estimator for homogenous specific risk

We introduce procedure for constructing an estimator $\hat{\rho}$ for the asymptotic oracle correction parameter $\bar{\rho}_N$ given in (19). It is based on estimates the specific variance δ^2 and c_1 from Assumption 2.2. From Yata & Aoshima (2012), we have the estimator given by,

$$\hat{\delta}^2 = \frac{\text{Tr}(\mathbf{S}) - \hat{\lambda}_1}{N - 1 - \frac{N}{T}}, \quad (21)$$

where \mathbf{S} is the sample covariance matrix for the data matrix \mathbf{R} and $\hat{\lambda}_1 = \hat{\sigma}^2$ is the first eigenvalue of \mathbf{S} . A natural estimate for the true eigenvalue $\lambda_1(\mathbf{\Sigma})$ is $\hat{\lambda}_1^S = \max\{\hat{\lambda}_1 - \hat{\delta}^2 N/T, 0\}$. For $\hat{\lambda}_1$ sufficiently large, the estimate of c_1 is given by,

$$\hat{c}_1 = \frac{N}{T\hat{\lambda}_1 - N\hat{\delta}^2}. \quad (22)$$

Given the estimates of δ^2 and c_1 , we need a precise value for ξ , as well as Ψ , in order to have a data-driven estimator. We approximate ξ by its expectation $\mathbb{E}[\xi] = 1$ to obtain a completely data driven correction parameter estimate $\hat{\rho}$,

$$\hat{\rho} = \frac{\hat{\Psi}\gamma_{\hat{\beta},z}}{1 - (\hat{\Psi}\gamma_{\hat{\beta},z})^2} (\hat{\Psi} - \hat{\Psi}^{-1}), \quad (23)$$

where $\hat{\Psi}^2 = 1 + \hat{\delta}^2 \hat{c}_1$.

We compute the factor variance as,

$$\hat{\sigma}_{\hat{\rho}}^2 = \Phi_{\hat{\rho}}^2 \hat{\sigma}^2, \quad \Phi_{\hat{\rho}} = \frac{\gamma_{\hat{\beta},z}}{\gamma_{\hat{\beta},z}}$$

where $\hat{\sigma}^2 = \hat{\lambda}_1$ is the first sample eigenvalue from \mathbf{S} .

4.2 An extension to heterogenous specific risk

Our analysis thus far has rested on the simplifying assumption that security return specific variances have a common value. Empirically, this is not the case, and the numerical experiments discussed below in Section 5 allow for the more complex and realistic case of heterogenous specific variances. To address the issue, we modify both the oracle estimator and the data-driven estimator by rescaling betas by specific variance.

Under heterogeneous specific variance, the oracle value ρ_N^* is given by the formula,

$$\rho_N^* = \frac{\gamma_{\hat{\beta},z}^{\hat{\Delta}} - \gamma_{\hat{\beta},\hat{\beta}}^{\hat{\Delta}} \gamma_{\hat{\beta},z}^{\hat{\Delta}}}{\gamma_{\hat{\beta},\hat{\beta}}^{\hat{\Delta}} - \gamma_{\hat{\beta},z}^{\hat{\Delta}} \gamma_{\hat{\beta},z}^{\hat{\Delta}}}, \quad (24)$$

where $\gamma_{x,y}^{\hat{\Delta}} = x^T \hat{\Delta}^{-1} y / \sqrt{\gamma_{x,x}^{\hat{\Delta}} \gamma_{y,y}^{\hat{\Delta}}}$ is a weighted inner product. Furthermore, the risk adjusted returns $\mathbf{R} \hat{\Delta}^{-1/2}$ have covariance $\tilde{\Sigma}$ given by,

$$\tilde{\Sigma} = \sigma^2 \hat{\Delta}^{-1/2} \beta \beta^T \hat{\Delta}^{-1/2} + \mathbf{I}.$$

For the risk adjusted returns, Theorem 3.2 holds. The oracle formula coupled with the risk adjusted returns suggest we use $\tilde{\mathbf{R}} = \mathbf{R} \hat{\Delta}^{-1/2}$ as the data matrix where $\hat{\Delta}$ is the specific risk estimate from the standard PCA method. Why should we expect this to work? The purpose of the scaling $\mathbf{R} \hat{\Delta}^{-1/2}$ is to make the specific return distribution isotropic, and $\tilde{\mathbf{R}}$ approximates that. Since we are only trying to obtain an estimator $\hat{\rho}$ that is close to ρ_N^* , this approximation ends up fine. And for ellipses specified by Δ with relatively low eccentricity, the estimator in (23) actually works in practice since the distribution is relatively close to isotropic. So for larger eccentricity, we require the following adjustment just to get the data closer to an isotropic specific return distribution.

The updated formulas for the heterogenous specific risk correction estimators are given below, and we use them in our numerical experiments. For an initial estimate of specific risk $\hat{\Delta}$, the modified quantities are,

$$\tilde{\beta} = \frac{\hat{\Delta}^{-1/2} \hat{\beta}}{\|\hat{\Delta}^{-1/2} \hat{\beta}\|_2}, \quad \tilde{z} = \frac{\hat{\Delta}^{-1/2} z}{\|\hat{\Delta}^{-1/2} z\|_2}, \quad (25)$$

$$\tilde{\mathbf{S}} = \hat{\Delta}^{-1/2} \mathbf{S} \hat{\Delta}^{-1/2}, \quad \tilde{\lambda}_1 = \hat{\lambda}_1 \|\hat{\Delta}^{-1/2} \hat{\beta}\|^2, \quad (26)$$

$$\tilde{\delta}^2 = \frac{\text{Tr}(\tilde{\mathbf{S}}) - \tilde{\lambda}_1}{N - 1 - \frac{N}{T}}, \quad \hat{c}_1 = \frac{N}{T \tilde{\lambda}_1 - N \tilde{\delta}^2}, \quad (27)$$

$$\hat{\rho} = \frac{\hat{\Psi} \gamma_{\tilde{\beta}, \tilde{z}}}{1 - (\hat{\Psi} \gamma_{\tilde{\beta}, \tilde{z}})^2} (\hat{\Psi} - \hat{\Psi}^{-1}). \quad (28)$$

We use the PCA estimate of the specific risk $\hat{\Delta} = \mathbf{diag}(\mathbf{S} - \hat{\sigma}^2 \hat{\beta} \hat{\beta}^T)$ as the initial estimator.

Once we have the estimated $\hat{\rho}$, we return to the original data matrix \mathbf{R} and apply the correction as before to $\hat{\beta}$, the first eigenvector of the sample covariance matrix. The method for correcting the sample eigenvalue remains the same and we opt to recompute the specific variances using the corrected factor exposures and variance.

