Reducing Your Reliance on Risk Models: Another Look at Active Share

Michael Hunstad
Director of Quantitative Research, Northern Trust Asset Management

Introduction

One common complaint concerning quantitative equity strategies is that they rely too heavily on standard risk models such as Barra or Axioma. Ostensibly this means risk metrics computed using these tools, most notably portfolio volatility and tracking error, are inaccurate if the future behavior of equities is not properly characterized by the model. This could result in exposure to considerably more risk than had been otherwise anticipated or desired.

While risk models have historically performed quite well, there have been periods when these models have failed to make reliable predictions. Although risk model errors are an unfortunate reality, in this article we show that not all portfolios are equally sensitive to these misspecifications. In particular, we demonstrate that, all else being equal, portfolios with higher active share are much more sensitive to model errors than those with lower active share. Therefore, confidence bands around risk metrics for high active share equity products are larger and, as a result, we have less faith in their accuracy.

Model Errors

Most risk models are comprised of two basic pieces: a set of factor exposures for each individual stock and a factor covariance matrix. With these two pieces in hand we can compute the expected volatility of a portfolio and/or the expected tracking error of the portfolio against a benchmark index. Expected volatility is computed as:

$$\sigma_P = \sqrt{(w'f)\Sigma(w'f)'} \quad (1)$$

Where σ_P is the volatility of the portfolio, w is a vector of portfolio weights, f is a matrix of factor exposures, and Σ is the factor covariance matrix. We can also compute tracking error as:

$$\sigma_E = \sqrt{(d'f)\Sigma(d'f)'} \quad (2)$$

Where σ_E is the portfolio tracking error and d is a vector of active weight deviations of the portfolio from a benchmark index.

Perhaps the most important risk model error is misspecification of the factor covariance matrix, such that the matrix Σ_M actually used in the model is not equal to the realized covariance matrix Σ_R. If this is the case, both portfolio volatility and tracking error estimates are biased. For example, the degree to which actual tracking error deviates from the risk model estimate is:

$$\Delta\sigma_E = \sqrt{(d'f)(\Sigma_R - \Sigma_M)(d'f)'} \quad (3)$$

For simplicity we can define the change matrix $\Sigma_C = \Sigma_R - \Sigma_M$ such that:

$$\Delta\sigma_E = \sqrt{(d'f)(\Sigma_C)(d'f)'} \quad (4)$$

And the factor deviation as $\delta = d'f$ and move the square to the left hand side such that:

$$\Delta\sigma_E^2 = \delta(\Sigma_C)\delta' \quad (5)$$

Since most risk models are designed such that factors are uncorrelated we note that Σ_R, Σ_M and hence Σ_C will have no non-zero off-diagonal elements. We can now write the above equation as simply:

$$\Delta\sigma_E^2 = \sum_{i=1}^{N} \Delta\sigma_i^2 \delta_i^2 \quad (6)$$

which is just the sum of the squared factor deviations times the change in their respective factor variances. Since $\delta = d'f$ and f remains unchanged, for any non-zero variance change $\Delta\sigma_i^2$, the squared change in tracking error $\Delta\sigma_E^2$ depends only on the square of the active weight vector d.

Equation 6 shows unambiguously that tracking error inaccuracies are magnified as the squared values of active weights are increased. This is directly equivalent to stating that tracking error changes are magnified as active share increases, since an increase in active share will *always* result in an increase in the sum of squared active weights.

Simulation

To confirm our assertion, we constructed a simple Monte Carlo simulation to study what happens to tracking error as the factor covariance matrix is perturbed. These perturbations are intended to reflect misspecifications in the covariance matrix and we will measure inaccuracies in tracking error from these misspecifications across a range of different active share levels. As we will see, tracking error inaccuracies rise polynomially with active share.

