Better Betas

Correcting a large error in minimum variance portfolios

Lisa Goldberg *(with Alex Papanicolaou & Alex Shkolnik)*
lrg@berkeley.edu

9th Western Conference on Mathematical Finance
November 16, 2018

Consortium for Data Analytics in Risk
University of California, Berkeley

Optimized portfolios and the impact of estimation error

Optimized portfolios

Since Markowitz (1952), quantitative investors have constructed portfolios with mean-variance optimization.

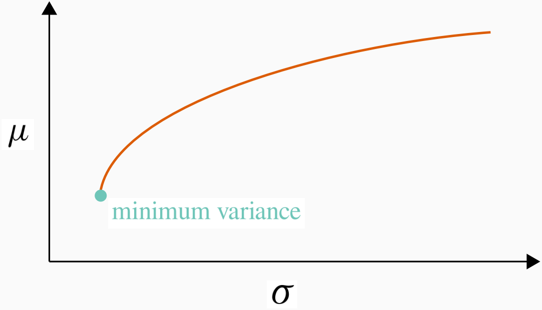

A simple quadratic program given a mean vector m and a covariance matrix $\mathbf{\Sigma}$.

The impact of estimation error

In practice, optimization relies on an estimate of the mean and covariance matrix ($\widehat{\mathbf{\Sigma}}$ estimates $\mathbf{\Sigma}$).

Estimation error leads to two types of errors:

- You get the wrong portfolio: Estimation error distorts portfolio weights so optimized portfolios are never optimal.
- And it's probably riskier than you think: A risk-minimizing optimization tends to materially underforecast portfolio risk.

We measure and correct both errors in simulation.

Two metrics for the impact of estimation error in simulation

Errors in weights

(Squared) tracking error of an optimized portfolio $\widehat{w}$ measures its distance from the optimal portfolio w_*:

$$\mathscr{T}_{\widehat{w}}^2 = (\widehat{w} - w_*)^\top \boldsymbol{\Sigma} (\widehat{w} - w_*)$$

Tracking error is the width of the distribution of return differences between w and $\widehat{w}$.

Ideally, tracking error should be close to 0.

Errors in risk forecasts

Variance forecast ratio measures the error in the risk forecast as:

$$\mathscr{R}_{\widehat{w}} = \frac{\widehat{w}^\top \widehat{\boldsymbol{\Sigma}} \widehat{w}}{\widehat{w}^\top \boldsymbol{\Sigma} \widehat{w}}$$

This is a ratio of the estimated portfolio risk over the actual risk of the estimated minimum variance portfolio $\widehat{w}$.

Ideally, the variance forecast ratio should be close to 1.

Error metrics in simulation

In simulation (given a model for $\mathbf{\Sigma}$),

- generate security returns and compute w_* using $\mathbf{\Sigma}$,
  ($\mathbf{\Sigma}$ is accessible in simulation)
- estimate $\mathbf{\Sigma}$ by $\widehat{\mathbf{\Sigma}}$ from observed returns and compute $\widehat{w}$,
- measure the error metrics $\mathscr{T}^2_{\widehat{w}}$ and $\mathscr{R}_{\widehat{w}}$.

Minimum variance

Why minimum variance?

Theory

Error amplification: Highly sensitive to estimation error.

Error isolation: Impervious to errors in expected return.

Insight into a general problem: Informs our understanding of how estimation error distorts portfolios and points to a remedy.

Practice

Large investments: For example, the iShares Edge MSCI Min Vol USA ETF (ticker USMV) had net assets of roughly $14 billion on Sept. 28, 2018.

True and estimated minimum variance portfolios

The true minimum variance portfolio w_* is the solution to:

$$\min_{w \in \mathbb{R}^N} w^\top \mathbf{\Sigma} w$$
$$w^\top \mathbf{1_N} = \mathbf{1}.$$

In practice, we construct an estimated minimum variance portfolio, $\widehat{w}$, that solves the same problem with $\widehat{\mathbf{\Sigma}}$ replacing $\mathbf{\Sigma}$.

Covariance matrix estimation

Structure and dimension reduction

The simplest candidate for $\widehat{\boldsymbol{\Sigma}}$ is the sample return covariance matrix, $\mathbf{S}$.