5 Numerical study

We use simulation to quantify the dispersion bias and its mitigation in minimum variance and equally portfolios. We design our simulations around a return generating process that is more realistic than the single-factor, homogenous-

specific-risk model featured in Theorems 2.6 and 3.2.²⁰ Returns to $N \in \mathbb{N}$ securities are generated by a multi-factor model with heterogenous specific risk,

$$R = \phi B + \epsilon, \quad (29)$$

where ϕ is the vector of factor returns, B is the matrix of factor exposures, and $\epsilon = (\epsilon^1, \epsilon^2, \dots, \epsilon^N)$ is the vector of diversifiable specific returns.

Our examples are based on a model with $K = 4$ factors. For consistency with Sections 2 and 3, we continue to adopt the mathematically convenient convention of scaling exposure vectors to have L_2 norm 1, and we denote the vector of exposures to the first factor (the first column of the exposure matrix B) by β . The recipe for constructing β with a target value $\gamma_{\beta,z}$ is to generate a random vector, rescale the vector to length 1, and then modify the component in the z direction to have the correct magnitude (while maintaining length 1). A similar approach would be to construct a random vector β' with average equal to 1 and variance equal to τ^2 , where τ^2 is the dispersion of the “market beta,” and then rescale β' to length 1 to obtain β . The parameter τ^2 would control the concentration of the market betas, just like $\gamma_{\beta,z}$, and tends to be greater in calmer regimes. The connection between τ^2 and the dispersion parameter $\gamma_{\beta,z}$ is given by $\gamma_{\beta,z}^2 = 1/\sqrt{1 + \tau^2}$.

The three remaining factors are fashioned from equity styles such as volatility, earnings yield and size.²¹ We draw the exposure of each security to each factor from a mean 0, variance 0.75 normal distribution and again normalize each vector of factor exposures to have L_2 length 1.

We calibrate the risk of the market factor in accordance with Clarke et al. (2011) and Goldberg, Leshem & Geddes (2014); both report the annualized volatility of the US market to be roughly 16%, based on estimates that rely on decades of data. We calibrate the risk of factors 2, 3 and 4 in accordance with Morozov, Wang, Borda & Menchero (2012, Table 4.3), by setting their annualized volatilities to be 8%, 4% and 4%. We assume that the returns to the factors are pairwise uncorrelated.

We draw annualized specific volatilities $\{\delta_n^2\}$ from a uniform distribution on [32%, 64%]. This range is somewhat broader than the estimated range in Clarke et al. (2011).²²

In each experiment, we simulate a year’s worth of daily returns, $T = 250$, to N securities. From this data set, we construct a sample covariance matrix, \mathbf{S} , from which we extract three estimators of the factor covariance matrix $\mathbf{\Sigma}$. The first is the data-driven estimator, the implementation specifics of which are discussed in Section 4.2 and precisely summarized of Appendix A. The second is

²⁰The simplistic setting for our theoretical results showcases the main theoretical tools used in their proofs without the distraction of regularity conditions required for generalization. Global equity risk models used by investors typically include more than 100 factors.

²¹Seminal references on the importance of volatility (or beta), earnings yield and size in explaining cross-sectional correlations is in Black, Jensen & Scholes (1972), Basu (1983) and Banz (1981).

²²For example, the empirically observed range in Clarke et al. (2011) is [25%, 65%].

the *oracle* estimator which is the same as the data-driven estimator but with the true value of $\gamma_{\beta,z}$, the projection of the true factor onto the z -vector, supplied.²³ Our third estimator is classical PCA, which is specified in detail in Section 2.2. We use the three estimated covariance matrices to construct minimum variance portfolios and to forecast their risk. We also use these covariance matrices to forecast the risk of an equally weighted portfolio.

In the experiments described below, we vary the number of securities, N , and the concentration of the market factor, $\gamma_{\beta,z}$. We run 50 simulations for each pair, $(N, \gamma_{\beta,z})$. Results are shown in Figure 3 for a fixed concentration and varying number of securities, and in Figures 4 and 5 for a fixed number of securities and varying concentration. Panels (a) and (b) in each figure show annualized tracking error and volatility for a minimum variance portfolio. Panels (c) and (d) show variance forecast ratio for a minimum variance portfolio and an equally weighted portfolio.

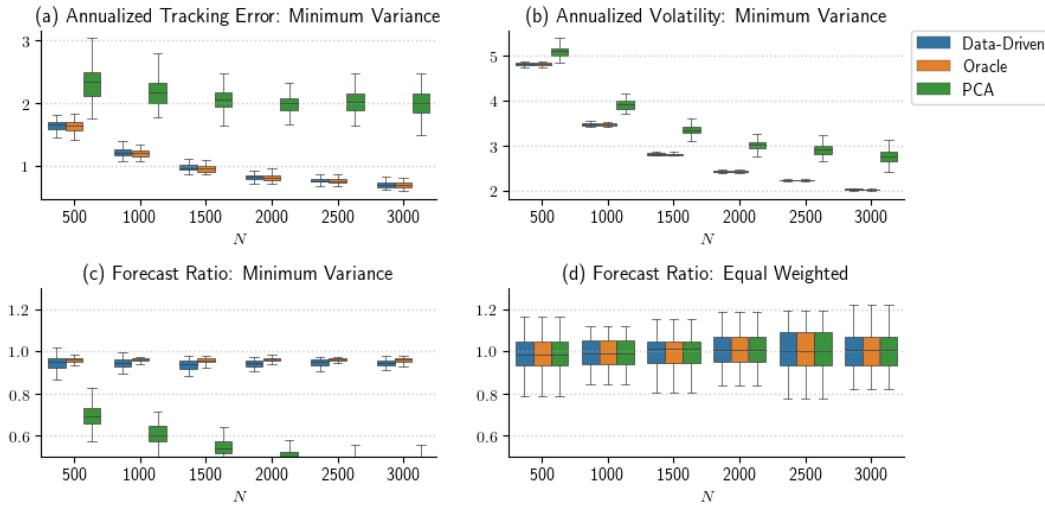


Figure 3: Performance statistics for 50 simulated data sets of $T = 250$ observations as N varies and $\gamma_{\beta,z} = 0.9$. Panel (a): Annualized tracking error for minimum variance portfolios. Panel (b): Annualized volatility for minimum variance portfolios. Panel (c): Variance forecast ratio for minimum variance portfolios. Panel (d): Variance forecast ratio for equally weighted portfolios.

In Figure 3, the concentration of the market factor is $\gamma_{\beta,z} = 0.9$. Panels (a), (b) and (c) show that for minimum variance portfolios optimized with dispersion bias-corrected PCA models, tracking error and volatility decline materially as N grows from 500 to 3000, ranges of outcomes compress, and variance forecast ratios are near 1 for all N considered. These desirable effects are less pronounced, or even absent, in a PCA model without dispersion bias mitigation. A comparison of panels (c) and (d) highlights the difference in accuracy of risk forecasts between a minimum variance portfolio and an equally weighted

²³For the data-driven estimator, this quantity is estimated from the observed data.

portfolio. Dispersion bias mitigation materially improves variance forecast ratio for the former and has no discernible impact for the latter.