For each active share increment of 0.1 between the theoretical minimum of 0 and the maximum of 1.0 we generate 10,000 iterations. Each iteration models Σ_M as an identity matrix and Σ_R as an identity matrix with random, normally-distributed perturbations with mean of zero and standard deviations as described below. Factor exposures to three common factors (meant to represent the Fama French factors for beta, value, and size) are also simulated as random draws from a standard normal distribution.

The starting active weights are also chosen randomly, but are bound by the prescribed level of active share. These active weights and their corresponding factor exposures are then used to compute expected tracking error using Σ_M and realized tracking error using Σ_R. The difference between these two measures is saved and the standard deviation of these differences is then computed following completion of all iterations. These standard deviations define the confidence intervals around our model tracking error estimates. For a perturbation standard deviation of 0.01, the confidence intervals are shown in Exhibit 1.

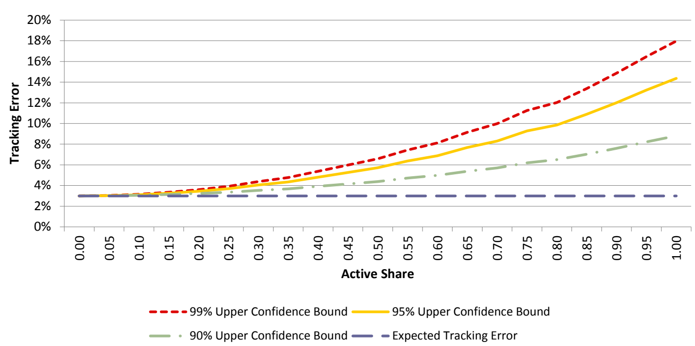

Exhibit 1 Confidence Bounds
Source: BARRA and author's calculations

Partial results for perturbation standard deviations of 0.01 and 0.02 are detailed in Exhibit 2. We base our analysis on an expected tracking error of 3% and show the 95% upper bound on realized tracking error as active share is increased. It is clear that at low levels of active share, errors in the risk model have very little impact on measured versus realized tracking error. However, as active share increases the errors are magnified such that realized tracking error could be significantly different from what was anticipated.

Are the 1% and 2% levels of perturbation realistic? The 1% value corresponds to an expected change in factor variance of about 10% and 2% to an expected change in factor variance of roughly 14% — not at all unlikely from a historical perspective. It seems that our concern about the reliability of high active share risk metrics is warranted.

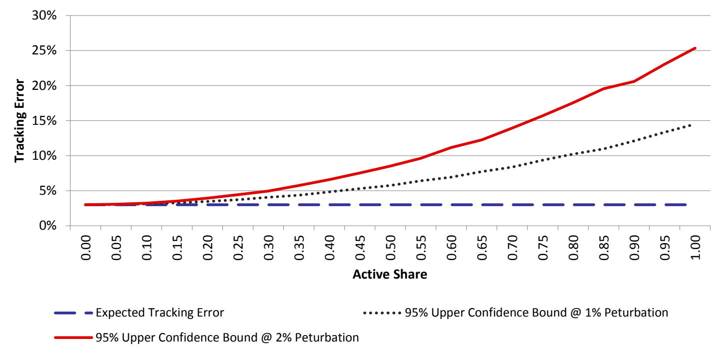

Exhibit 2 Perturbed Confidence Bounds Under
Source: BARRA and author's calculations

While we have focused on perturbations of the factor covariance matrix, we get very similar results when we perturb the individual stock factor exposures. Since the matrix f is simply a multiplier on the deviation vector d, we can clearly see how larger active shares once again produce significant biases in risk metric estimates when factor exposures are misspecified. Importantly, note that these errors are multiplicative and not additive. If both the factor covariance matrix and the factor exposures are misspecified, then the confidence interval around risk metrics is even more extreme.

Finally, we note that errors in either the factor covariance matrix or the factor loadings are directly synonymous with errors in the individual stock return covariance matrix. To illustrate this simply, note from Equation 2 that we can use the factor covariance matrix Σ and individual stock factor loadings f to recover the covariance matrix of individual stock returns we'll call Σ_{SR} where:

$$\Sigma_{SR} = f\left(\Sigma f'\right) \qquad (7)$$

Individual Factor Misspecification

Up until now, we have assumed risk model errors are equally likely for any factor. In other words, misspecification of factor loadings or factor covariances are, for example, just as probable for size as they are for value or beta.

