ESTIMATION ERROR AND WHAT TO DO ABOUT IT

Northfield's 29th Annual Research Conference

Mark Kritzman
March 20, 2017

Contents

- Traditional approaches to estimation error
- Types of estimation error
  - Small-sample error
  - Independent-sample error
  - Mapping error
  - Interval error
- Stability-adjusted return distribution
- Implications for portfolio construction
- Asset allocation examples

Traditional approaches to estimation error

1. Equal Weighting

   This approach avoids estimation altogether and builds an equally weighted portfolio.

2. Bayesian shrinkage

   This approach compresses estimates toward a prior belief such as the cross-sectional average.

3. Resampling

   This approach repeatedly draws random samples from the data, generates efficient portfolios for each sample, and averages the weights for the portfolios at a chosen risk level.

4. Robust optimization

   This approach considers a wide range of expected returns and risk and selects the portfolio that suffers the least in the most adverse scenario.

Types of estimation error

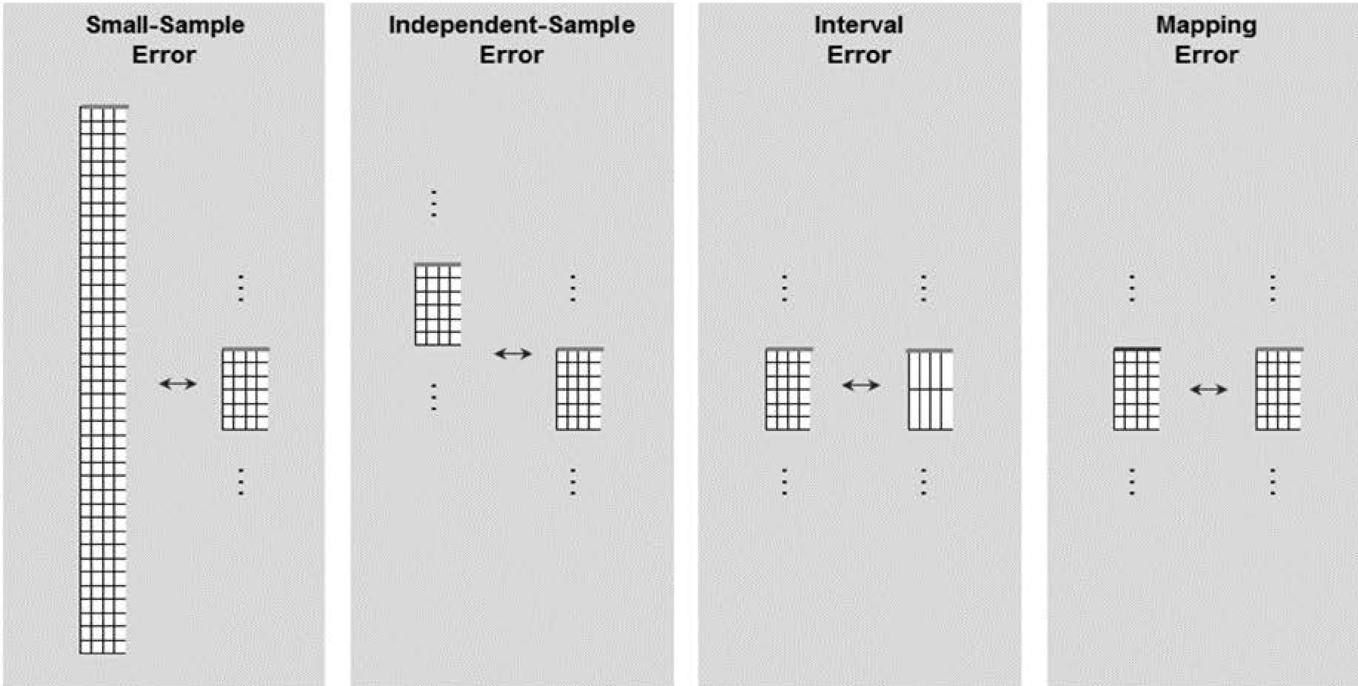

Vary: **Small samples**

Hold constant:
- Forecasting sample
- Factor mapping
- Measurement interval

Vary: **Forecasting sample**

Hold constant:
- Sample size
- Factor mapping
- Measurement interval

Vary: **Measurement interval**

Hold constant:
- Sample size
- Forecasting sample
- Factor mapping

Vary: **Factor mapping**

Hold constant:
- Sample size
- Forecasting sample
- Measurement interval

Small-sample error

The realization of parameters from a small sample will likely differ from the parameter values of a large sample from which it is selected.