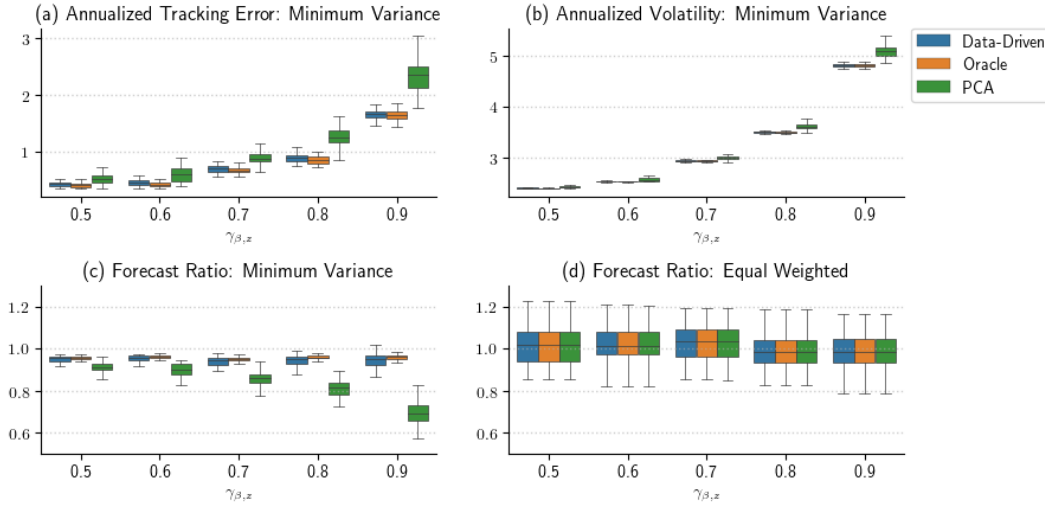


Figure 4: Performance statistics for 50 simulated data sets of $T = 250$ observations as $\gamma_{\beta,z}$ varies and $N = 500$. Panel (a): Annualized tracking error for minimum variance portfolios. Panel (b): Annualized volatility for minimum variance portfolios. Panel (c): Variance forecast ratio for minimum variance portfolios. Panel (d): Variance forecast ratio for equally weighted portfolios.

In Figure 4, the number of securities is $N = 500$. Panels (a), (b) and (c) show that for minimum variance portfolios optimized with PCA models, tracking error and volatility increase materially as $\gamma_{\beta,z}$ grows from 0.5 to 0.9, variance forecast ratio diminishes, and the ranges of outcomes expand. These undesirable effects are diminished when the dispersion bias is mitigated.

Results for values $\gamma_{\beta,z} \in [0.9, 1.0]$ are shown separately in Figure 5. The severe decline in performance for all risk metrics for $\gamma_{\beta,z}$ in this range is a consequence of the sea change in the composition of a minimum variance portfolio that can occur when the dominant eigenvalue of the covariance matrix is sufficiently concentrated. Here, the true minimum variance portfolio tends to lose its short positions.²⁴ Errors in estimation of the dominant factor lead to long-short optimized portfolios approximating long-only optimal portfolios. The market factor is hedged in the former but not in the latter, and this discrepancy propagates to the error metrics.

A comparison of panels (c) and (d) in Figures 4 and 5 highlights, again, the difference in accuracy of risk forecasts between a minimum variance portfolio and an equally weighted portfolio. Dispersion bias mitigation materially improves variance forecast ratio for the former and has no discernible impact on the latter.

²⁴If specific variances are all equal, the minimum variance portfolio becomes equally weighted when the dominant eigenvector is dispersionless.

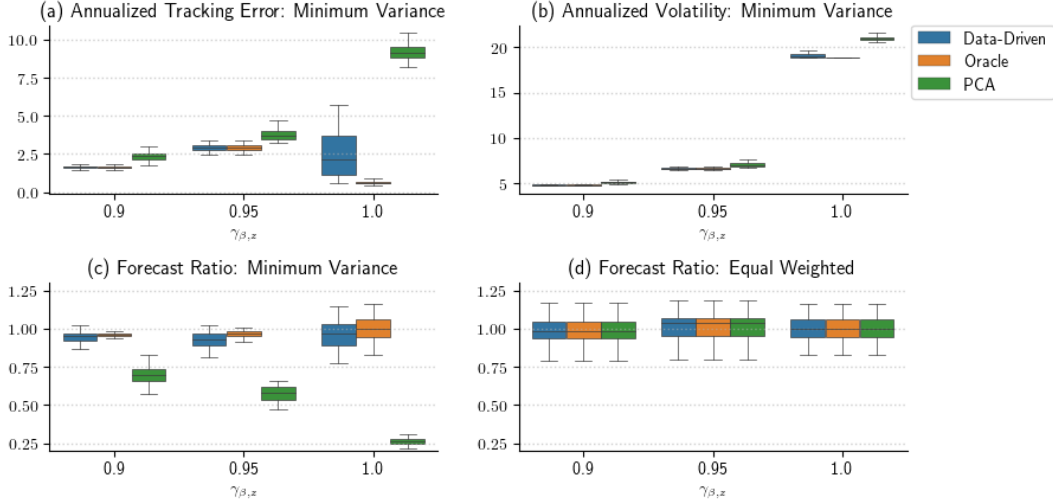


Figure 5: Performance statistics for 50 simulated data sets of $T = 250$ observations as $\gamma_{\beta,z}$ varies and $N = 500$. Panel (a): Annualized tracking error for minimum variance portfolios. Panel (b): Annualized volatility for minimum variance portfolios. Panel (c): Variance forecast ratio for minimum variance portfolios. Panel (d): Variance forecast ratio for equally weighted portfolios.

A casual inspection of the figures suggests that the data-driven estimator performs nearly as well as the oracle. However, there can be substantial differences between the two estimators in some cases, as shown in tracking error and variance forecast ratio of minimum variance Figure 5 when $\gamma_{\beta,z} \in [0.9, 1.0]$. The origin of the differences can be seen Lemma 3.1. In order for the tracking error and variance forecast ratio to have good asymptotic properties, the estimator $\hat{\beta}$ must lie in the null space \mathcal{S}_β . This condition is guaranteed for the oracle estimators but will not generally be satisfied for data-driven estimators.

6 Summary

In this article, we develop a correction for bias in PCA-based covariance matrix estimators. The bias is excess dispersion in a dominant eigenvector, and the form of the correction is suggested by formulas for estimation error metrics applied to minimum variance portfolios.

We identify an oracle correction that optimally shrinks the sample eigenvector along a spherical geodesic toward the distinguished zero-dispersion vector, and we provide asymptotic guarantees that oracle shrinkage reduces both types of error. These findings are especially relevant to equity return covariance matrices, which feature a dominant factor whose overwhelmingly positive exposures tend to be overly dispersed by PCA estimators.

Our results fit into two streams of academic literature. The first is the large- N -fixed- T branch of random matrix theory, which belongs to statistics. The

second is empirical finance, which features results about Bayesian adjustments to estimated betas.

