The Impact of Estimation Error on Latent Factor Model Forecasts of Portfolio Risk

(**forthcoming in** *The Journal of Portfolio Management*)

Stephen W. Bianchi, Lisa R. Goldberg, Allan Rosenberg

Stephen W. Bianchi is a Lecturer in Economics at the University of California, Berkeley, and a consultant at Berkeley Associates, LLC, in Berkeley, CA. `swbianchi@berkeley.edu`

Lisa R. Goldberg is an Adjunct Professor of Economics and Statistics at the University of California, Berkeley, and a founding partner of Berkeley Associates, LLC, in Berkeley, CA. `lrg@berkeley.edu`

Allan Rosenberg is a Senior Quantitative Analyst at State Street Global Exchange in San Francisco, CA. `arosenberg@statestreetgx.com`

Abstract

In this article, the authors measure the impact of estimation error on latent factor model forecasts of portfolio risk and factor exposures. In markets simulated with a Gaussian return generating process, the authors measure errors in forecasts for equally weighted and long-only minimum variance portfolios constructed from a universe of 500 securities. They find that an estimation period of 250 days may be adequate to accurately forecast risk and factor exposures for an equally weighted portfolio. In contrast, the risk of a long-only minimum variance portfolio is substantially under-forecast even with an estimation period of 1000 days. This underscores the importance of testing risk models on optimized portfolios.

Quantitative investors rely on factor models of portfolio risk to make decisions. Factor models generate forecasts of volatility, expected tail loss, and other measures that are used for risk management, regulatory reporting and portfolio construction. Factor models also forecast exposures of portfolios to risk factors, which are used to construct hedges and tilts and to control unintended bets.

The importance of factor models to quantitative investors cannot be overstated, but these models are not foolproof. Factor models are estimated from data, so factor model forecasts are inevitably affected by estimation error. How much data are required to effectively control estimation error? We address this question by measuring the impact of estimation error on investment decisions in simulated markets, where security returns follow a known return generating process. This allows us to assess minimal data requirements for a model that is accurate enough to be useful.

In the experiments described below, we evaluate latent factor models, which make minimal assumptions about market structure and allow data to "speak." We concentrate on portfolios whose returns depend linearly on returns to the risk factors, and we focus on volatility forecasts and factor exposures. While this is a narrow program, it provides salient, clean assessments of the impact of model estimation error on metrics that guide investment decisions. The restriction to linear portfolios allows us to express model forecasts as simple functions of model parameters and to construct optimal portfolios with simple programs.[1]

Return Generating Process

A standard assumption in financial economics is that returns to securities are driven by a relatively small number of risk factors, plus security specific returns. In general, the relationship between security returns and factor returns is non-linear. However, to identify the factors and measure estimation error, we rely on a universe of securities whose returns depend linearly on returns to factors. The security return generating process is

$$R = \psi Y + \epsilon, \tag{1}$$

where R is a $T \times N$ matrix of security returns, ψ is a $T \times K$ matrix of simulation factor returns, Y is a $K \times N$ matrix of simulation factor exposures and ϵ is a is a $T \times N$ matrix of security specific returns. The simulation factor returns ψ and specific returns ϵ have mean zero and are uncorrelated over time. The specific returns ϵ are pairwise uncorrelated over time. Under these assumptions, the security covariance matrix can be expressed as

$$\Sigma = Y^\top F Y + \Delta, \tag{2}$$

where F is a $K \times K$ factor covariance matrix and Δ is an $N \times N$ diagonal specific risk matrix.

Simulation Factors and Latent Factors

Formula (1) is useful for generating simulation data since the factors can be calibrated to known features that drive markets. For example in simulating an equity market, we may want to make the first factor "market-like," and then add factors that represent industries, countries, currencies, or investment styles. The returns on these factors may be correlated. In contrast, the output of our estimation process is a set of *latent factors*, whose security exposure vectors are orthogonal, and whose covariance matrix is the identity.

[1] Many risk management systems use simulated return distributions to forecast and attribute risk. Errors stemming from a simulated return distribution adversely affect risk forecasts even if the risk model driving the simulation of the return distribution is perfectly specified. In practice, risk models are never perfectly specified, and it is one aspect of imperfect model specification that is the subject of this note. Specifically, we are concerned with errors in forecasts of risk and factor exposures that arise from model estimation error. In order to isolate the impact of model estimation error on risk forecasts, we restrict our attention to a single risk measure, volatility (or equivalently, its square, variance) and we consider only portfolios whose returns are linear functions of factor returns. When portfolio return depends linearly on factor return, volatility forecasts and factor exposures are simple functions of model parameters so simulated return distributions are not required in order to forecast volatility. Consequently, our measures of estimation error are not confounded by errors arising from a simulated return distribution.