- But it is useless for portfolio construction since it is highly singular.
- The number of parameters tends to (vastly) exceed the number of observations.
- Still, it can serve as a starting point for constructing better estimates.

Factor models add (empirically sound) structure to a covariance matrix and reduce the number of required parameters to a manageable level.

One-factor model of security returns

The return generating process for N securities is specified by

$$R = \phi\beta + \epsilon$$

where ϕ is the return to a market factor, β is the N-vector of factor exposures, ϵ is the N-vector of diversifiable specific returns.

The (ϕ, ϵ) are latent variables. We observe T i.i.d. returns to N securities, i.e., $R_1, R_2, \ldots, R_T$.

One-factor model covariance matrix

When the ϕ and ϵ are uncorrelated (as we assume), the security covariance matrix can be expressed as

$$\mathbf{\Sigma} = \sigma^2 \beta\beta^\top + \mathbf{\Delta},$$

where σ^2 is the variance of the market factor and the diagonal entries of $\mathbf{\Delta}$ are specific variances, δ^2.

Assumption 1. *$\sigma^2/N \to \mu_\infty \in (0, \infty)$ and $\mathbf{\Delta} = \delta^2 \mathbf{I}$.*

Assumption 2. *$\{R_i\}_{i=1}^T$ are i.i.d. with $R_1 \sim \mathcal{N}(0, \mathbf{\Sigma})$.*

Assumption 3. *β always has some dispersion.*

Estimation error in factor models

In practice, we have only the estimates $\hat{\sigma}$, $\hat{\beta}$ and $\hat{\delta}$.

$$\widehat{\mathbf{\Sigma}} = \hat{\sigma}^2 \hat{\beta} \hat{\beta}^\top + \widehat{\mathbf{\Delta}}$$

We measure the errors in estimated parameters, of course.

But our focus is how errors in parameter estimates affect portfolio metrics: (squared) tracking error and variance forecast ratio.

Factor model estimation with PCA

PCA: compute the sample covariance matrix $\mathbf{S}$ and set:

- $\hat{\beta}$ – first eigenvector of $\mathbf{S}$,
- $\hat{\sigma}^2$ – largest eigenvalue of $\mathbf{S}$,
- $\hat{\delta}^2$ – OLS regression of returns on the estimated factor.

PCA approximates true factors well for N large and $\widehat{\mathbf{\Sigma}} = \mathbf{\Sigma}$.

Sample eigenvectors behave differently for N large and T fixed *(i.e., in the $N \uparrow \infty$ and T fixed asymptotic regime)*.

Current techniques adjust only the eigenvalue $\hat{\sigma}^2$ (typically biased upward). No (direct) corrections of the eigenvector, $\hat{\beta}$, are available.

Motivating example

Parameter errors → portfolio metric errors

A large literature on random matrix theory identifies and corrects biases in estimated eigenvalues.

It turns out, however, that in a simple PCA model, portfolio metrics for a minimum variance portfolio are insensitive to errors in eigenvalues.

But errors in the dominant eigenvector make a material difference.

Selective error correction in minimum variance

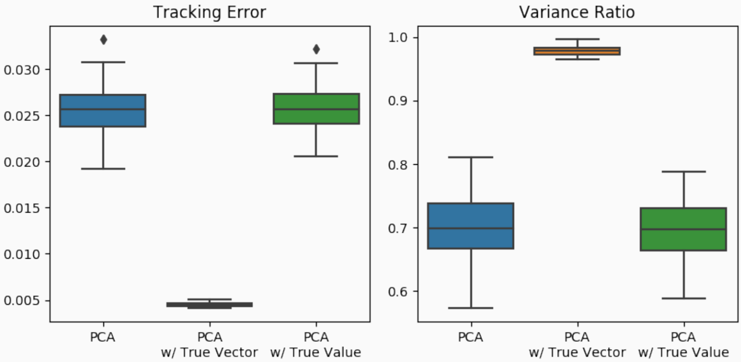

Results communicated by Stephen Bianchi

The dispersion bias

No reference frame to detect eigenvector bias...