We call this small-sample error.

$$SSE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}} - \rho_{AB,m}\sqrt{\sigma_{A,m}\sigma_{B,m}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2}$$

When A and B are the same asset, this formula will measure the error in the standard deviation of that asset.

Small-sample error

The realization of parameters from a small sample will likely differ from the parameter values of a large sample from which it is selected.

We call this small-sample error.

$$SSE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}} - \rho_{AB,m}\sqrt{\sigma_{A,m}\sigma_{B,m}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2}$$

m, j indicates monthly estimates from a 36-month testing subsample

m alone indicates monthly estimates from the full sample

When A and B are the same asset, this formula will measure the error in the standard deviation of that asset.

Independent-sample error

The realization of parameters from a future sample will likely differ from the parameter values of an independent historical sample.

We call this independent-sample error.

$$ISE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}} - \rho_{AB,\hat{m},j}\sqrt{\sigma_{A,\hat{m},j}\sigma_{B,\hat{m},j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2 - SSE(A,B)^2}$$

Notes: In rare instances, small sampler error exceeds the mean squared error between two independent samples. In these cases, we record small sample error as the mean squared error between the two samples, and record independent sample error as zero.

Independent-sample error

The realization of parameters from a future sample will likely differ from the parameter values of an independent historical sample.

We call this independent-sample error.

$$ISE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}} - \rho_{AB,\hat{m},j}\sqrt{\sigma_{A,\hat{m},j}\sigma_{B,\hat{m},j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2 - SSE(A,B)^2}$$

m, j indicates monthly estimates from a 36-month testing subsample

$\hat{m}, j$ indicates monthly estimates from a 36-month independent subsample immediately preceding the testing subsample

Notes: In rare instances, small sampler error exceeds the mean squared error between two independent samples. In these cases, we record small sample error as the mean squared error between the two samples, and record independent sample error as zero.

Mapping error

Assets define the opportunity set for investing. A desired factor exposure must be mapped onto a portfolio of assets to be investable.

The mapping that best tracks a factor in the future will likely differ from the mapping of an independent historical sample.

We call this mapping error.

Mapping error applies only to factors.

The calculations for small sample error and independent sample error for factors do not reflect mapping error.

Mapping error

Mapping error can be isolated by comparing the covariance of the best-fit factor-mimicking portfolio to the covariance of a factor-mimicking portfolio estimated from an independent sample, holding the evaluation period constant.

$$ME(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{\hat{A}\hat{B},m,j}\sqrt{\sigma_{\hat{A},m,j}\sigma_{\hat{B},m,j}} - \rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2}$$

Mapping error

Mapping error can be isolated by comparing the covariance of the best-fit factor-mimicking portfolio to the covariance of a factor-mimicking portfolio estimated from an independent sample, holding the evaluation period constant.

$$ME(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{\hat{A}\hat{B},m,j}\sqrt{\sigma_{\hat{A},m,j}\sigma_{\hat{B},m,j}} - \rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2}$$

$\hat{A}$, $\hat{B}$ indicate factor mappings derived from the independent subsample

A, B indicate factor mappings derived from the testing subsample

Interval error

Parameters estimated from monthly or higher-frequency returns often differ from those estimated from lower-frequency returns, within the same sample.

We call this interval error.