The estimated latent risk factors are computationally convenient, but they may not have a ready interpretation and they cannot be compared directly to the simulation factors specified in formula (1). This is not a serious issue, however, but rather a matter of presentation. It is always possible to transform the simulation returns ψ to latent returns ϕ whose covariance matrix is the identity. This transformation maps the simulation factor exposures Y to latent factor exposures X that are pairwise orthogonal.

Specifically, there is an invertible $K \times K$ matrix M for which the covariance matrix of $\phi = \psi M^{-1}$ is the identity, and the rows of

$$M^{-1}X = Y. \tag{3}$$

are pairwise orthogonal. Then

$$\begin{aligned} R &= \psi M^{-1} M Y + \epsilon \\ &= \phi X + \epsilon. \end{aligned} \tag{4}$$

The matrix M is unique provided that a technical condition is satisfied. Specifically, if $F^{1/2}$ is the symmetric square root of F, the eigenvalues of $F^{1/2} Y Y^\top F^{1/2}$ must be distinct. The uniqueness allows direct comparison of estimated latent factors to the true latent factors, thereby facilitating the measurement of estimation error. A precise expression for M is in Appendix B.

Formula (3) provides a two-way translation between the more intuitive presentation of risk in terms of simulation factor exposures Y and the more computationally convenient presentation in terms of latent factor exposures X. In the latent factor basis, formula (2) simplifies to

$$\Sigma = X^\top X + \Delta. \tag{5}$$

Latent Factor Model Estimation

Many latent factor model estimation algorithms have their roots in principal component analysis (PCA). We estimate factors with a modified PCA, principal factor analysis (PFA), which iterates PCA weighted by the inverse of an estimated specific variance matrix. The process continues until the factors and specific variance estimates stabilize. Unlike basic PCA, PFA is compatible with the assumption that specific risk can vary across securities.[2]

Risk Forecasting and Portfolio Construction

Quantitative investment management relies heavily on two measures of risk. Volatility (Vol) is the standard deviation of portfolio return, and is used for portfolio construction. Expected tail loss (ETL) is the average loss given that a specified value at risk [3] is

[2]Details are in Stroyny and Rowe [2002].

[3]Value at risk is a quantile of a portfolio return distribution.

breached; its applications include risk management and regulatory reporting.[4] Quantitative investment also relies heavily on factor exposures, which are used to make controlled bets, to avoid unintended bets, and to hedge.

As we discuss in Appendix A, if the dependence of security returns on factor returns is non-linear, simulation is generally required to forecast volatility. And if security returns are not Gaussian, simulation is generally required to forecast expected tail loss even if the relationship between security returns and factor returns is linear. In the special case where portfolio returns depend linearly on returns to factors, volatility estimates and factor exposures can be expressed in terms of simple formulas that depend on the factor model parameters. This focus allows a clean assessment of error arising from estimation of the factors and specific variances since it is not corrupted by estimation error arising from simulated factor return distributions.

Here, we outline simple metrics for the impact of estimation error on forecasts of risk and factor exposures. Throughout, we specify a portfolio by its vector of weights w.

Forecasting the Volatility and Variance of a Linear Portfolio

For transparent statistical interpretation and consistency with other studies (such as Bender, Lee, Stefek, and Yao [2009]), we work with variance (the square of volatility) instead of volatility. If the returns to the securities in a portfolio with weights w are linear in the sense that they follow formula (1), the true variance of that portfolio is given by

$$\begin{aligned}
\mathrm{Var}(w) &= (\mathrm{Vol}(w))^2 \\
&= w^\top \Sigma w && (6)\\
&= w^\top \left(Y^\top F Y + \Delta\right) w && (7)\\
&= w^\top \left(X^\top X + \Delta\right) w && (8)\\
&= w^\top \left(X^\top X\right) w + w^\top \Delta w \\
&= \mathrm{CFV}(w) + \mathrm{SV}(w), && (9)
\end{aligned}$$

where formula (8) follows from formula (7) since formulas (2) and (5) give expressions for Σ, and $\mathrm{CFV}(w) = w^\top \left(X^\top X\right) w = |Xw|^2$ is the common factor variance of w and $\mathrm{SV}(w) = w^\top \Delta w$ is the specific variance of w. If we assume the factor returns ϕ and specific returns ϵ in formula (1) are jointly Gaussian, expected tail loss at quantile α is gotten by scaling volatility with a known constant,[5] $\mathrm{ETL}_\alpha = C(\alpha)\mathrm{Vol}$. It follows that expected tail loss can be calculated directly from volatility, so a single set of experiments can be used to assess estimation error in the two risk measures.