To enable practitioners to use our results, we develop a data-driven estimator of the oracle. Simulation experiments support the practical value of our correction, but much work remains to be done. That includes the development of estimates of the size and likelihood of the exposure bias in finite samples, the identification and correction of biases in other risk factors, and empirical studies. Explicit formulas for error metrics in combination with the geometric perspective in this article provide a way to potentially improve construction and risk forecasts of investable optimized portfolios.

A Algorithm

Our corrected covariance matrix algorithm is given in Algorithm 1 where the input is a data matrix of returns R .

Algorithm 1 1-Factor Bias Corrected PCA Covariance Estimator

Require: $\mathbf{R} = (R_1, \dots, R_T)$

Require: $z = 1_N / \sqrt{N}$

- 1: **procedure** BIAS CORRECTED PCA COVARIANCE(\mathbf{R})
 - 2: $\mathbf{S} \leftarrow \frac{1}{T} \mathbf{R} \mathbf{R}^T$
 - 3: $\hat{\mathbf{U}} \leftarrow [\hat{u}_1, \dots, \hat{u}_N]$, $\hat{\Lambda} \leftarrow \text{Diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_N)$ \triangleright Eigendecomposition of \mathbf{S}
 - 4: $\hat{\beta} \leftarrow \text{sign}(\hat{u}_1^T z) \hat{u}_1$ \triangleright Orient such that $\hat{\beta}^T z > 0$.
 - 5: $\hat{\sigma}^2 \leftarrow \hat{\lambda}_1$
 - 6: $\hat{\Delta} \leftarrow \text{Diag}(\mathbf{S} - \hat{\mathbf{L}})$, $\hat{\mathbf{L}} \leftarrow \hat{\sigma}^2 \hat{\beta} \hat{\beta}^T$ \triangleright Initial PCA estimate.
 - 7: $\tilde{\mathbf{S}} \leftarrow \hat{\Delta}^{-1/2} \mathbf{S} \hat{\Delta}^{-1/2}$, $\tilde{\lambda}_1 \leftarrow \hat{\lambda}_1 \|\hat{\Delta}^{-1/2} \hat{\beta}\|^2$
 $\tilde{\beta} \leftarrow \frac{\hat{\Delta}^{-1/2} \hat{\beta}}{\|\hat{\Delta}^{-1/2} \hat{\beta}\|_2}$, $\tilde{z} \leftarrow \frac{\hat{\Delta}^{-1/2} z}{\|\hat{\Delta}^{-1/2} z\|_2}$
 - 8: $\hat{\Psi} \leftarrow 1 + \tilde{\delta}^2 \hat{c}_1$, $\tilde{\delta}^2 \leftarrow \frac{\text{Tr}(\tilde{\mathbf{S}}) - \tilde{\lambda}_1}{N - 1 - \frac{N}{T}}$, $\hat{c}_1 \leftarrow \frac{N}{T \tilde{\lambda}_1 - N \tilde{\delta}^2}$
 - 9: $\hat{\beta}_{\hat{\rho}} \leftarrow \frac{\hat{\beta} + \hat{\rho} z}{\sqrt{1 + 2 \hat{\rho} \gamma_{\hat{\beta}, z} + \hat{\rho}^2}}$, $\hat{\rho} \leftarrow \frac{\hat{\Psi} \gamma_{\tilde{\beta}, \tilde{z}}}{1 - (\hat{\Psi} \gamma_{\tilde{\beta}, \tilde{z}})^2} (\hat{\Psi} - \hat{\Psi}^{-1})$
 - 10: $\hat{\Delta}_{\hat{\rho}} \leftarrow \text{Diag}(\mathbf{S} - \hat{\mathbf{L}}_{\hat{\rho}})$, $\hat{\mathbf{L}}_{\hat{\rho}} \leftarrow \hat{\sigma}_{\hat{\rho}}^2 \hat{\beta}_{\hat{\rho}} \hat{\beta}_{\hat{\rho}}^T$, $\hat{\sigma}_{\hat{\rho}}^2 \leftarrow \frac{\gamma_{\tilde{\beta}, \tilde{z}}^2}{\gamma_{\hat{\beta}_{\hat{\rho}}, z}^2} \hat{\sigma}^2$
 - 11: **return** $\hat{\Sigma} \leftarrow \hat{\mathbf{L}}_{\hat{\rho}} + \hat{\Delta}_{\hat{\rho}}$ \triangleright The 1-factor bias corrected PCA estimator.
 - 12: **end procedure**
-

B Tables

N	500	1000	1500	2000	2500	3000
PCA	2.35%	2.18%	2.06%	2.01%	2.03%	2.00%
Oracle	1.64%	1.21%	0.97%	0.83%	0.77%	0.70%
Data-Driven	1.66%	1.21%	0.98%	0.84%	0.77%	0.71%

(a) Annualized Tracking Error: Minimum Variance

N	500	1000	1500	2000	2500	3000
PCA	5.10%	3.92%	3.34%	3.03%	2.91%	2.75%
Oracle	4.82%	3.47%	2.80%	2.42%	2.22%	2.01%
Data-Driven	4.82%	3.47%	2.81%	2.42%	2.23%	2.02%

(b) Annualized Volatility: Minimum Variance

N	500	1000	1500	2000	2500	3000
PCA	0.695	0.604	0.541	0.492	0.452	0.413
Oracle	0.963	0.962	0.960	0.961	0.963	0.961
Data-Driven	0.951	0.946	0.943	0.945	0.953	0.947

(c) Forecast Ratio: Minimum Variance

N	500	1000	1500	2000	2500	3000
PCA	0.983	0.991	1.014	1.011	1.001	1.006
Oracle	0.983	0.991	1.015	1.011	1.001	1.006
Data-Driven	0.983	0.991	1.015	1.011	1.001	1.006

(d) Forecast Ratio: Equal Weighted

Table 1: Median values derived from Figure 3. Performance statistics for 50 simulated data sets of $T = 250$ observations as $\gamma_{\beta,z}$ varies and $N = 500$. Panel (a): Annualized tracking error for minimum variance portfolios. Panel (b): Annualized volatility for minimum variance portfolios. Panel (c): Variance forecast ratio for minimum variance portfolios. Panel (d): Variance forecast ratio for equally weighted portfolios.