Interval error

U.S. and Emerging Markets Stocks:

Monthly Returns - January 1988 through December 2015

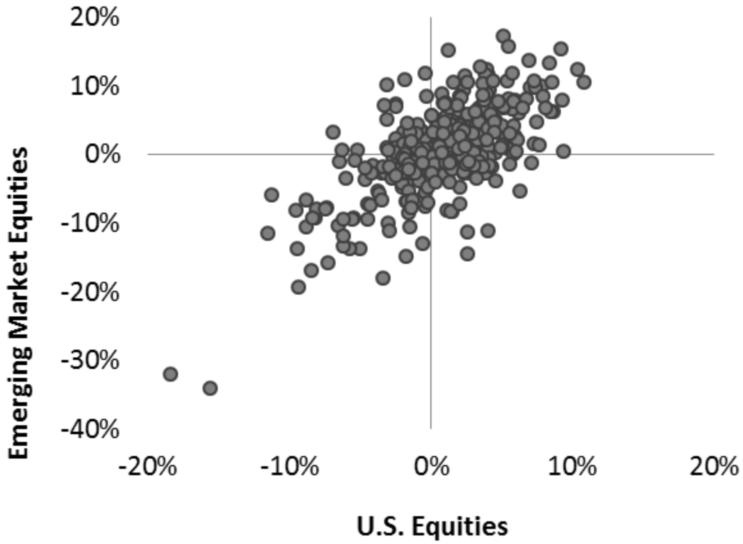

Correlation = 0.68

U.S. and Emerging Markets Stocks:
Five-Year Returns - January 1988 through December 2015

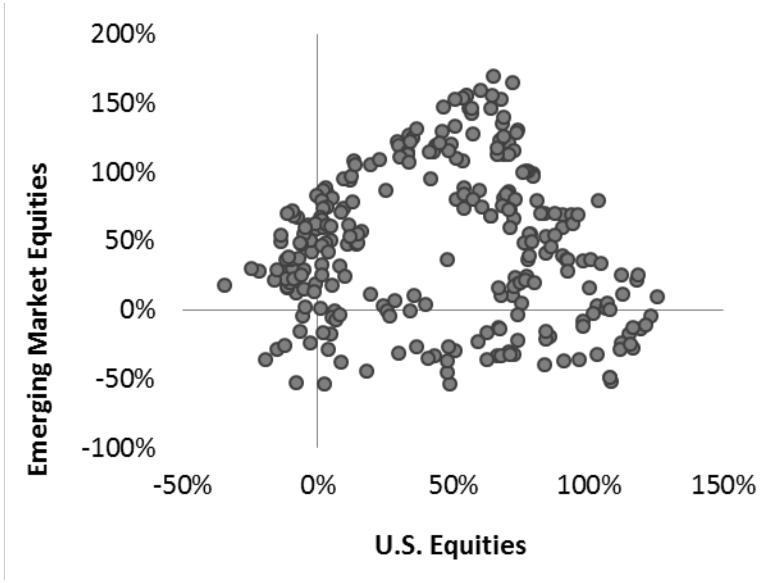

Correlation = -0.04

Interval error: Long-horizon and short-horizon volatility

The volatility of the cumulative continuous returns of x over q periods is given by:

$$\sigma(x_t + \cdots + x_{t+q-1}) = \sigma_x \sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t, x_{t+k}}}$$

Interval error: Long-horizon and short-horizon volatility

The volatility of the cumulative continuous returns of x over q periods is given by:

$$\sigma(x_t + \cdots + x_{t+q-1}) = \sigma_x \sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t, x_{t+k}}}$$

This term reflects annualization in the absence of lagged effects

Interval error: Long-horizon and short-horizon volatility

The volatility of the cumulative continuous returns of x over q periods is given by:

$$\sigma(x_t + \cdots + x_{t+q-1}) = \sigma_x \sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t, x_{t+k}}}$$

This term captures the impact of auto-correlation

Interval error: Long-horizon and short-horizon correlation

The correlation between the cumulative returns of x and the cumulative returns of y over q periods is given by:

$$\rho(x_t + \cdots + x_{t+q-1}, y_t + \cdots + y_{t+q-1}) =$$

$$\frac{q\rho_{x_t, y_t} + \sum_{k=1}^{q-1}(q-k)(\rho_{x_{t+k}, y_t} + \rho_{x_t, y_{t+k}})}{\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t, x_{t+k}}}\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{y_t, y_{t+k}}}}$$