[4] Historically, value at risk, and not expected tail loss, has played a central role in risk management and regulation. However, expected tail loss is preferred, on both mathematical and economic grounds, and it is becoming the new standard for extreme risk measurement. More information is in Acerbi and Tasche [2002] and Acerbi and Székely [2014].

[5] Formulas for the relationship between volatility and ETL at different quantiles can be found in Goldberg, Miller, and Weinstein [2008].

Forecasting the Factor Exposures of a Linear Portfolio

Using the transformation in formula (3), the K-vectors of exposures of a portfolio with weights w to simulation factors Y and latent factors X are related by

$$Xw = MYw, \tag{10}$$

so

$$Yw = M^{-1}Xw. \tag{11}$$

Constructing a Long Only Minimum Variance Portfolio

Modern portfolio theory began when Markowitz [1952] framed portfolio construction as a tradeoff between mean return and variance (volatility squared).[6] More than sixty years later, mean-variance optimization is still a standard tool for portfolio construction, but we are just beginning to understand the impact of model estimation error on risk forecasts on optimized portfolios.

While there is a wide spectrum of portfolios that are constructed with mean-variance optimization, we focus on minimum variance equity portfolios in the US for several reasons. First minimum variance portfolios are popular with investors.[7] Second, the construction of minimum variance portfolios does not require estimates of mean return, so analysis based on minimum variance is not corrupted by estimation error in mean return. Finally, minimum variance portfolios are extremely sensitive to model estimation error, as discussed in Menchero, Wang, and Orr [2008]. In the analysis below, we consider a long-only minimum variance portfolio, which is obtained by solving the following optimization problem:

$$\min_{w} \; w^T \Sigma w \tag{12}$$
$$\text{subject to } w^\top 1 = 1, w \geq 0.$$

Using Simulation to Gauge the Impact of Estimation Error

To gauge the impact on estimation error on a latent factor model, we repeat the following steps over many simulations:

- Use a known process to generate a data set of security returns over a fixed estimation period.
- Use the data set of security returns to estimate a latent factor model.

[6]More recently, methods to use expected tail loss forecasts in portfolio construction have been developed. Further information is in Rockafellar and Uryasev [2000], Bertsimas, Lauprete, and Samarov [2004], and Goldberg, Hayes, and Mahmoud [2014].

[7]Minimum variance portfolios have outperformed the market on the basis of absolute return and on a volatility-adjusted basis over the past four decades, as discussed in Goldberg, Leshem, and Geddes [2014].

- Use the estimated latent factor model to forecast portfolio risk and factor exposures of unoptimized and optimized portfolios.
- Compute statistics that measure the difference between risk and factor exposure forecasts from the estimated model with the true values, which can be calculated using the parameters of the return generating process.

The distributions of these statistics over many simulations provide metrics for the impact of estimation error on forecasts of risk and factor exposures.

In the studies described below, we impart the estimation process with knowledge of the number of factors, K, used to generate the data set. In practice, however, the true numbers are not known. Since estimation error can only increase if the algorithm is given the extra burden of determining the number of factors, the results shown here must be regarded as optimistic.

Sample Simulation of an Estimation Universe

Our analysis relies on a parametric simulation based on a known return generating process. This facilitates a precise comparison of estimated and true quantities.[8] A prescription for simulating a stationary, Gaussian model that follows formula (1) is given here.

- Choose the number of factors K, the number of securities N, and the number of days T.
- Specify:
  - A $K \times N$ factor exposure matrix Y,
  - A diagonal $K \times K$ factor covariance matrix F, and
  - A diagonal $N \times N$ specific covariance matrix Δ.
- Draw T $(K + N)$ vectors (ψ_t, ϵ_t) with mean 0 and covariance matrix equal to the diagonal matrix generated by F and Δ.[9]
- Compute T simulated returns
$$r_t = \psi_t Y + \epsilon_t.$$
- Use the T simulated returns to estimate a sample covariance matrix $\hat{S}$.
- Use a latent factor methodology to generate

[8]There is inevitably a discrepancy between the data generating process that drives the simulation and the process that drives empirical data. Complementary to a parametric simulation is an empirical simulation, also known as an empirical bootstrap. In this case, the empirical estimate plays the role of the truth, and the bootstrap leads to confidence intervals around estimated quantities. Here, there is an implicit assumption that the daily observations are independent and identically distributed.