We know, however, that this is not the case and the likelihood of errors in a specific factor increases as its return volatility becomes less stable or less consistent through time. If factor volatilities are changing, then estimates of factor covariances and loadings are probably biased.

We can get a rough gauge of factor return stability by measuring how its volatility changes through time. This measure, known as the "volatility of the volatility" or "vol. of the vol.," is simply the annual standard deviation of rolling 12-month return volatilities. The higher the "vol. of the vol." the more the factor return volatility changes over time and the less confidence we have in risk model components associated with that factor.

Exhibit 3 details the "vol. of the vol." estimates for several BARRA risk factors. These figures were calculated from January 1974 to December 2014 and show that different factors do indeed have different "vol. of vol." measures. The particularly high "vol. of vol." for momentum and volatility are strongly intuitive. Numerous studies have highlighted the temporal inconsistency of equity volatilities and demonstrated that the volatility of momentum is equally episodic.

BARRA Factor	Annualized Vol. of Vol.	Skew
Momentum	2.6%	1.97
Volatility	2.5%	1.51
Size	1.0%	.087
Earnings Yield	0.9%	1.88
Value	0.6%	1.34
Dividend Yield	0.6%	0.92
Leverage	0.5%	1.16
Earnings Variability	0.5%	0.90
Growth	0.5%	0.09

Exhibit 3 Factor Vol. of Vol. 1974 to 2014

Source: BARRA and author's calculations

These results are confirmed by examining the historical time series of factor standard deviations from the BARRA covariance matrices themselves. From Exhibit 4 it is clear that few of the factors have standard deviations that are consistent through time. As with our previous analysis, volatility and momentum show the highest degree of instability, with standard deviations that fluctuate between 30% and 100%. More importantly, it shows our perturbation assumptions in our Monte Carlo study are entirely realistic.

Another measure we can use to assess the inconsistency of factor volatility is skewness. In this case, skewness measures the relative frequency of volatility spikes within the factor returns. The higher the skewness, the more likely a factor is to have a volatility spike and, hence, the more likely the factor will have inconsistent volatility. Exhibit 4 shows that momentum and volatility are particularly prone to volatility spikes. A histogram of momentum volatilities is shown in Exhibit 5, which clearly suggests a long right hand tail (strong positive skewness) to the distribution such as that approximated by a lognormal fit.

These findings suggest that portfolios targeted as specific factors, particularly high momentum and high volatility (i.e., high beta), are more exposed to factor model misspecifications and, hence, the confidence bands around their risk metrics are particularly wide. Fundamental equity strategies that typically have a high active share along with a relatively high exposure to momentum are particularly prone to risk metric bias.

[1]In statistics this phenomenon is known as heteroskedasticity.

[2]Tests for heteroskedasticity were also conducted using the Breush-Pagan (1979), Breusch-Pagan-Koenker modification (1980) and White (1980) tests. All factors were found to be heteroskedastic with the exception of Value.

Skewness is defined as the third central moment about the mean:

$$\frac{\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^3}{\left(\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^2\right)^{3/2}}$$