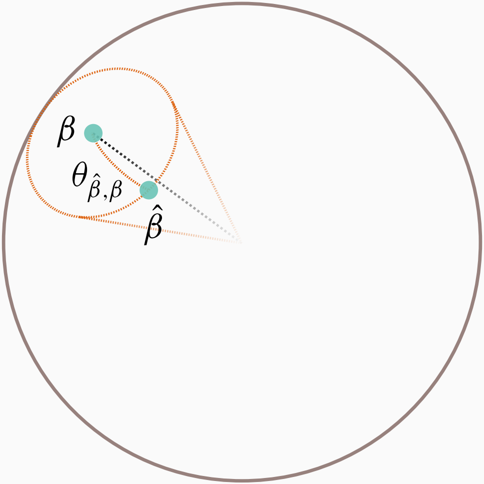

..unless the eigenvector is (mostly) positive

Let $z = 1_N/\sqrt{N}$ (a vector on the unit N-sphere).

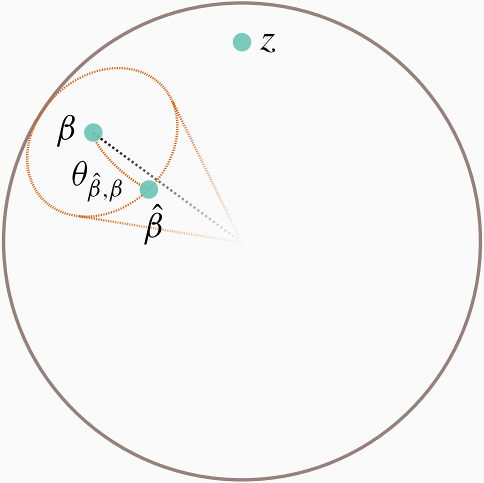

Let $z = 1_N/\sqrt{N}$ (a vector on the unit N-sphere).

This is the unique (up to negation) *dispersionless* unit vector.

Theorem

Let $\hat{\beta}$ be a PCA-estimate of β. Then,

$$\cos\theta_{\beta,z} \overset{a.s.}{\sim} \psi_T \cos\theta_{\hat{\beta},z} \qquad (N \uparrow \infty) \tag{1}$$

with $\psi_T > 1$. In other wordswords, $\theta_{\hat{\beta},z}$ is larger than $\theta_{\beta,z}$ with high probability for N large.

PCA bias illustration

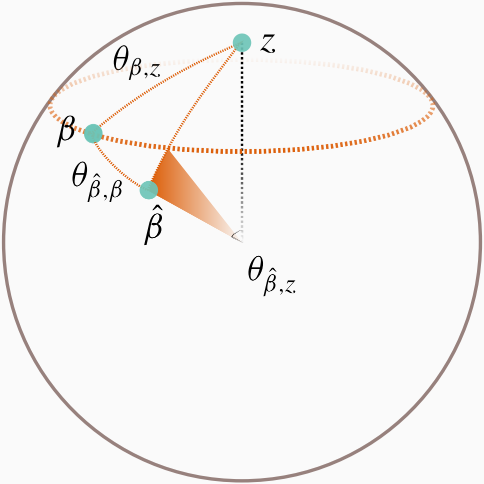

Error metrics for the minimum variance portfolio

Define $\gamma_{x,y} = x^\top y$ (on unit sphere $\gamma_{x,y} = \cos\theta_{x,y}$) and

$$\mathscr{E}_{\hat{\beta}} = \frac{\gamma_{\beta,z} - \gamma_{\beta,\hat{\beta}}\gamma_{\hat{\beta},z}}{\sin\theta_{\hat{\beta},z}}. \tag{2}$$

The variable $\mathscr{E}_{\hat{\beta}}$ drives all the error in our metrics.

We prove, $\mathscr{E}_{\hat{\beta}} > 0$ for the PCA-estimate $\hat{\beta}$.