[9]It is a common assumption that the vectors (ψ_t, ϵ_t) are independent, identical and jointly Gaussian. However, this assumption is at odds with empirical findings, and that should be considered when interpreting the results of this experiment.

  - estimates $\hat{\phi}$ of factor returns whose covariance matrix is equal to the identity,
  - estimates $\hat{X}$ of (pairwise orthogonal) factor exposures and
  - an estimated specific risk matrix $\hat{\Delta}$.
- Find M as outlined in Appendix B.
- Set $\hat{Y} = M^{-1}\hat{X}$.
- Assess the impact of using the estimated factor exposures $\hat{Y}$ or $\hat{X}$ and the estimated covariance matrix $\hat{\Sigma} = \hat{X}^\top \hat{X} + \hat{\Delta}$ on forecast of risk and factor exposures.

Even in this simple setting, it is possible to include some empirically observed features of financial markets, such as a dominant first factor to which most securities in the estimation universe are positively exposed. In practice, it is desirable to consider more realistic simulations that take account of market regimes, memory, and industry, currency and country factors.

Estimates of Latent Factor Model Parameters and Applications

The goal of latent factor estimation is to recover the unobservable components of a factor model from observable data using purely statistical methods.[10] In our framework, the elements of the latent factor model are estimated from data simulated with return generating process given in formula (1).

Below, we use PFA to find $\hat{X}$, an estimate of latent factor exposures X, and $\hat{\Delta}$, an estimate of the specific variance matrix Δ. Substituting these estimated quantities for the true quantities in formula (5), we can generate portfolio volatility and variance forecasts, as well the breakdown of forecast variance into common factor and specific components. Substituting the estimate of latent factor exposures X into formula (3) gives $\hat{Y}$, an estimate of the simulation factor exposures Y.

Estimation Error Metrics

As above, we specify a portfolio by its vector of weights w, and we assume the returns to the securities depend linearly on the returns to the factors as in formula (1). Since we do not know the true factor exposures X, the true specific covariance matrix Δ or the true security covariance matrix Σ, we base our risk forecasts, factor exposure forecasts and portfolio construction routines on the latent factor estimates $\hat{X}$, $\hat{\Delta}$ and $\hat{\Sigma}$, as described earlier in this article.

Errors in Volatility and Variance Forecasts

Substituting the latent factor estimate $\hat{\Sigma}$ of the covariance matrix Σ into formula (6) gives estimated variance $\widehat{\mathrm{Var}}(w) = w^\top \hat{\Sigma} w$, while the true variance is $\mathrm{Var}(w) = w^\top \Sigma w$.

[10]For a multi-asset class risk model, the observables may include security prices, yield curves, capitalization weights, trading volume, and previous model estimates.

The quotient of estimated variance by true variance is the Variance Forecasting Ratio (VFR),[11]

$$\mathrm{VFR}(w) = \frac{\widehat{\mathrm{Var}}(w)}{\mathrm{Var}(w)}. \tag{13}$$

Formula (9) expresses the risk of portfolio w as a sum of a common factor and specific components. Substituting the latent factor estimates $\hat{\mathrm{CFV}}$ and $\hat{\mathrm{SV}}$ of CFV and SV lets us define the Factor Variance Forecasting Ratio

$$\mathrm{FVFR}(w) = \frac{\widehat{\mathrm{CFV}}(w)}{\mathrm{CFV}(w)} \tag{14}$$

and the Specific Variance Forecasting Ratio

$$\mathrm{SVFR}(w) = \frac{\widehat{\mathrm{SV}}(w)}{\mathrm{SV}(w)} \tag{15}$$

Values of VFR, FVFR and SVFR that are closer to 1 indicate greater accuracy.