$\gamma_{\beta,z}$	0.50	0.60	0.70	0.80	0.90	0.95	1.00
PCA	0.51%	0.59%	0.88%	1.26%	2.35%	3.72%	9.11%
Oracle	0.41%	0.42%	0.67%	0.86%	1.64%	2.91%	0.63%
Data-Driven	0.43%	0.45%	0.70%	0.89%	1.66%	2.94%	2.18%

(a) Annualized Tracking Error: Minimum Variance

$\gamma_{\beta,z}$	0.50	0.60	0.70	0.80	0.90	0.95	1.00
PCA	2.42%	2.57%	2.99%	3.62%	5.10%	7.02%	20.95%
Oracle	2.40%	2.54%	2.94%	3.50%	4.82%	6.63%	18.88%
Data-Driven	2.41%	2.54%	2.95%	3.51%	4.82%	6.65%	18.99%

(b) Annualized Volatility: Minimum Variance

$\gamma_{\beta,z}$	0.50	0.60	0.70	0.80	0.90	0.95	1.00
PCA	0.915	0.904	0.860	0.816	0.695	0.579	0.264
Oracle	0.958	0.965	0.949	0.959	0.963	0.967	1.001
Data-Driven	0.955	0.959	0.945	0.949	0.951	0.928	0.967

(c) Forecast Ratio: Minimum Variance

$\gamma_{\beta,z}$	0.50	0.60	0.70	0.80	0.90	0.95	1.00
PCA	1.019	1.015	1.035	0.984	0.983	1.04	1.002
Oracle	1.020	1.015	1.035	0.984	0.983	1.04	1.002
Data-Driven	1.020	1.015	1.035	0.984	0.983	1.04	1.002

(d) Forecast Ratio: Equal Weighted

Table 2: Median values derived from Figures 4 and 5. Performance statistics for 50 simulated data sets of $T = 250$ observations as $\gamma_{\beta,z}$ varies and $N = 500$. Panel (a): Annualized tracking error for minimum variance portfolios. Panel (b): Annualized volatility for minimum variance portfolios. Panel (c): Variance forecast ratio for minimum variance portfolios. Panel (d): Variance forecast ratio for equally weighted portfolios.

C Proof of main results

We start off with some foundational asymptotic results from the literature. Let $\mathbf{X} \sim \mathcal{N}(0, \mathbf{\Lambda})$ where $\mathbf{\Lambda} = \mathbf{diag}(\lambda_1, \lambda_2, \dots, \lambda_2)$ is a diagonal matrix satisfying with λ_1 satisfying Assumption C1 in Shen et al. (2016). Let $\mathbf{S}_\mathbf{X} = \frac{1}{T}\mathbf{X}\mathbf{X}^T$ be the sample covariance for \mathbf{X} with eigendecomposition $\mathbf{S}_\mathbf{X} = \hat{\mathbf{V}}\hat{\mathbf{\Lambda}}\hat{\mathbf{V}}^T$. Further let \hat{v}_1 be the first sample eigenvector given by,

$$\hat{v}_1 = \begin{bmatrix} \hat{v}_{11} \\ \dots \\ \hat{v}_{N1} \end{bmatrix},$$

and define $\tilde{v} = [\hat{v}_{21} \dots \hat{v}_{N1}]^T$. By Paul (2007, Theorem 6), $\tilde{v} \sim \text{Unif}(B(N-2))$ where $B(N-2)$ is a unit $N-2$ sphere.

Via a simple scaling by λ_2 , by Shen et al. (2016, Theorem 6.3) we have $e_1^T \hat{v}_1 = \hat{v}_{11} \xrightarrow{a.s.} \Psi^{-1}$ where $\Psi^2 = 1 + \lambda_2 c_1 / \xi$ and $\xi = \chi_T^2 / T$.

We introduce the following lemma, which we will use in the proofs of our results. We leave its proofs until the end.

Lemma C.1. *If $X_n \sim \text{Unif}(B(n-1))$ and $Y_n \in B(n-1)$ is a sequence independent of X_n , then $X_n^T Y_n \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$.*

PROOF OF THEOREM 2.5. Let $\mathbf{X} = \mathbf{U}\mathbf{R}$ where \mathbf{U} is the matrix of eigenvectors (β, u_2, \dots, u_N) of $\mathbf{\Sigma}$ so that $\text{Cov}(\mathbf{X}) = \mathbf{\Lambda}$ as introduced in the beginning of this section. Also as before let $\mathbf{S}_\mathbf{X} = \frac{1}{T}\mathbf{X}\mathbf{X}^T = \hat{\mathbf{V}}\hat{\mathbf{\Lambda}}\hat{\mathbf{V}}^T$ be the sample covariance of \mathbf{X} and its eigendecomposition. Then $\mathbf{S} = \frac{1}{T}\mathbf{R}\mathbf{R}^T = \mathbf{U}\hat{\mathbf{V}}\hat{\mathbf{\Lambda}}\hat{\mathbf{V}}^T\mathbf{U}^T$ so that the first sample eigenvector of \mathbf{S} , $\hat{\beta}$, is given by,

$$\hat{\beta} = \hat{v}_{11}\beta + \sum_{j=2}^N \hat{v}_{j1}u_j.$$

Then we have $\gamma_{\hat{\beta}, \beta} = \hat{v}_{11}$ and,

$$\gamma_{\hat{\beta}, z} = \hat{v}_{11}\gamma_{\beta, z} + \tilde{v}^T \omega_N,$$

where $\omega_N = [u_2^T z \dots u_N^T z]$. As noted before, by Shen et al. (2016, Theorem 6.3), $\hat{v}_{11} \xrightarrow{a.s.} \Psi^{-1}$. We know that both $\|\tilde{v}\|_2$ and $\|\omega_N\|_2$ are bounded as,

$$\|\tilde{v}\|_2 \leq 1, \quad 1 = \|U^T z\|_2^2 = (\beta^T z)^2 + \sum (u_j^T z)^2 = \gamma_{\beta, z}^2 + \|\omega_N\|_2^2.$$

Therefore, from Lemma C.1 for $X_N = \frac{\tilde{v}}{\|\tilde{v}\|_2}$ and $Y_N = \frac{\omega_N}{\|\omega_N\|_2}$ we have $\tilde{v}^T \omega_N$ converges almost surely to 0 so we conclude the result. ■

PROOF OF THEOREM 3.2.

- i. It is easy to verify that ρ_N^* solves $\gamma_{\beta,z} - \gamma_{\hat{\beta}(\rho),z} \gamma_{\hat{\beta}(\rho),\beta}$ and maximizes $\gamma_{\beta,\hat{\beta}_\rho}$.
- ii. For the oracle value $\rho_N^*, \hat{\beta}_{\rho_N^*} \in \mathcal{S}_\beta$ so the result follows by Lemma 3.1.
- iii. Convergence of ρ_N^* to $\bar{\rho}_N$ stems from Theorem 2.5. It is also clear that $\bar{\rho}_N$ if $\gamma_{\beta,z} > 0$ since $\Psi - \Psi^{-1} > 0$.