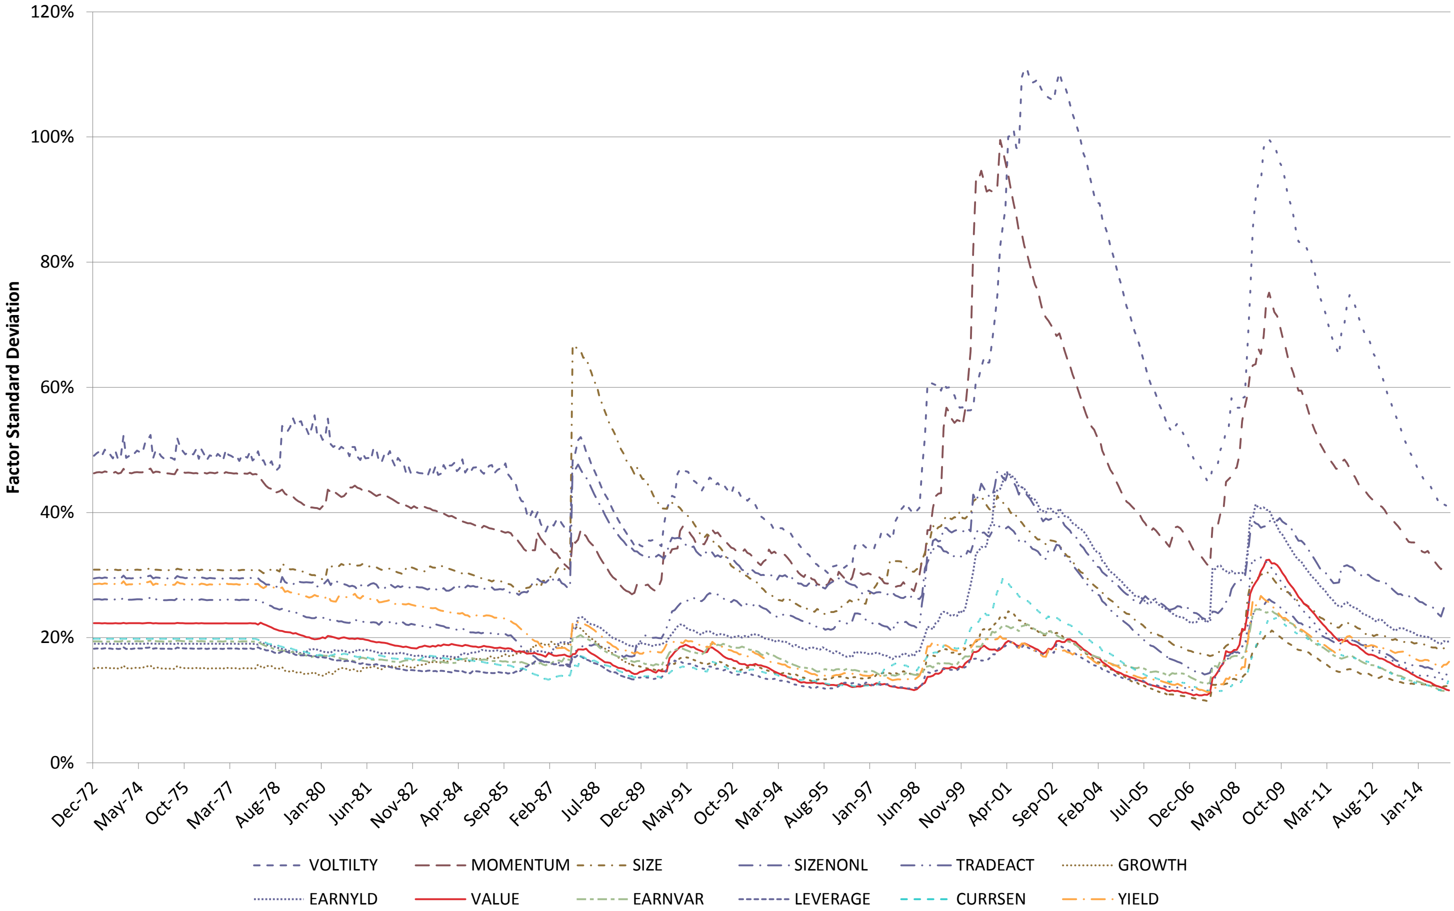

Exhibit 4 BARRA Factor Standard Division

Source: BARRA and author's calculations

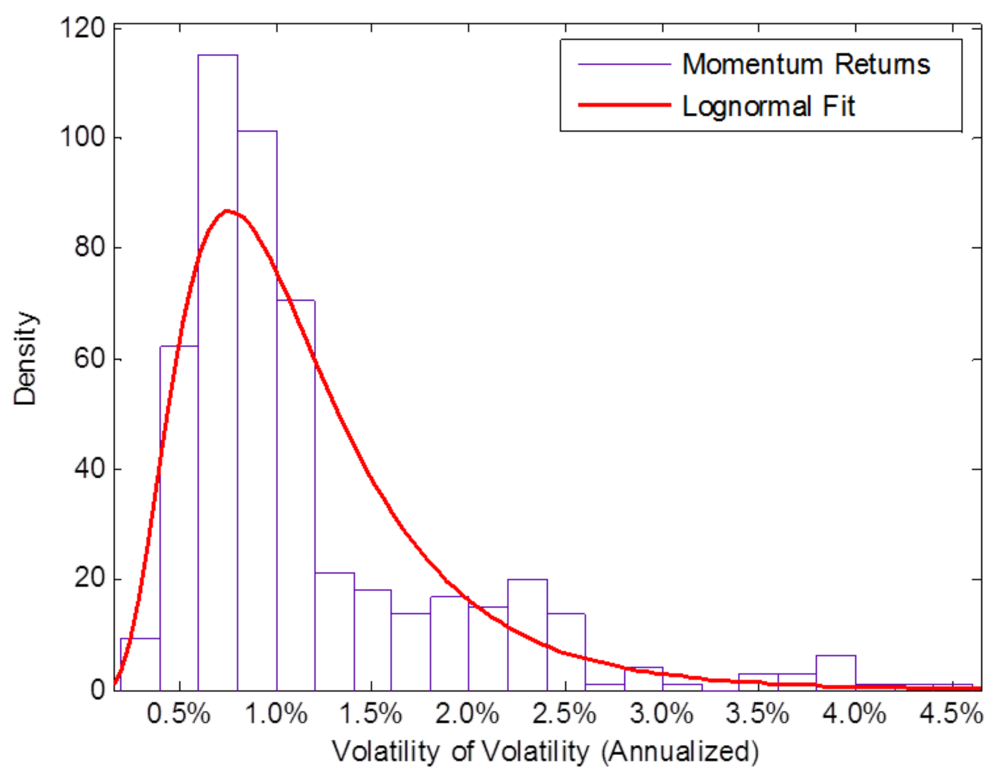

Exhibit 5 Distribution of Momentum Volatilities