Errors in Factor Exposures

The Latent Factor Exposure Error (LFEE) of a portfolio w is given by

$$\mathrm{LFEE}(w) = \left(\hat{X} - X\right) w, \tag{16}$$

where LFEE(w) is a K-vector whose kth entry is the error in the exposure of portfolio w to latent factor k. [12]

Since the latent factor exposures may not be easy to interpret, we use the transformation M to transform the true and estimated latent factor exposures to the intuitive setting used to simulate the data. The Simulation Factor Exposure Error (SFEE) of a portfolio w is given by

$$\mathrm{SFEE}(w) = \left(\hat{Y} - Y\right) w \tag{17}$$

$$= M^{-1}\mathrm{LFEE}(w). \tag{18}$$

SFEE(w) is a K-vector whose kth entry is the error in the exposure of portfolio w to simulation factor k. We are interested in the L_2 lengths as well as individual components of these error vectors. Values of LFEE(w) and SFEE(w) that are closer to 0 indicate greater accuracy. In the experiments outlined below, we report SFEE(w) since it is easier to interpret.

[11] A similar metric is used in Bender et al. [2009].

[12] PFA simultaneously estimates a set of factor returns (columns of $\hat{\phi}$) and a set of factor exposures (rows of $\hat{X}$). To the extent that the columns of $\hat{\phi}$ span the same space as the columns of ϕ, formula (16) is a reasonable estimate of the error in the latent factor exposures. The columns of $\hat{\phi}$ approach the columns of ϕ as T goes to infinity. For a finite sample, however, estimation error can cause the columns of $\hat{\phi}$ to span a different space than the columns of ϕ.

Empirical Experiments

In each of the four experiments presented below, we simulate 1000 data sets, each composed of time series of returns to 500 securities. We assume these returns follow the linear process in formula (1) and we apply PFA to each simulated data set to estimate a latent factor model. Next, we use the latent factor model to forecast risk and factor exposures for a particular portfolio. Since the data sets are simulated, we know the true risk and the true factor exposures of equally weighted and minimum variance portfolios. So we can compare true risk and factor exposures to their counterparts generated by models estimated from data set, We measure errors using VFR, FVFR, SVFR and SFEE, which are averaged over simulations.

The experiments are designed to measure the extent to which increasing the observations used to estimate a latent factor model lowers estimation error. We consider both an equally weighted portfolio, which is constructed without reference to the estimated factor model, and a minimum variance portfolio, whose construction relies on the estimated factor model. A summary of our four experiments is in Exhibit 1.

Portfolio	**Number of Daily Observations**
Equally Weighted	250
Equally Weighted	1000
Minimum Variance	250
Minimum Variance	1000

Exhibit 1: Parameters of the four experiments used to gauge the impact of model estimation error on forecasts of risk and factor exposures.

Model Calibration

Following the data simulation prescription outlined above, we assume $N = 500$ securities, $K = 2$ factors, and either $T = 250$ or $T = 1000$ days. The simulation factor returns ψ are normal and uncorrelated, and they have mean 0. Further,

- Factor 1 is market-like, meaning that most securities have positive exposure and the factor has an annualized volatility of 16%.

- Factor 2 is long/short with an annualized volatility of 8%.

Specifically, factor exposures are drawn from a normal distribution with mean $(1, 0)^\top$ and covariance matrix

$$\begin{pmatrix} 0.25 & 0 \\ 0 & 0.75 \end{pmatrix}.$$

Note that factor exposures (unlike factor returns) are not random in our model. Here, we are using the normal distribution as a convenient means of constructing exposures of securities to factors.

The specific variance matrix Δ is diagonal, and annualized specific volatilities are drawn from a uniform distribution on $[32\%, 64\%]$.

Equally Weighted Portfolio Results

Exhibit 2 shows the impact of estimation error on forecasts of variance and factor exposures for the equally weighted portfolio. Panel A of Exhibit 2 shows the VFR, the FVFR and the SVFR for 1000 latent factor models estimated from simulated data. We consider both $T = 250$ and $T = 1000$ daily observations, and the results are materially similar in the two cases. VFR and FVFR are close to 1, while SVFR indicates that specific risk is overforecast by roughly 12%. Additional observations do not improve results.

Panel B of Exhibit 2 shows the SFEE for the equally weighted portfolio. Factor exposures are underforecast, and the expansion of the dataset from $T = 250$ to $T = 1000$ diminishes the impact of estimation error. The dominant "market" factor exposure is underforecast by roughly 1.5% of the exposure, and the second long/short factor exposure is underforecast by roughly 16.0% when $T = 250$. When we increase the size of the data set to $T = 1000$ daily observations, the factors were underforecast by 0.9% and 11.7%.

