MIT Sloan School of Management

MIT Sloan School Working Paper 5128-15

Facts About Factors

Paula Cocoma, Megan Czasonis, Mark Kritzman, David Turkington

April 6, 2015

FACTS ABOUT FACTORS

This Version: April 6, 2015

Paula Cocoma
pcocoma@statestreet.com

Megan Czasonis
mczasonis@statestreet.com

Mark Kritzman
kritzman@mit.edu

David Turkington
dturkington@statestreet.com

Abstract

It has become fashionable to allocate portfolios to factors rather than to assets. The often stated motivation for this approach is that factors are less correlated with each other than assets; therefore, factors afford greater opportunity for diversification. This argument is specious, of course, because ultimately the portfolio must be invested in assets. It is, therefore, impossible to produce a better in-sample portfolio by describing the portfolio as a set of factors than assets. There are several potentially legitimate arguments, though, for favoring factor stratification over asset stratification. It could be that factors are easier to forecast than assets, because investors are better able to relate current information to future factor behavior than to future asset behavior. Unfortunately, we have no way of testing this conjecture generically. But there are several testable conjectures. Perhaps risk estimated from high-frequency returns predicts risk over longer horizons more reliably for factors than for assets. Or the statistical properties of large samples may predict the statistical properties of small samples more reliably for factors than for assets. Or, for the same sample size, the statistical properties of factors may be more stationary from one sample to the next than they are for assets. Finally, it may be that reducing the dimensionality of a large set of assets to a smaller set of factors reduces noise more effectively than reducing dimensionality to a smaller set of assets. We offer empirical evidence of the validity, or lack thereof, of these testable conjectures.

FACTS ABOUT FACTORS

It has recently become fashionable to allocate portfolios across factors instead of assets, but the motivation for doing so is often misguided. For example, some claim that factors are less correlated with each other than assets; hence, they enable investors to achieve greater diversification. What is often overlooked, however, is that the asset combinations used to mimic factors include both long and short positions. It is the inclusion of short positions that give factors the appearance of lower correlations. The reality is that it is impossible to produce more efficient portfolios, in sample, by expressing exposures as factors instead of assets, as long as the investable units are the same in both cases.[1]

Is there any advantage to factor allocation? Yes, potentially. Investors may be able to predict factor behavior more dependably than asset behavior for several reasons. Perhaps some investors are better able to relate current information to future factor performance than they are to asset performance. Unfortunately, we are unable to test this conjecture generically, as this skill varies from investor to investor. Or it may be that risk estimated from high-frequency returns gives more reliable estimates of risk over longer horizons for factors than for assets.[2] Or perhaps the statistical properties of large samples may better predict the statistical properties of small samples more dependably for factors than for assets. Or, for the same sample size, the statistical properties of factors may be more stationary from one sample to the next than they are for assets. Finally, it may be that reducing the dimensionality of a large set of assets to a small set of factors reduces noise more effectively than reducing dimensionality to a small set of assets. If these conjectures are true, factor stratification will dominate asset stratification, because factors will produce more reliable results across return intervals or from large samples to small samples or across independent samples. Otherwise, factor allocation may be yet another fad in a long history of investment fads.

We begin by repudiating the notion that factors offer superior diversification benefits. We then introduce our metrics for evaluating stationarity. Next we define the assets and factors used in our analysis, and we describe the data. We then present our results, which quantify four types of estimation error: interval error, small-sample error, independent-sample error, and noise arising from redundant dimensionality. We conclude with a summary of the relative merits of asset and factor allocation.

The Diversification Benefits of Factors – or Not

Some investors believe that factors offer superior diversification benefits relative to assets because factors are less correlated with each other. This argument is specious if the factors represent regroupings of the assets, even if these regroupings are less correlated with each other than the component assets. The factors would be less correlated only because they would include some short exposures to the assets. Exhibit 1 reveals that assets deliver the same degree of efficiency as factors.