Interval error: Long-horizon and short-horizon correlation

The correlation between the cumulative returns of x and the cumulative returns of y over q periods is given by:

$$\rho(x_t + \cdots + x_{t+q-1}, y_t + \cdots + y_{t+q-1}) =$$

$$\frac{q\rho_{x_t,y_t} + \sum_{k=1}^{q-1}(q-k)(\rho_{x_{t+k},y_t} + \rho_{x_t,y_{t+k}})}{\sqrt{q+2\sum_{k=1}^{q-1}(q-k)\rho_{x_t,x_{t+k}}}\sqrt{q+2\sum_{k=1}^{q-1}(q-k)\rho_{y_t,y_{t+k}}}}$$

This term captures the lagged cross-correlation between x and y

The correlation between the cumulative returns of x and the cumulative returns of y over q periods is given by:

$$\rho(x_t + \cdots + x_{t+q-1}, y_t + \cdots + y_{t+q-1}) =$$

$$\frac{q\rho_{x_t,y_t} + \sum_{k=1}^{q-1}(q-k)(\rho_{x_{t+k},y_t} + \rho_{x_t,y_{t+k}})}{\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t,x_{t+k}}}\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{y_t,y_{t+k}}}}$$

This term captures the auto-correlation of x

This term captures the auto-correlation of y

Interval error

Interval error can be isolated by comparing a covariance matrix estimated from low-frequency returns to a covariance matrix estimated from high-frequency returns.

$$IE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,ann,j}\sqrt{\sigma_{A,ann,j}\sigma_{B,ann,j}}/12 - \rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2}$$

Interval error

Interval error can be isolated by comparing a covariance matrix estimated from low-frequency returns to a covariance matrix estimated from high-frequency returns.

$$IE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,ann,j}\sqrt{\sigma_{A,ann,j}\sigma_{B,ann,j}}/12 - \rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2}$$

tri, j indicates implied 3-year estimates from a 36-month testing subsample

m, j indicates monthly estimates from the same 36-month testing sample

Composite instability score

The four sources of estimation error are independent from one another, which means we can sum the variances of each error and then take the square root of this sum to compute a composite instability score.

$$Composite\ Instability\ Score\ (CIS) = \sqrt{SSE^2 + ISE^2 + ME^2 + IE^2}$$

Stability-adjusted return distribution

Begin with a long sample of asset returns					
Estimate small sample covariance matrices	Σ_{s1}	Σ_{s2}	Σ_{s3}	…	Σ_{sn}
Compute error matrix for each small sample versus its complementary sample	$\Sigma_{e1} = \Sigma_{s1} - \Sigma_{c1}$	$\Sigma_{e2} = \Sigma_{s2} - \Sigma_{c2}$	$\Sigma_{e3} = \Sigma_{s3} - \Sigma_{c3}$	…	$\Sigma_{en} = \Sigma_{sn} - \Sigma_{cn}$
Add each error matrix to the baseline covariance matrix	$\Sigma_1 = \Sigma_{e1} + \Sigma_{base}$	$\Sigma_2 = \Sigma_{e2} + \Sigma_{base}$	$\Sigma_3 = \Sigma_{e3} + \Sigma_{base}$	…	$\Sigma_n = \Sigma_{en} + \Sigma_{base}$
Draw sample returns from normal distributions				…	
Combine to form a composite non-normal distribution					