Tentative conclusions about latent model forecasts of risk and factor exposures for the equally weighted portfolio are:

- At $T = 250$, latent factor factor variance forecasts are accurate, while specific variance forecasts are somewhat high.
- Additional observations do not improve specific variance forecasts. There may be some systematic bias in the estimation method.
- The latent factor model underforecasts factor exposures, but an increase in the number of observations mitigates the problem.

Panel A:

Equal Weighted Portfolio (500 Assets)	Mean	Std/Mean	True Vol
T = 250			
Variance Forecasting Ratio (VFR)	1.0021	0.0831	0.1622
Factor Variance Forecasting Ratio (FVFR)	0.9998	0.0847	0.1607
Specific Variance Forecasting Ratio (SVFR)	1.1198	0.0085	0.0220
T = 1000			
Variance Forecasting Ratio (VFR)	0.9992	0.0449	0.1622
Factor Variance Forecasting Ratio (FVFR)	0.9969	0.0457	0.1607
Specific Variance Forecasting Ratio (SVFR)	1.1220	0.0047	0.0220

Panel B:

Equal Weighted Portfolio (500 Assets)	Mean	Std/Mean	True Exp
T = 250			
Factor 1 Exposure Error (SFEE)	-0.0098	0.0438	1.0131
Factor 2 Exposure Error (SFEE)	-0.0019	0.0520	0.0195
T = 1000			
Factor 1 Exposure Error (SFEE)	-0.0106	0.0217	1.0131
Factor 2 Exposure Error (SFEE)	-0.0008	0.0625	0.0195

Exhibit 2: Forecasting errors for an equally weighted portfolio. Panel A displays errors in risk forecasts for models estimated from $T = 250$ and $T = 1000$ daily observations. The Mean column shows the average of VFR(w), FVFR(w) and SVFR(w) over the 1000 simulated data sets, and the Std/Mean column shows the absolute value of the standard deviation divided by the mean (larger values indicate more uncertainty in the estimate of the mean). The True Vol column shows the true (annualized) volatility of the equally weighted portfolio, as well as the true common factor volatility and true specific volatility. Panel B displays errors in factor exposures. The Mean column shows the average of SFEE(w) over the 1000 simulated data sets, and the Std/Mean column shows the absolute value of the standard deviation divided by the mean. The True Exp column shows the true factor exposures of the equally weighted portfolio.

Minimum Variance Portfolio Results

Exhibit 3 shows the impact of estimation error on variance and factor exposures for the minimum variance portfolio constructed by solving the optimization problem given by formula (12). There are material differences between the results for the minimum variance portfolio, which is constructed with an estimated risk model, and the the results for the equally weighted portfolio, whose weights do not depend on the estimated risk model.

When we forecast with a latent factor model estimated from 250 observations, we find that VFR ≈ 0.61, which means that the variance forecast of the minimum variance portfolio is roughly 61% of the true variance on average. An increase to 1000 observations raises VFR to 88%. When we disaggregate, we find that specific variance forecasts are accurate on average, but common factor variance is severely underforecast. For $T = 250$, forecast common factor variance is 49% of true common factor variance on average, but that increases to 82% when the model is estimated from 1000 observations.

Panel B of Exhibit 2 shows the Simulation Factor Exposure Errors (SFEE) for the minimum variance portfolio. Factor exposures are also consistently and materially underforecast by our latent factor model. The dominant "market" factor exposure is underforecast by roughly 31%. and the second long/short factor exposure is underforecast by roughly 38% when $T = 250$. When we increase the size of the data set to $T = 1000$ daily observations, the factor exposures were underforecast by 10% and 12%.

Tentative conclusions about latent model forecasts of risk and factor exposures for the minimum variance portfolio are:

- When $T = 250$ observations are used to estimate the latent factor model underforecast variance by 40% on average. When the number of observations is increased to $T = 1000$ the model underforecast variance by 12%.

- Risk underforecasts are attributable to common factors.

- Factor exposures are underforecast by 30-40% when $T = 250$. When the number of observations was increased to $T = 1000$, factor exposure were substantially diminished.