Exhibit 1: Efficient Frontiers of Factors and Assets

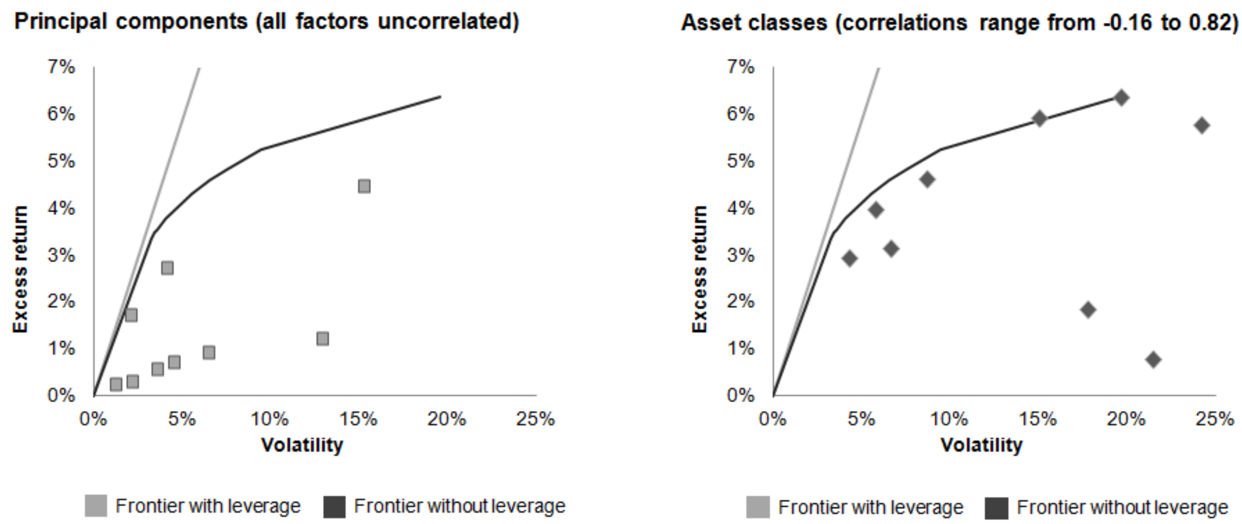

Source: State Street Global Exchange.
Notes: This analysis incorporates the following asset classes: U.S. large cap, U.S. small cap, EAFE equities, emerging equities, global sovereigns, U.S. government bonds, U.S. corporate bonds, commodities, and hedge funds. Based on monthly returns over the period Jan 1990 through Dec 2013. Excess returns represent the return over the risk-free rate. All data obtained from DataStream.

The left panel shows efficient frontiers with and without leverage composed from principal components, which by construction, are uncorrelated with each other. The right panel shows efficient frontiers with and without leverage composed from the assets that produced the principal components. The efficient frontiers on the left and right are identical to each other.[3] This outcome must always be the case, in sample, given three conditions: 1) assets define the opportunity set; 2) both sets of frontiers are subject to the same constraints, and 3) the results are shown in the same return units as the inputs. Moreover, this result will prevail as well for efficient frontiers composed from fundamental factors or security attributes that are constructed from assets.

Sources of Estimation Error

Investors often rely on long samples of historical data to forecast returns, standard deviations, and correlations over shorter future periods such as a few years. These forecasts are subject to three sources of estimation error. First, interval error arises when high-frequency estimates of standard deviation and correlation, such as those derived from monthly returns, are used to forecast lower-frequency standard deviations and correlations of multiple-year returns. Second, small-sample error arises when return and risk parameters from a long sample are used to forecast the outcome of a specific smaller sample. Even though the true parameter of a long sample is known, the realization of that parameter in a shorter sub-sample may be meaningfully different. Third, independent-sample error arises when known parameters from one sample are projected onto a separate, independent sample. We next describe how we measure each of these distinct sources of error.

Interval Error

Financial analysts typically estimate standard deviations and correlations from monthly or higher-frequency returns when determining the optimal composition of a portfolio, which is intended to be held for horizons as long as several years. Analysts implicitly assume that the standard deviations and correlations they estimate from monthly returns pertain as well to longer periodicities. Specifically, they assume that standard deviations scale with the square root of time and correlations estimated from high-frequency returns are similar to longer-interval correlations. However, these properties hold only if asset returns are independently distributed across time, which means that auto-correlations and lagged cross-correlations are zero. Evidence reveals that lagged correlations are significantly non-zero.[4]

Equation (1) shows how high-frequency standard deviations are related to low-frequency standard deviations. The left-hand side of Equation (1) is the standard deviation of the cumulative continuous returns of x over q periods, $x_t+\ldots+x_{t+q-1}$, where σ_x on the right-hand side is the standard deviation of x measured over single-period intervals and $\rho_{x_t,x_{t+k}}$ is the auto-correlation of x at various lags.

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

If these auto-correlations are non-zero, the standard deviation of the longer-interval returns will differ from the product of the shorter-interval returns and the square root of the number of shorter-intervals within the longer interval. If the auto-correlations at all lags equal zero, the standard deviation of x will scale with the square root of the horizon, q.

Equation (2) relates the correlation between the cumulative returns of x and y over q periods (left-hand side) to the correlation between x and y measured over single-period intervals (right-hand side).