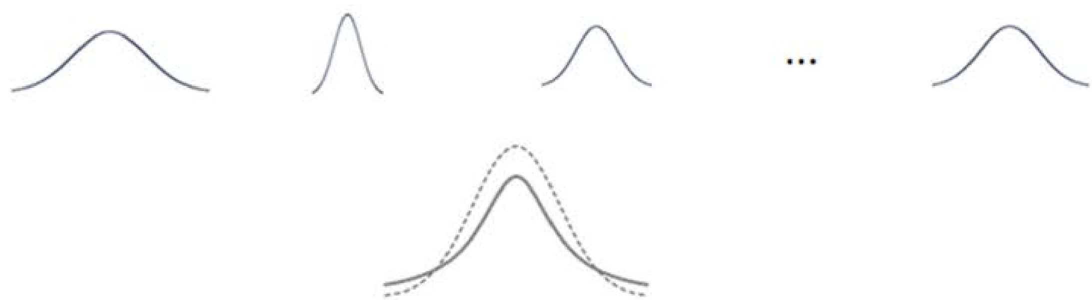

Implications for portfolio construction

- Mean-variance analysis assumes either that investors have preferences that can be well approximated by mean and variance or that returns are elliptically distributed.
- If sub-samples within a large sample have different correlations or if kurtosis is not uniform across asset classes, the return distribution will not be elliptical.
- Some investors face thresholds and therefore have preferences that are better represented by a kinked utility function which displays sharp aversion to losses below a given threshold. Mean and variance are not sufficient to describe kinked utility functions.
- If returns are not elliptical and investors have preferences than cannot be approximated by mean and variance, it may be preferable to employ full-scale optimization to identify the optimal portfolio.

Implications for portfolio construction

Full-scale optimization: Illustrative example

	Stock Return	Bond Return	Stock Weight	Bond Weight	Utility Calculation	Utility
1993	10.1%	16.2%	100%	0%	Ln [(1 + 10.1%) x 100% + (1 + 16.2%) x 0%] =	0.0962
1994	1.3%	-7.1%	100%	0%	Ln [(1 + 1.3%) x 100% + (1 – 7.1%) x 0%] =	0.0129
1995	37.5%	30.0%	100%	0%	Ln [(1 + 39.5%) x 100% + (1 + 30.0%) x 0%] =	0.3185
1996	22.9%	0.1%	100%	0%	Ln [(1 + 22.9%) x 100% + (1 + 0.1%) x 0%] =	0.2062
1997	33.3%	14.5%	100%	0%	Ln [(1 + 33.3%) x 100% + (1 + 14.5%) x 0%] =	0.2874
1998	28.6%	11.8%	100%	0%	Ln [(1 + 28.6%) x 100% + (1 + 11.8%) x 0%] =	0.2515
1999	20.9%	-7.6%	100%	0%	Ln [(1 + 20.9%) x 100% + (1 – 7.6%) x 0%] =	0.1898
2000	-9.1%	16.1%	100%	0%	Ln [(1 – 9.1%) x 100% + (1 + 16.2%) x 0%] =	-0.0954
2001	-11.9%	7.3%	100%	0%	Ln [(1 – 11.9%) x 100% + (1 + 7.3%) x 0%] =	-0.1267
2002	-22.1%	14.8%	100%	0%	Ln [(1 – 22.1%) x 100% + (1 + 14.8%) x 0%] =	-0.2497

Average utility = **0.0891**

Implications for portfolio construction

Full-scale optimization: Illustrative example

	Stock Return	Bond Return	Stock Weight	Bond Weight	Utility Calculation	Utility
1993	10.1%	16.2%	50%	50%	Ln [(1 + 10.1%) x 50% + (1 + 16.2%) x 50%] =	0.1235
1994	1.3%	-7.1%	50%	50%	Ln [(1 + 1.3%) x 50% + (1 – 7.1%) x 50%] =	-0.0294
1995	37.5%	30.0%	50%	50%	Ln [(1 + 39.5%) x 50% + (1 + 30.0%) x 50%] =	0.2908
1996	22.9%	0.1%	50%	50%	Ln [(1 + 22.9%) x 50% + (1 + 0.1%) x 50%] =	0.1089
1997	33.3%	14.5%	50%	50%	Ln [(1 + 33.3%) x 50% + (1 + 14.5%) x 50%] =	0.2143
1998	28.6%	11.8%	50%	50%	Ln [(1 + 28.6%) x 50% + (1 + 11.8%) x 50%] =	0.1840
1999	20.9%	-7.6%	50%	50%	Ln [(1 + 20.9%) x 50% + (1 – 7.6%) x 50%] =	0.0644
2000	-9.1%	16.1%	50%	50%	Ln [(1 – 9.1%) x 50% + (1 + 16.2%) x 50%] =	0.0344
2001	-11.9%	7.3%	50%	50%	Ln [(1 – 11.9%) x 50% + (1 + 7.3%) x 50%] =	-0.0233
2002	-22.1%	14.8%	50%	50%	Ln [(1 – 22.1%) x 50% + (1 + 14.8%) x 50%] =	-0.0372