Panel A:

Minimum Variance Portfolio (500 Assets)	Mean	Std/Mean	**E**[True Vol]
T = 250			
Variance Forecasting Ratio (VFR)	0.6166	0.1021	0.1182
Factor Variance Forecasting Ratio (FVFR)	0.4914	0.1475	0.1027
Specific Variance Forecasting Ratio (SVFR)	1.0091	0.0627	0.0585
T = 1000			
Variance Forecasting Ratio (VFR)	0.8803	0.0465	0.1097
Factor Variance Forecasting Ratio (FVFR)	0.8198	0.0660	0.0936
Specific Variance Forecasting Ratio (SVFR)	1.0430	0.0097	0.0572

Panel B:

Minimum Variance Portfolio (500 Assets)	Mean	Std/Mean	**E**[True Exp]
T = 250			
Factor 1 Exposure Error (SFEE)	-0.1970	0.1815	0.6440
Factor 2 Exposure Error (SFEE)	-0.0300	1.5502	0.0736
T = 1000			
Factor 1 Exposure Error (SFEE)	-0.0591	0.2997	0.5871
Factor 2 Exposure Error (SFEE)	-0.0102	2.4199	0.0748

Exhibit 3: Forecasting Errors for a Minimum Variance Portfolio. Panel A displays errors in risk forecasts for models estimated from $T = 250$ and $T = 1000$ daily observations. The Mean column shows the average of VFR($\hat{w}$), FVFR($\hat{w}$) and SVFR($\hat{w}$) over the 1000 simulated data sets, and the Std/Mean column shows the absolute value of the standard deviation divided by the mean (larger values indicate more uncertainty in the estimate of the mean). The E[True Vol] column shows the average true (annualized) volatility of the equally weighted portfolio, as well as the average true common factor volatility and average true specific volatility. Panel B displays errors in factor exposures. The Mean column shows the SFEE($\hat{w}$) averaged over the 1000 simulated data sets, and the Std/Mean column shows the absolute value of the standard deviation divided by the mean. The E[True Exp] column shows the average true factor exposures of the minimum variance portfolio.

Conclusions

Factor models of portfolio risk are estimated from data, so factor model forecasts are affected by estimation error. In this note, we develop methods to gauge the impact of estimation error on forecasts of portfolio variance and on exposures of portfolios to factors. Our method relies on three measures of error in risk forecasts: the Variance Forecasting Ratio, the Factor Variance Forecasting Ratio, and the Specific Variance Forecasting Ratio, as well as a measure of error in factor exposures, the Simulated Factor Exposure Error. We apply these measures to data sets composed of 500 security returns simulated from a Gaussian two-factor model. In each simulated data set, we use principal factor analysis to estimate the factor model, and then compare true to estimated forecasts for an equally weighted portfolio and a long-only minimum variance portfolio.

The results for the two portfolios are qualitatively different. Forecasts of variance and factor exposures are reasonably accurate for the equally weighted portfolio. In addition, common factor risk is estimated more accurately than specific risk. For the long only minimum variance portfolio, however, both the common factor component of risk and the factor exposures are consistently underforecast.

The difference in results for the two portfolios stems from the fact that the estimated risk model is used to construct a long-only minimum variance portfolio. In contrast, an equally weighted portfolio does not require a risk model. Our study is consistent with results in Marčenko and Pastur [1967] and El Karoui [2013], and it underscores the importance of testing factor risk models on optimized portfolios, where model weaknesses are most severe.

A Estimating the Risk of Empirical Portfolios: Return Distribution Simulation

For the purpose of estimating the risk of empirical portfolios, more realistic situations occur when

- Security returns depend in a non-linear way on factor returns (in particular, when securities include derivatives or credit instruments), or
- Factor return distributions are non-Gaussian.[13]

In either case, parametric formulas for risk may be unavailable and simulation of a portfolio return distribution is a practical way to estimate both portfolio volatility and expected tail loss. The simulation of a portfolio return distribution for the purpose of estimating risk is not the same as the factor model simulations described in this article. The latter are used to gauge the impact of estimation error on model forecasts.

To estimate risk with simulation, we specify future states of the world by drawing factor returns and specific returns that follow a specified generating process.[14] In each

[13] When factor return distributions are non-Gaussian, empirical or parametric assumptions are required to proceed. Examples of this approach are described in Dubikovsky, Goldberg, Hayes, and Liu [2010] and Goldberg et al. [2013].

[14] The return generated function can be empirical or parametric.

state of the world, portfolio returns are estimated using security pricing formulas. The resulting distribution of portfolio returns is used to estimate volatility and expected tail loss with sample statistics.