For the asymptotic improvement due to shrinkage, we rely on Björck & Golub (1973, Theorem 1), which shows that principal angle can be derived from the singular value decomposition of,

$$\beta^T \begin{bmatrix} z & \frac{\hat{\beta} - \gamma_{\hat{\beta},z} z}{\sqrt{1 - \gamma_{\hat{\beta},z}^2}} \end{bmatrix} = \begin{bmatrix} \gamma_{\beta,z} & \frac{\gamma_{\beta,\hat{\beta}} - \gamma_{\hat{\beta},z} \gamma_{\beta,z}}{\sqrt{1 - \gamma_{\hat{\beta},z}^2}} \end{bmatrix}.$$

By maximizing $\gamma_{\beta,\hat{\beta}_\rho}$ (or equivalently minimizing the angle between β and $\hat{\beta}_\rho$), we are directly choosing ρ_N^* such that $\hat{\beta}_{\rho_N^*}$ is the principal vector with corresponding principal angle to $\hat{\beta}$. Finding the vector in terms of correction quantity is easier through direct maximization of $\gamma_{\beta,\hat{\beta}_\rho}$, and finding the improvement is easier through the principal angle computation, despite the results being equivalent.

From the above product, the squared cosine of the principal angle and its asymptotic value is,

$$\begin{aligned} \gamma_{\beta,\hat{\beta}_{\rho_N^*}}^2 &= \left\| \begin{bmatrix} \gamma_{\beta,z} & \frac{\gamma_{\beta,\hat{\beta}} - \gamma_{\hat{\beta},z} \gamma_{\beta,z}}{\sqrt{1 - \gamma_{\hat{\beta},z}^2}} \end{bmatrix} \right\|_2^2 \\ &= \gamma_{\beta,z}^2 + \frac{(\gamma_{\beta,\hat{\beta}} - \gamma_{\hat{\beta},z} \gamma_{\beta,z})^2}{1 - \gamma_{\hat{\beta},z}^2} \\ &\stackrel{a.s.}{\sim} \gamma_{\beta,z}^2 + \frac{\Psi^{-2}(1 - \gamma_{\beta,z}^2)^2}{1 - (\frac{\gamma_{\beta,z}}{\Psi})^2} \\ &= \gamma_{\beta,\hat{\beta}_{\bar{\rho}_N}}^2 \end{aligned}$$

Since $\sin^2 \theta_{\beta,\hat{\beta}} = 1 - \gamma_{\beta,\hat{\beta}}^2 \xrightarrow{a.s.} 1 - \Psi^{-2}$, we see clearly that,

$$\frac{\sin^2 \theta_{\beta,\hat{\beta}_{\rho_N^*}}}{\sin^2 \theta_{\beta,\hat{\beta}}} \stackrel{a.s.}{\sim} \frac{1 - \gamma_{\beta,z}^2}{1 - (\frac{\gamma_{\beta,z}}{\Psi})^2} = \frac{\sin^2 \theta_{\beta,\hat{\beta}_{\bar{\rho}_N}}}{\sin^2 \theta_{\beta,\hat{\beta}}}$$

■

PROOF OF LEMMA C.1. By orthogonal invariance of X_n and the independence of Y_n ,

$$X_n^T Y_n = X_n^T Q_n^T Q_n Y_n = (Q_n X_n)_1 \stackrel{\mathcal{D}}{=} X_{n1}$$

where Q_n is an orthogonal matrix such that $Q_n Y_n = e_1$, e_1 is the first canonical vector, and X_{n1} is the first entry of X_n . We know from Muller (1959) for X_{n1} ,

$$X_{n1} \stackrel{\mathcal{D}}{=} \frac{Z_1}{\sqrt{Z_1^2 + \chi_{n-1}^2}},$$

where $Z_1 \sim \mathcal{N}(0, 1)$ and is independent of χ_{n-1}^2 . We have,

$$\mathbb{P}(|X_{n1}| \geq \epsilon) \leq \frac{\mathbb{E} X_{n1}^4}{\epsilon^4} \leq \frac{1}{\epsilon^4 n^2} \mathbb{E}[Z_1^4] \mathbb{E}\left[\left(\frac{n}{\chi_n^2}\right)^2\right] \leq \frac{C}{n^2},$$

where the inverse chi-squared distribution has finite moment for n large enough and C is some constant related to the moments of the standard normal distribution and the inverse chi-squared distribution. By the Borel-Cantelli lemma, we conclude the result. ■

PROOF OF PROPOSITION 3.5. For the equally weighted portfolio $w = \frac{1}{N} 1_N$ using $\hat{\beta}_{\rho_N^*}$,

$$\begin{aligned} \mathcal{R}_w &= \frac{\hat{\sigma}_{\rho_N^*}^2 N \gamma_{\hat{\beta}_{\rho_N^*}, z}^2 + \delta^2 \frac{1}{N}}{\sigma^2 N \gamma_{\beta, z}^2 + \delta^2 \frac{1}{N}} \\ &\stackrel{a.s.}{\sim} \frac{\hat{\sigma}_{\rho_N^*}^2 \gamma_{\hat{\beta}_{\rho_N^*}, z}^2}{\sigma^2 \gamma_{\beta, z}^2}, \end{aligned}$$

where $\hat{\sigma}_{\rho_N^*}^2 = \left(\frac{\gamma_{\hat{\beta}, z}}{\gamma_{\hat{\beta}_{\rho_N^*}, z}}\right)^2 \hat{\sigma}^2$ is the corrected factor variance. Using Theorem 2.5, we get,

$$\mathcal{R}_w \stackrel{a.s.}{\sim} \frac{\sigma^2 \xi \Psi^2 \Psi^{-2} \gamma_{\beta, z}^2}{\sigma^2 \gamma_{\beta, z}^2} \stackrel{a.s.}{\sim} \xi. \quad \blacksquare$$

D Asymptotic estimates

Throughout, we assume Assumption 2.2. The estimates $(\hat{\sigma}, \hat{\beta}, \hat{\delta})$, are general (but with $\|\hat{\beta}\| = 1$) and not necessarily from PCA. In our setting, formula (7) for the minimum variance portfolio simplifies to (cf. Clarke et al. (2011))

$$w_* \propto \beta_{\text{MV}z} - \beta \quad \text{where} \quad \beta_{\text{MV}} = \frac{\sigma^2 + \delta^2}{\sigma^2 \gamma_{\beta,z}}.$$

Analogously, \hat{w} and $\hat{\beta}_{\text{MV}}$ are defined in terms of the estimates $(\hat{\sigma}, \hat{\beta}, \hat{\delta})$. The normalizing factor for the portfolio weights that ensures $1_N^\top w_* = 1$ is given by

$$W = \sqrt{N} (\beta_{\text{MV}} - \gamma_{\beta,z}). \quad (30)$$

We develop $(N \uparrow \infty)$ -asymptotics for the tracking error \mathcal{T} and variance ratio \mathcal{R} in equations (8)–(9) for our simple one-factor model. It is easy to see that

$$\mathcal{R}_{\hat{w}} = \frac{\hat{\sigma}^2 (\hat{\beta}^\top \hat{w})^2 + \hat{\delta}^2 \|\hat{w}\|_2^2}{\sigma^2 (\beta^\top \hat{w})^2 + \delta^2 \|\hat{w}\|_2^2}, \quad (31)$$

$$\mathcal{T}_{\hat{w}}^2 = \sigma^2 (\beta^\top \hat{w} - \beta^\top w_*)^2 + \delta^2 \|w_* - \hat{w}\|_2^2. \quad (32)$$