Average utility = **0.0930**

Full-scale optimization: Illustrative example

	Stock Return	Bond Return	Stock Weight	Bond Weight	Utility Calculation	Utility
1993	10.1%	16.2%	55%	45%	Ln [(1 + 10.1%) x 55% + (1 + 16.2%) x 45%] =	0.1208
1994	1.3%	-7.1%	55%	45%	Ln [(1 + 1.3%) x 55% + (1 – 7.1%) x 45%] =	-0.0251
1995	37.5%	30.0%	55%	45%	Ln [(1 + 39.5%) x 55% + (1 + 30.0%) x 45%] =	0.2936
1996	22.9%	0.1%	55%	45%	Ln [(1 + 22.9%) x 55% + (1 + 0.1%) x 45%] =	0.1190
1997	33.3%	14.5%	55%	45%	Ln [(1 + 33.3%) x 55% + (1 + 14.5%) x 45%] =	0.2219
1998	28.6%	11.8%	55%	45%	Ln [(1 + 28.6%) x 55% + (1 + 11.8%) x 45%] =	0.1910
1999	20.9%	-7.6%	55%	45%	Ln [(1 + 20.9%) x 55% + (1 – 7.6%) x 45%] =	0.0777
2000	-9.1%	16.1%	55%	45%	Ln [(1 – 9.1%) x 55% + (1 + 16.2%) x 45%] =	0.0222
2001	-11.9%	7.3%	55%	45%	Ln [(1 – 11.9%) x 55% + (1 + 7.3%) x 45%] =	-0.0331
2002	-22.1%	14.8%	55%	45%	Ln [(1 – 22.1%) x 55% + (1 + 14.8%) x 45%] =	-0.0565
					Average utility =	**0.0931**