Source: BARRA and author's calculations

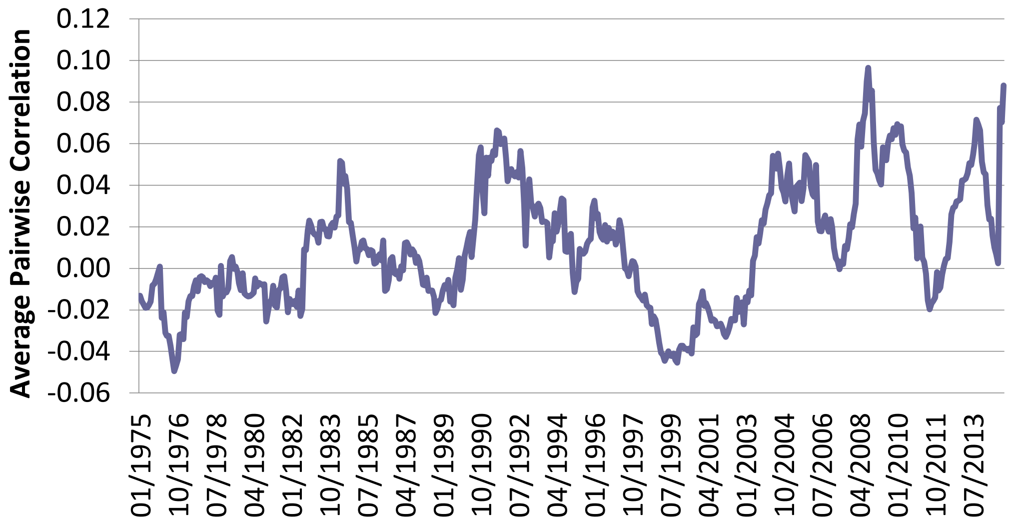

Exhibit 6 Average Pairwise BARRA Factor Return Correlation
Source: BARRA and author's calculations