While the complexity of empirical portfolios demands that portfolio return simulation be used to estimate risk in practice, this complicates the problem of assessing model error. When we estimate risk with a simulated return distribution, there are two sources of model error in the forecasts of Vol and ETL: errors in the simulation parameter estimates as well as errors introduced through the simulation of the return distribution. The latter would affect risk forecasts even if we were to use the true simulation parameters. The magnitude of the error that comes from simulating a return distribution depends on the number of sample paths, and it can be mitigated by using enough paths so that error due to portfolio simulation is small compared to estimation error in the covariance matrix. Additional corruption arises from the imperfect nature of security pricing formulas, and that can be mitigated (to some extent) by calibrating pricing formulas to the market.[15]

In view of these difficulties, it is essential to establish baseline estimates of estimation error in volatility of portfolios whose returns depend linearly on factors. The baseline relies on formula (6). Studies including portfolios that dependent non-linearly on factors or assessment of ETL depend on nested simulations, and these studies should be taken up in a second step.

B Transforming Simulation Factors to Latent Factors

Formula (1) specifies security returns in terms of linear factors, which we can calibrate to known market factors. In order to compare the true and estimated factors, however, we need to change basis. This is done as follows. Given the return generating process

$$R = \psi Y + \epsilon$$

specified in formula (1), there is a $K \times K$ invertible matrix M so that the covariance matrix F of $\phi = \psi M^{-1}$ is the identity, and the rows of $X = MY$ are pairwise orthogonal. The matrix M is unique provided that the eigenvalues of $F^{1/2}YY^\top F^{1/2}$ are distinct.

The proof is as follows. Let $F^{1/2}$ be its symmetric square root of the $K \times K$ factor covariance matrix F. The matrix $F^{1/2}YY^\top F^{1/2}$ is symmetric, so it can be diagonalized with an orthogonal matrix O. This orthogonal matrix is unique up to order of factors if the eigenvalues of $F^{1/2}YY^\top F^{1/2}$ are distinct. Set $M = OF^{1/2}$ and check that the covariance matrix of $\phi = \psi M^{-1}$ is the identity and the rows of $X = MY$ are orthogonal.

[15]In some cases, the security pricing formulas are, themselves, simulations.

References

Acerbi, C., B. Székely. "Back-testing Expected Shortfall." *Risk*, December 2014, pp. 3-8.

Acerbi, C., D. Tasche. "On the Coherence of Expected Shortfall." *Journal of Banking and Finance*, Vol. 26, No. 7 (2002), pp. 1487-1503.

Bender, J., J. Lee, D. Stefek, J. Yao. "Forecast Risk Bias in Optimized Portfolios." MSCI Barra Research Insight, 2009.

Bertsimas, D., G.J. Lauprete, A. Samarov. "Shortfall as a Risk Measure: Properties, Optimization and Applications." *Journal of Economic Dynamics and Control*, 28 (2004), pp. 1353-1381.

Dubikovsky, V., L.R. Goldberg, M.Y. Hayes, M. Liu. "Forecasting Extreme Risk of Equity Portfolios with Findamental Factors." In A. Berd ed., *Lessons from the Financial Crisis*, RiskBooks, 2010, Chapter 10.

El Karoui, N. "On the Realized RIsk of High-Dimensional Markowitz Portfolios." *SIAM Journal of Financial Mathematics*, 4 (2013), pp. 737-783.

Goldberg, L., R. Leshem, P. Geddes. "Restoring Value to Minimum Variance." *Journal of Investment Management*, Vol. 12, No. 2 (2014), pp. 32-39.

Goldberg, L.R., M.Y. Hayes, O. Mahmoud. "Minimizing Shortfall." *Quantitative Finance*, Vol. 13, No. 10 (2013), pp. 1533-1545.

Goldberg, L.R., G. Miller, J. Weinstein. "Beyond Value at Risk: Forecasting Portfolio Loss at Multiple Horizons." *Journal of Investment Management*, Vol. 32, No. 2 (2008), pp. 1-26.

Markowitz, H. "Portfolio Selection." *Journal of Finance*, Vol. 7, No. 1 (1952), pp. 77-91.

Marčenko, A., L.A. Pastur. "Distribution of Eigenvalues in Certain Sets of Random Matrices." *Mat. Sb. (N.S.)*, Vol. 72, No. 114 (1967), pp. 507-536.

Menchero, J., J. Wang, D.J. Orr. "Improving Risk Forecasts of Optimized Portfolios." *Financial Analysts Journal*, Vol. 64, No. 2 (2008), pp. 40-49.

Rockafellar, R.T., S. Uryasev. "Optimization of Conditional Value-at-Risk." *Journal of Risk*, 2 (2000), pp. 493-517.

Stroyny, A.L., D.B. Rowe. "A Re-Examination of Some Popular Latent Factor Estimation Methods." Working Paper, Marquette University, 2002.