Implications for portfolio construction

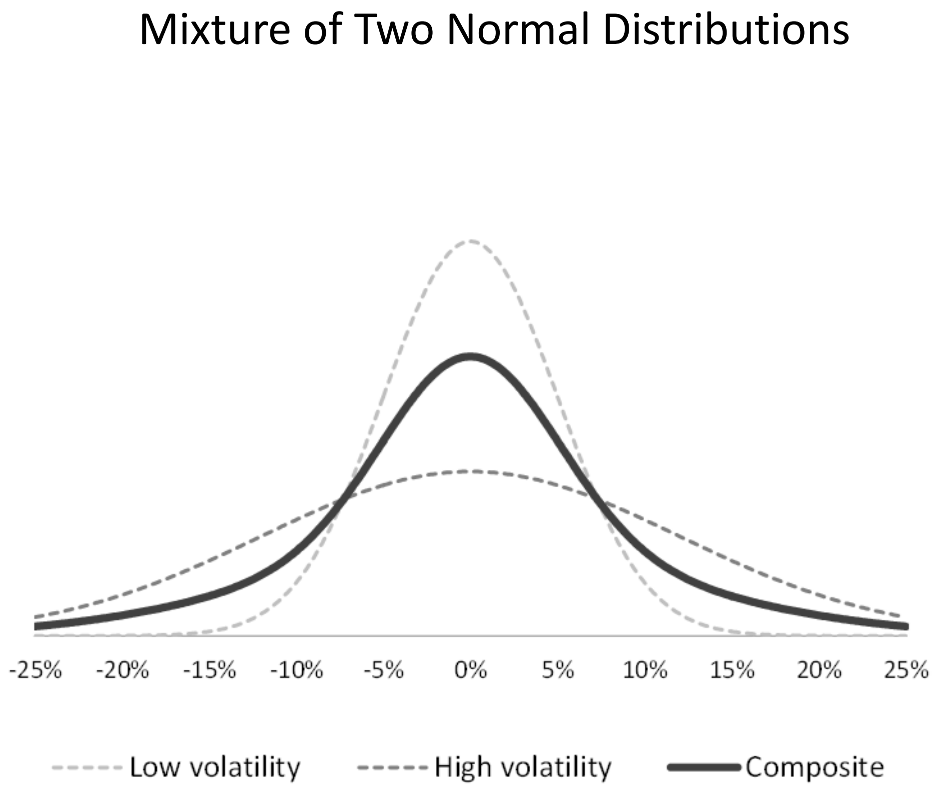

Multivariate Mixture of Asset Classes with Unstable Correlation

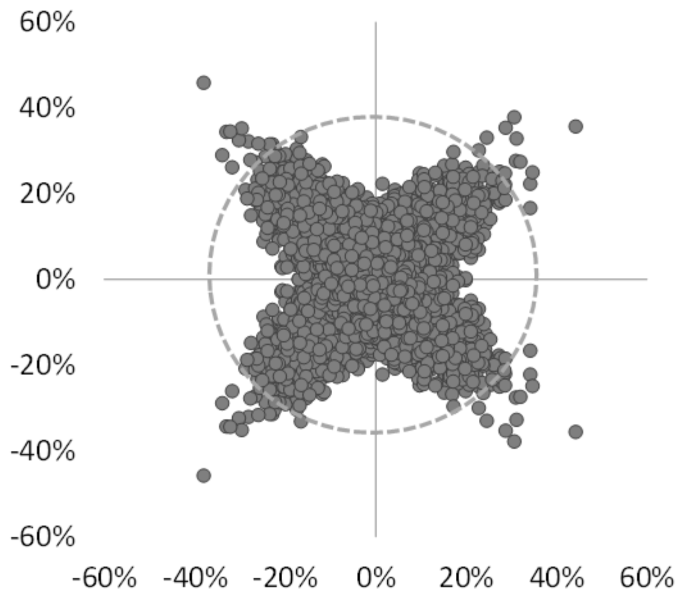

Asset allocation examples - methodology

- We select the first five-year subsample from the full sample and set it aside.
- We then build a stability-adjusted return sample using the remaining data in the original sample.
- Next we use this complementary sample to build six portfolios: a portfolio that ignores errors, one that applies Bayesian shrinkage, and one formed from the stability-adjusted return sample, all using full-scale optimization, and then again using mean-variance analysis.
- We repeat steps 1 through 3 for all 36 testing samples, which are overlapping periods ending in December.
- Using a variety of metrics, we evaluate each portfolio in the subsample that was held out of the complementary sample used to form it.
- We shrink the standard deviations by blending them equally with their cross-sectional mean, and we do the same for the correlations.

Asset allocation using full-scale optimization

Average Optimal Weights	Ignoring Errors	Bayesian Shrinkage	Stability Adjusted
U.S. Equities	25.0%	20.1%	33.3%
Foreign Developed Market Equities	19.6%	18.9%	16.7%
Emerging Market Equities	13.6%	17.2%	11.8%
Treasury Bonds	20.1%	13.9%	27.9%
U.S. Corporate Bonds	15.6%	15.3%	7.9%
Commodities	3.8%	5.7%	1.4%
Cash Equivalents	2.4%	8.9%	1.0%
10 Percentile Worst Outcome Across Testing Samples			
12-Month Volatility	16.9%	16.9%	15.8%
12-Month Value at Risk (10% significance)	-24.1%	-23.9%	-22.9%
12-Month Value at Risk (5% significance)	-28.5%	-28.3%	-25.2%
Worst 12-Month Return	-32.7%	-32.5%	-29.3%