Error metrics for the minimum variance portfolio

As $N \uparrow \infty$, (for any $\hat{\beta}$ such that $\inf_N \mathscr{E}^2_{\hat{\beta}} > 0$)

$$\mathscr{T}^2_{\hat{w}} \sim \frac{\mu^2 \, \mathscr{E}^2_{\hat{\beta}}}{\sin^2 \theta_{\hat{\beta},z}} \qquad \mathscr{R}_{\hat{w}} \sim \frac{\hat{\delta}^2 N^{-1}}{\mu^2 \, \mathscr{E}^2_{\hat{\beta}}}. \tag{3}$$

Remarkable: *No dependence on the eigenvalue estimate $\hat{\sigma}^2$.*
(above, $\mu^2 = \sigma^2/N$)

Bias correction

Dispersion bias correction

We propose a correction $\hat{\beta}^*$ of the form,

$$\hat{\beta}^* \propto \hat{\beta} + \rho z \qquad \rho \in \mathbb{R}.$$

We consider two estimators (i.e., values of ρ)

$$\rho_1 = \frac{\gamma_{\beta,z} - \gamma_{\hat{\beta},z}\gamma_{\hat{\beta},\beta}}{\gamma_{\hat{\beta},\beta} - \gamma_{\hat{\beta},z}\gamma_{\beta,z}} \qquad \text{(oracle, } \mathscr{E}_{\hat{\beta}^*} = 0\text{)}. \tag{4}$$

$$\rho_2 = \frac{q\gamma_{\hat{\beta},z}}{1 - (q\gamma_{\hat{\beta},z})^2}(q - q^{-1}) \qquad \text{(data-driven, } \mathscr{E}_{\hat{\beta}^*} \approx 0\text{)}. \tag{5}$$

where q is computed from observed data only.

Dispersion bias correction

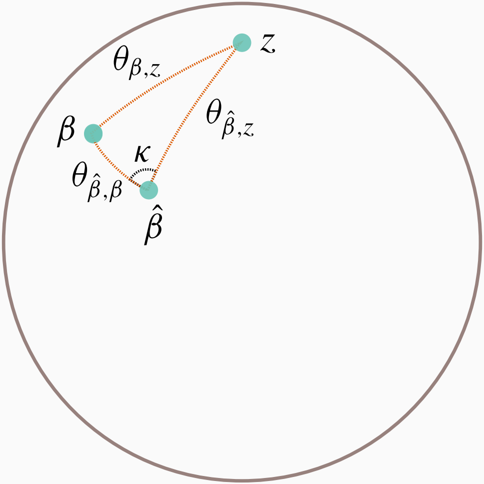

Dispersion bias correction

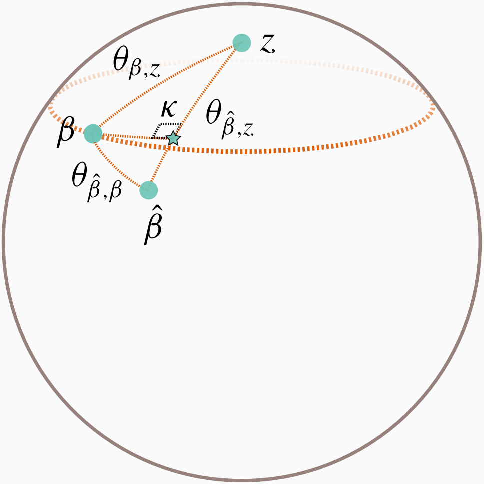

Main result

Theorem

Under our assumptions, the oracle estimator achieves

$$\mathscr{R}_{\hat{w}} \overset{a.s.}{\sim} \hat{\delta}^2/\delta^2 \qquad \mathscr{T}_{\hat{w}}^2 \overset{a.s.}{\sim} O(N^{-1}). \tag{6}$$

Complementary results available in paper (online).

Numerical simulation

Study outline

Even though our theory is worked out only for a one-factor model, we experiment with security returns simulated with an empirically guided four-factor model.

Calibrate to calm and stressed regimes.

Compare the accuracy of long-only minimum variance portfolio weights and risk forecasts using the dispersion bias correction and standard models, including vanilla PCA, beta shrinkage based on Blume (1975), and covariance matrix shrinkage based on Ledoit & Wolf (2004) and Ledoit & Wolf (2017).