Misspecification of Correlations

We mentioned previously that most risk models are designed such that factors are uncorrelated and, hence, Σ_R, Σ_M, and Σ_C will have no non-zero off-diagonal elements. While this assumption typically holds, there can be periods in which off-diagonal covariances are non-zero which can also influence tracking error inaccuracies. To see this we can rewrite Equation 6 with covariance terms as:

$$\Delta\sigma_E^2 = \sum_{i=1}^{N} \Delta\sigma_i^2 \delta_i^2 + 2 \sum_{1 \le i < j \le N} \Delta \operatorname{cov}(x_i, x_j) \delta_i^2 \delta_j^2 \qquad (8)$$

Where x_i and x_j represent the return series of individual factors i and j used to compute Σ_M. Note that, once again, for any non-zero covariance change $\Delta \operatorname{cov}(x_i, x_j)$ the squared change in tracking error $\Delta\sigma_E^2$ depends only on the square of the active weight deviation d_i (or d_j) since $\delta_i = d_i f_i$ and f_i remains unchanged.

Exhibit 6 shows the average rolling 24 month pair-wise correlation among the 13 Barra factor return series. Although the average correlation is approximately zero over time, there are distinct periods where correlations and hence covariances are significantly positive. For example, between July of 2007 and August of 2008, the average pairwise correlation jumps from 0.00 to more than 0.10. Although this is a relatively small correlation in absolute terms, it represents a significant change in covariances that can influence tracking error estimates materially. Note that most correlation spikes occur in periods of recession where asset correlations in general tend to increase.

Like individual factor volatilities, misspecification of factor covariances causes tracking error changes to be magnified as active share increases. Although the effect is somewhat less extreme than those shown in Exhibits 1 and 2, the impact can still be meaningful.

Conclusions

The results of this paper show:

- While all portfolio risk metrics are sensitive to errors in the risk model, some portfolios are more sensitive than others.
- The sensitivity of a portfolio to risk model errors rises with active share.
- Monte Carlo simulation shows that at realistic levels of risk model error the confidence bounds on risk metrics grow dramatically with active share and, therefore, these metrics may lose credibility as active share increases.
- The likelihood of risk model error depends on portfolio factor exposure. For example, the higher the exposure to momentum and volatility factors, the larger the confidence band around portfolio risk metrics.
- To minimize reliance on risk models, one should choose an equity portfolio that meets return and risk objectives, but otherwise minimizes active share and exposure to specific factors like momentum and volatility.

References

1. Breusch, T. S.; Pagan, A. R. (1979). "A Simple Test for Heteroscedasticity and Random Coefficient Variation." *Econometrica* 47 (5): 1287-1294.JSTOR 1911963.

2. Koenker, R., (1981). "A Note on Studentizing a Test for Heteroscedasticity." *Journal of Econometrics* 17, p. 107–112.

3. White, H., (1980). "A Heteroskedasticity-Consistent Covariance Matrix and a Direct Test for Heteroskedasticity." *Econometrica* 48, p. 817–838.

Author's Bio

Michael Hunstad, Ph.D.

Michael Hunstad, Ph.D., is the Head of Quantitative Research within the Global Equity Group of Northern Trust Asset Management. Prior to joining Northern Trust, Dr. Hunstad was head of research at Breakwater Capital, a proprietary trading firm and hedge fund. Previously, he was head of quantitative asset allocation at Allstate Investments, LLC, and a quantitative analyst with a long-short equity hedge fund. Dr. Hunstad holds a Ph.D. in applied mathematics from the Illinois Institute of Technology as well as an MBA in finance and an MA in econometrics.