Asset allocation using mean-variance analysis

	Long Only			Long-Short		
Average Optimal Weights	Ignoring Errors	Bayesian Shrinkage	Stability Adjusted	Ignoring Errors	Bayesian Shrinkage	Stability Adjusted
U.S. Equities	24.7%	21.2%	32.5%	24.6%	21.2%	33.4%
Foreign Developed Market Equities	19.4%	18.8%	18.0%	19.3%	18.8%	20.2%
Emerging Market Equities	14.1%	16.7%	10.9%	13.8%	16.7%	13.0%
Treasury Bonds	22.4%	13.7%	24.6%	30.7%	13.7%	56.7%
U.S. Corporate Bonds	13.9%	14.8%	12.3%	12.2%	14.8%	-13.1%
Commodities	3.6%	5.9%	1.3%	4.0%	5.9%	-5.4%
Cash Equivalents	1.9%	9.0%	0.3%	-4.7%	9.0%	-4.8%
10 Percentile Worst Outcome Across Testing Samples						
12-Month Volatility	17.1%	17.1%	15.8%	18.7%	17.1%	16.1%
12-Month Value at Risk (10% significance)	-25.3%	-24.9%	-23.5%	-29.9%	-24.9%	-21.8%
12-Month Value at Risk (5% significance)	-29.7%	-28.9%	-25.9%	-34.6%	-28.9%	-23.9%
Worst 12-Month Return	-33.0%	-32.6%	-30.1%	-35.7%	-32.6%	-28.2%

Stability-adjusted benefit compared to error-blind optimization

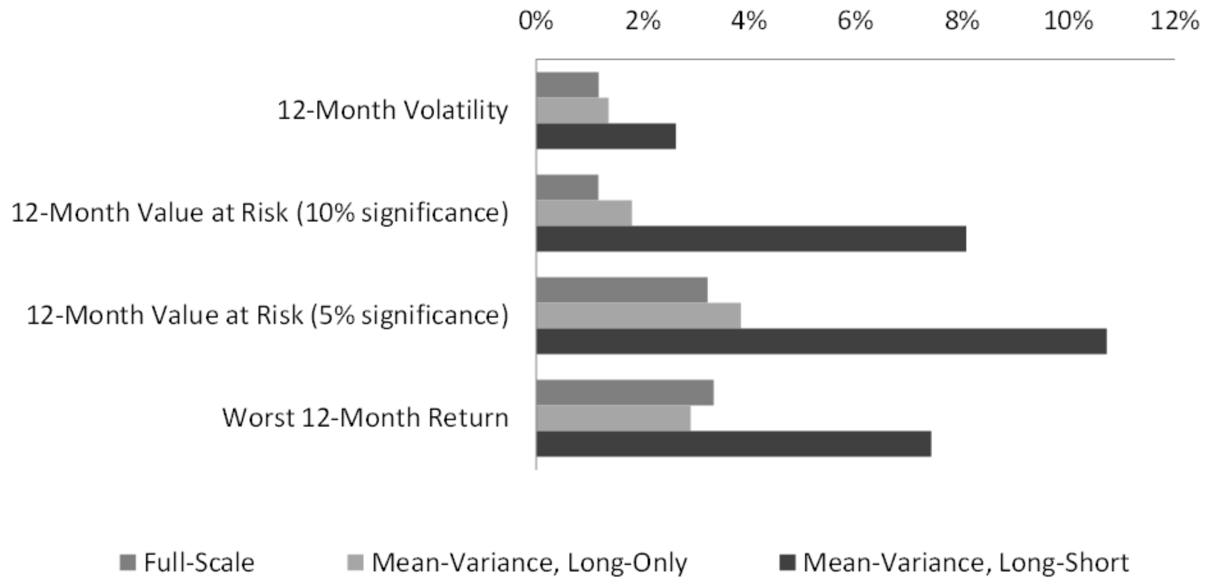

Summary

- When investors estimate covariances from historical returns they face four types of estimation error: small-sample error, independent-sample error, mapping error, and interval error.
- Small-sample error arises because the investor's investment horizon is typically shorter than the historical sample from which covariances are estimated.
- Independent-sample error arises because the investor's investment horizon is independent of history.
- Interval error arises because investors estimate covariances from higher frequency returns than the return frequency they care about. If returns have non-zero auto-correlations, standard deviation does not scale with the square root of time. If returns have non-zero auto-correlations or non-zero lagged cross correlations, correlation is not invariant to the return interval used to measure it.
- Common approaches for controlling estimation error, such as Bayesian shrinkage and resampling, make portfolios less sensitive to estimation error.
- A new approach, called stability-adjusted optimization, assumes that some covariances are reliably more stable than other covariances. It delivers portfolios that rely more on relatively stable covariances and less on relatively unstable covariances.