Calibrating a four-factor model

Parameter	Calm	Stressed
market volatility	16%	32%
style volatility	$[4\%, 8\%]$	$[4\%, 8\%]$
specific volatility	$[32\%, 64\%]$	$[48\%, 96\%]$
beta dispersion	0.95	0.93

Accuracy of long-only minimum variance portfolio weights

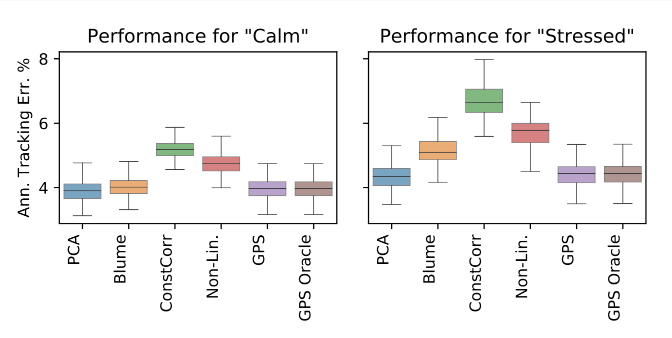

Simulation based on 100 paths, $N = 500$, $T = 250$

Risk forecast accuracy for long-only minimum variance portfolios

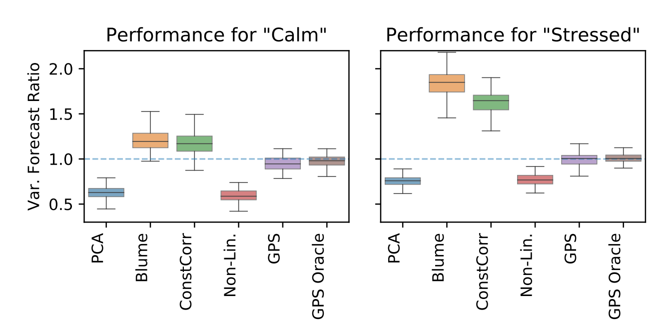

Simulation based on 100 paths, $N = 500$, $T = 250$

True volatility of long-only minimum variance portfolios

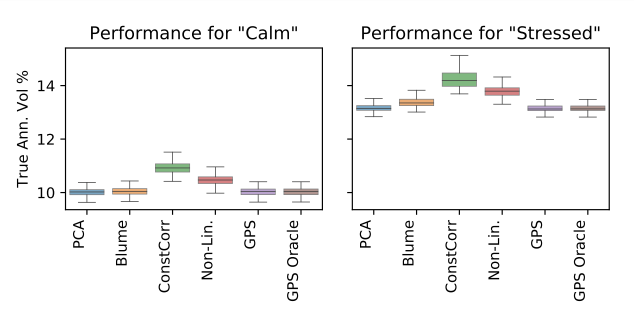

Simulation based on 100 paths, $N = 500$, $T = 250$

Summary

We provided a characterization of a systematic (dispersion) bias in PCA factors (sample eigenvectors).

– *applicable in many other settings.*

Developed and tested oracle and data-driven corrections to mitigate this dispersion bias (distinct from literature).

Emphasized portfolio-based (practitioner-oriented) metrics for evaluating the impact of our bias correction.

Our results can be viewed as an extension and formalization to PCA of ideas that have been known by practitioners since the 1970s.

Ongoing research

Extend theory to account for what we see in simulation: allow for multiple factors, non-Gaussian distributions and temporal dependency.

Evaluate the dispersion bias correction on empirical data.

Explore bias corrections for non-dominant PCA factors.

Thank You

References i

Blume, M. E. (1975), 'Betas and their regression tendencies', *The Journal of Finance* **30**(3), 785–795.

Clarke, R., De Silva, R. & Thorley, S. (2011), 'Minimum-variance portfolio composition', *Journal of Portfolio Management* **2**(37), 31–45.

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.

Ledoit, O. & Wolf, M. (2004), 'Honey, i shrunk the sample covariance matrix', *The Journal of Portfolio Management* **30**, 110–119.

Ledoit, O. & Wolf, M. (2017), 'Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets goldilocks', *The Review of Financial Studies* **30**(12), 4349–4388.

Markowitz, H. (1952), 'Portfolio selection', *The Journal of Finance* **7**(1), 77–91.

Shen, D., Shen, H., Zhu, H. & Marron, S. (2016), 'The statistics and mathematics of high dimensional low sample size asympotics', *Statistica Sinica* **26**(4), 1747–1770.

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.