Note that all quantities involved depend on N . We suppress this dependence to ease the notation. Recall \mathcal{E} in (10) defined as

$$\mathcal{E} = \frac{\gamma_{\beta,z} - \gamma_{\beta,\hat{\beta}} \gamma_{\hat{\beta},z}}{1 - \gamma_{\hat{\beta},z}^2}.$$

Proposition D.1 (Asymptotics). *Suppose Assumption 2.2 holds. Let $\mu^2 = \sigma^2/N$ and $\hat{\mu}^2 = \hat{\sigma}^2/N$ and assume that $(\hat{\mu}, \hat{\delta})$ are bounded in N .*

(i) *Suppose that $\sup_N \gamma_{\beta,z} < 1$ and $\sup_N \gamma_{\hat{\beta},z} < 1$. Then,*

$$\begin{aligned} \mathcal{R}_{\hat{w}} &\sim \frac{\hat{\delta}^2}{\delta^2 + \mu^2 \mathcal{E}^2 N \sin^2 \theta_{\hat{\beta},z}} \\ \mathcal{T}_{\hat{w}}^2 &\sim \mu^2 \mathcal{E}^2 + \delta^2 N^{-1} \frac{(\gamma_{\hat{\beta},z}^2 - \gamma_{\beta,z}^2)}{(1 - \gamma_{\beta,z}^2)(1 - \gamma_{\hat{\beta},z}^2)} \end{aligned} \quad (N \uparrow \infty). \quad (33)$$

(ii) *Suppose that $\sup_N \gamma_{\beta,z} < 1$ and $\gamma_{\hat{\beta},z} = 1$ eventually in N . Then,*

$$\begin{aligned} \mathcal{R}_{\hat{w}} &\sim \frac{\hat{\mu}^2}{\mu^2 \gamma_{\beta,z}^2} \\ \mathcal{T}_{\hat{w}}^2 &\sim \mu^2 \gamma_{\beta,z}^2 \end{aligned} \quad (N \uparrow \infty). \quad (34)$$

(iii) Suppose that $\gamma_{\beta,z} = 1$ eventually in N and $\sup_N \gamma_{\hat{\beta},z} < 1$. Then,

$$\begin{aligned}\mathcal{R}_{\hat{w}} &\sim N^{-1} \left(\frac{\hat{\delta}^2 / \mu^2}{\sin^2 \theta_{\hat{\beta},z}} \right) \\ \mathcal{F}_{\hat{w}}^2 &= N^{-1} \left(\frac{\delta^2 \gamma_{\hat{\beta},z}^2}{\sin^2 \theta_{\hat{\beta},z}} \right)\end{aligned} \quad (N \uparrow \infty). \quad (35)$$

(iv) Suppose that both $\gamma_{\beta,z} = 1$ and $\gamma_{\hat{\beta},z} = 1$ eventually in N . Then, $\mathcal{F}_{\hat{w}}^2 = 0$ eventually in N and $\mathcal{R}_{\hat{w}} \sim \hat{\sigma}^2 / \sigma^2$ as $N \uparrow \infty$.

PROOF. All claims follow from the collection of Lemmas below. ■

It is not difficult to show the following identities in our setting.

$$\begin{aligned}\sigma^2 (\beta^\top w_*)^2 &= \mu^2 \left(\frac{\delta^2 \gamma_{\beta,z}}{\delta^2 + \sigma^2 (1 - \gamma_{\beta,z}^2)} \right)^2 \\ \sigma^2 (\beta^\top \hat{w})^2 &= \mu^2 \left(\frac{\gamma_{\beta,z} \hat{\delta}^2 + \hat{\sigma}^2 (\gamma_{\beta,z} - \gamma_{\beta,\hat{\beta}} \gamma_{\hat{\beta},z})}{\hat{\delta}^2 + \hat{\sigma}^2 (1 - \gamma_{\hat{\beta},z}^2)} \right)^2 \\ w_*^\top w_* &= \frac{(\sigma^2 + \delta^2)^2 - (\sigma^2 + 2\delta^2) \sigma^2 \gamma_{\beta,z}^2}{N(\delta^2 + \sigma^2 (1 - \gamma_{\beta,z}^2))^2} \\ \hat{w}^\top w_* &= \frac{\delta^2 + \sigma^2 - \sigma^2 \gamma_{\beta,z} E}{(\delta^2 + \sigma^2 (1 - \gamma_{\beta,z}^2)) N} \\ E &= \frac{(\hat{\delta}^2 + \hat{\sigma}^2) \gamma_{\beta,z} - \hat{\sigma}^2 \gamma_{\hat{\beta},z} \gamma_{\hat{\beta},\beta}}{\hat{\delta}^2 + \hat{\sigma}^2 (1 - \gamma_{\hat{\beta},z}^2)}\end{aligned} \quad (36)$$

These are sufficient to prove the following Lemmas. As in Proposition D.1 we set $\mu^2 = \sigma^2 / N$ and $\hat{\mu}^2 = \hat{\sigma}^2 / N$ and assume that $(\hat{\mu}, \hat{\delta})$ are bounded in N .

Lemma D.2 (True portfolio variance due to factor). *When $\sup_N \gamma_{\beta,z} < 1$ and $\sup_N \gamma_{\hat{\beta},z} < 1$, the factor components of the true variance of w_* and \hat{w} satisfy*

$$\begin{aligned}\sigma^2 (\beta^\top w_*)^2 &\sim \frac{\delta^4}{\mu^2 N^2} \left(\frac{\gamma_{\beta,z}}{1 - \gamma_{\beta,z}^2} \right)^2 \\ \sigma^2 (\beta^\top \hat{w})^2 &\sim \mu^2 \mathcal{E}^2\end{aligned} \quad (N \uparrow \infty). \quad (37)$$

Lemma D.3 (Portfolio weights). *Suppose $\sup_N \gamma_{\beta,z} < 1$ and $\sup_N \gamma_{\hat{\beta},z} < 1$.*

$$\begin{aligned} w_*^\top w_* &\sim N^{-1} \left(\frac{1}{1 - \gamma_{\beta,z}^2} \right) \\ \hat{w}^\top w_* &\sim N^{-1} \left(\frac{1 - \gamma_{\beta,z} \mathcal{E}}{1 - \gamma_{\beta,z}^2} \right) \end{aligned} \quad (N \uparrow \infty). \quad (38)$$

References

- Bai, J. & Ng, S. (2008), ‘Large dimensional factor analysis’, *Foundations and Trends® in Econometrics* **3**(2), 89–163.
- Banz, R. W. (1981), ‘The relationship between return and market value of common stock’, *Journal of Financial Economics* **9**, 3–18.
- Basu, S. (1983), ‘The relationship between earnings’ yield, market value and return for nyse common stocks’, *Journal of Financial Economics* **12**, 129–156.
- Bianchi, S. W., Goldberg, L. R. & Rosenberg, A. (2017), ‘The impact of estimation error on latent factor models of portfolio risk’, *The Journal of Portfolio Management* **43**(5), 145–156.
- Bickel, P. J. & Levina, E. (2008), ‘Covariance regularization by thresholding’, *The Annals of Statistics* pp. 2577–2604.
- Björck, Ø. & Golub, G. H. (1973), ‘Numerical methods for computing angles between linear subspaces’, *Mathematics of computation* **27**(123), 579–594.
- Black, F., Jensen, M. C. & Scholes, M. (1972), The capital asset pricing model: some empirical tests, in M. C. Jensen, ed., ‘Studies in the Theory of Capital Markets’, Praeger.
- Blume, M. E. (1975), ‘Betas and their regression tendencies’, *The Journal of Finance* **30**(3), 785–795.
- Britten-Jones, M. (1999), ‘The sampling error in estimates of mean-variance efficient portfolio weights’, *The Journal of Finance* **54**(2), 655–671.
- Bun, J., Bouchaud, J. & Potters, M. (2016), ‘Cleaning correlation matrices’, *Risk Magazine* .
- Clarke, R., De Silva, R. & Thorley, S. (2011), ‘Minimum-variance portfolio composition’, *Journal of Portfolio Management* **2**(37), 31–45.

- Connor, G. & Korajczyk, R. A. (1986), 'Performance measurement with the arbitrage pricing theory: A new framework for analysis', *Journal of financial economics* **15**, 373–394.
- Connor, G. & Korajczyk, R. A. (1988), 'Risk and return in equilibrium apt: Application of a new test methodology', *Journal of financial economics* **21**, 255–289.
- DeMiguel, V., Garlappi, L., Nogales, F. J. & Uppal, R. (2009), 'A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms', *Management Science* **55**(5), 798–812.
- DeMiguel, V., Garlappi, L. & Uppal, R. (2007), 'Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?', *The review of Financial studies* **22**(5), 1915–1953.
- El Karoui, N. (2013), 'On the realized risk of high-dimensional markowitz portfolios', *SIAM Journal on Financial Mathematics* **4**(1), 737–783.
- El Karoui, N. E. (2008), 'Spectrum estimation for large dimensional covariance matrices using random matrix theory', *The Annals of Statistics* pp. 2757–2790.
- El Karoui, N. et al. (2010), 'High-dimensionality effects in the markowitz problem and other quadratic programs with linear constraints: risk underestimation', *The Annals of Statistics* **38**(6), 3487–3566.
- Fan, J., Liao, Y. & Mincheva, M. (2013), 'Large covariance estimation by thresholding principal orthogonal complements', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75**(4), 603–680.
- Frost, P. A. & Savarino, J. E. (1986), 'An empirical bayes approach to efficient portfolio selection', *Journal of Financial and Quantitative Analysis* **21**(3), 293–305.
- Gillen, B. J. (2014), 'An empirical bayesian approach to stein-optimal covariance matrix estimation', *Journal of Empirical Finance* **29**, 402–420.
- Goldberg, L., Leshem, R. & Geddes, P. (2014), 'Restoring value to minimum variance', *Journal of Investment Management* **12**(2), 32–39.
- Goldfarb, D. & Iyengar, G. (2003), 'Robust portfolio selection problems', *Mathematics of operations research* **28**(1), 1–38.
- Green, R. C. & Hollifield, B. (1992), 'When will mean-variance efficient portfolios be well diversified?', *The Journal of Finance* **47**(5), 1785–1809.
- Jagannathan, R. & Ma, T. (2003), 'Risk reduction in large portfolios: Why imposing the wrong constraints helps', *The Journal of Finance* **58**(4), 1651–1683.

- Jobson, J. D. & Korkie, B. (1980), ‘Estimation for markowitz efficient portfolios’, *Journal of the American Statistical Association* **75**(371), 544–554.
- Lai, T. L. & Xing, H. (2008), *Statistical models and methods for financial markets*, Springer.
- Lai, T. L., Xing, H. & Chen, Z. (2011), ‘Mean-variance portfolio optimization when means and covariances are unknown’, *The Annals of Applied Statistics* pp. 798–823.
- Ledoit, O. & Péché, S. (2011), ‘Eigenvectors of some large sample covariance matrix ensembles’, *Probability Theory and Related Fields* **151**(1–2), 233–264.
- Ledoit, O. & Wolf, M. (2003), ‘Improved estimation of the covariance matrix of stock returns with an application to portfolio selection’, *Journal of empirical finance* **10**(5), 603–621.
- Ledoit, O. & Wolf, M. (2004), ‘Honey, i shrunk the sample covariance matrix’, *The Journal of Portfolio Management* **30**, 110–119.
- Markowitz, H. (1952), ‘Portfolio selection’, *The Journal of Finance* **7**(1), 77–91.
- Michaud, R. O. & Michaud, R. O. (2008), *Efficient asset management: a practical guide to stock portfolio optimization and asset allocation*, Oxford University Press.
- Morozov, A., Wang, J., Borda, L. & Menchero, J. (2012), The barra global equity model (gem3). MSCI Barra model documentation.
- Muller, M. E. (1959), ‘A note on a method for generating points uniformly on n-dimensional spheres’, *Communications of the ACM* **2**(4), 19–20.
- Onatski, A. (2012), ‘Asymptotics of the principal components estimator of large factor models with weakly influential factors’, *Journal of Econometrics* **168**(2), 244–258.
- Pástor, L. (2000), ‘Portfolio selection and asset pricing models’, *The Journal of Finance* **55**(1), 179–223.
- Paul, D. (2007), ‘Asymptotics of sample eigenstructure for a large dimensional spiked covariance model’, *Statistica Sinica* pp. 1617–1642.
- Ross, S. A. (1976), ‘The arbitrage theory of capital asset pricing’, *Journal of economic theory* **13**(3), 341–360.
- Sharpe, W. F. (1964), ‘Capital asset prices: A theory of market equilibrium under conditions of risk’, *The Journal of Finance* **19**(3), 425–442.

- Shen, D., Shen, H., Zhu, H. & Marron, S. (2016), ‘The statistics and mathematics of high dimensional low sample size asymptotics’, *Statistica Sinica* **26**(4), 1747–1770.
- Treynor, J. L. (1962), Toward a theory of market value of risky assets. Presented to the MIT Finance Faculty Seminar.
- Vasicek, O. A. (1973), ‘A note on using cross-sectional information in bayesian estimation of security betas’, *The Journal of Finance* **28**(5), 1233–1239.
- Wang, W. & Fan, J. (2017), ‘Asymptotics of empirical eigenstructure for high dimensional spiked covariance’, *The Annals of Statistics* **45**(3), 1342–1374.
- Yata, K. & Aoshima, M. (2012), ‘Effective pca for high-dimension, low-sample-size data with noise reduction via geometric representations’, *Journal of multivariate analysis* **105**(1), 193–215.