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

The numerator equals the covariance of the assets taking lagged cross-correlations into account, whereas the denominator equals the product of the assets' standard deviations as specified in Equation (1).[5] Equation (2) reveals that longer-interval correlations (left-hand side) will differ from shorter-interval correlations to the extent the auto-correlations or the lagged cross correlations of x and y differ from zero.

We capture the error due to non-zero lagged correlations, which we term "interval error," as follows. We first estimate the parameter (standard deviation or correlation)[6] using the full-sample, one-month returns, which we denote as x_m. Next we estimate the parameter using the full-sample, three-year rolling returns, which we denote as $x_{tri}/\sqrt{36}$.[7] We then compute the absolute difference of these estimates and scale this difference by the full-sample standard deviation of three-year returns, as shown in Equation (3).[8] We perform these calculations within the same sample in order to isolate this effect from errors that arise from having different estimation and realization samples.

$$IE = \frac{|x_m - x_{tri}/\sqrt{36}|}{\sigma_{tri}/\sqrt{36}} \tag{3}$$

If x_m equals 0.04 and $\sigma_{tri}/\sqrt{36}$ equals 0.05, the interval error would equal 0.20:

$$IE = \left(\frac{|0.04 - 0.05|}{0.05}\right) = 0.20$$

This means that the parameter value estimated from monthly returns is 0.20 standardized units away from the parameter value estimated from three-year returns, expressed in monthly units. We use this measure to determine whether high-frequency estimates of asset or factor standard deviations and correlations yield more reliable estimates of longer-interval risk.

Small-Sample Error

Investors typically construct portfolios based on inputs of expected return and risk estimated from a history of returns which is longer than their investment horizon. They are thus subject to small-sample error.

To isolate small-sample error we first estimate the parameter (return, standard deviation, or correlation) from our full sample of data, and denote this value as x_m. Next we re-estimate the parameter from all realization sub-samples of 36 months, which we denote as $x^*_{m,r}$. These sub-samples are not independent of the full sample; hence we are isolating the small-sample effect. We then calculate the root of the average squared difference between the parameter estimated from the full sample and the parameter estimated from the realization samples, standardized by full-sample triennial standard deviation (except in the case of correlation), as shown in Equation (4).

$$SSE = \frac{\sqrt{\frac{1}{n}\Sigma\left(x^*_{m,r} - x_m\right)^2}}{\sigma_{tri}/\sqrt{36}} \qquad (4)$$

Suppose that x_m equals 0.04 and $\sigma_{tri}/\sqrt{36}$ equals 0.05. Also suppose that $x^*_{m,r}$ from one sub-sample equals 0.06 and from another sub-sample it equals 0.035. The small-sample error in this case would equal 0.425:

$$SSE = \frac{\sqrt{(0.02^2 - 0.005^2)/2}}{0.05} = 0.425$$

This means that small-sample realizations introduce a standardized error of 0.425 relative to the true value.

Independent-Sample Error

Investors also face independent-sample error because they project historical estimates of return, standard deviation, and correlation onto a future period which is independent of the historical sample. We capture this error by first estimating the parameter from all estimation samples of 36 months, which we denote as $\hat{x}_m$. Next we re-estimate the parameter from all independent, contiguous realization samples of 36 months, which we denote as $x^*_{m,r}$. We then calculate the root of the average squared difference between the parameter estimated from the estimation samples and from the realization samples, and we standardize the average squared difference for return and standard deviation but not correlation. Because both the estimation samples and the realization samples are subject to small-sample error, to isolate the incremental impact of independent samples, we subtract the previously estimated small-sample error, as shown in Equation (5).[9]

$$ISE = \frac{\sqrt{\frac{1}{n}\sum (x^*_{m,r} - \hat{x}_m)^2}}{\sigma_{tri}/\sqrt{36}} - \frac{\sqrt{\frac{1}{n}\sum (x^*_{m,r} - x_m)^2}}{\sigma_{tri}/\sqrt{36}} \qquad (5)$$

If we assume that $\hat{x}_m$ equals 0.04 and $x^*_{m,r}$ equals 0.06 in one case, and they equal 0.055 and 0.045 respectively in another case, and $\sigma_{tri/\sqrt{36}}$ equals 0.05 and small-sample error equals 0.425, then the independent-sample error would equal 0.075:

$$ISE = \frac{\sqrt{(0.02^2 - 0.01^2)/2}}{0.05} - 0.425 = .075$$

Hence, independent-sample errors introduce an incremental standardized error of 0.075 relative to the true parameter.

Exhibit 2 presents a visualization of these stationarity metrics.[10]

Exhibit 2: Visualization of Stationarity Metrics

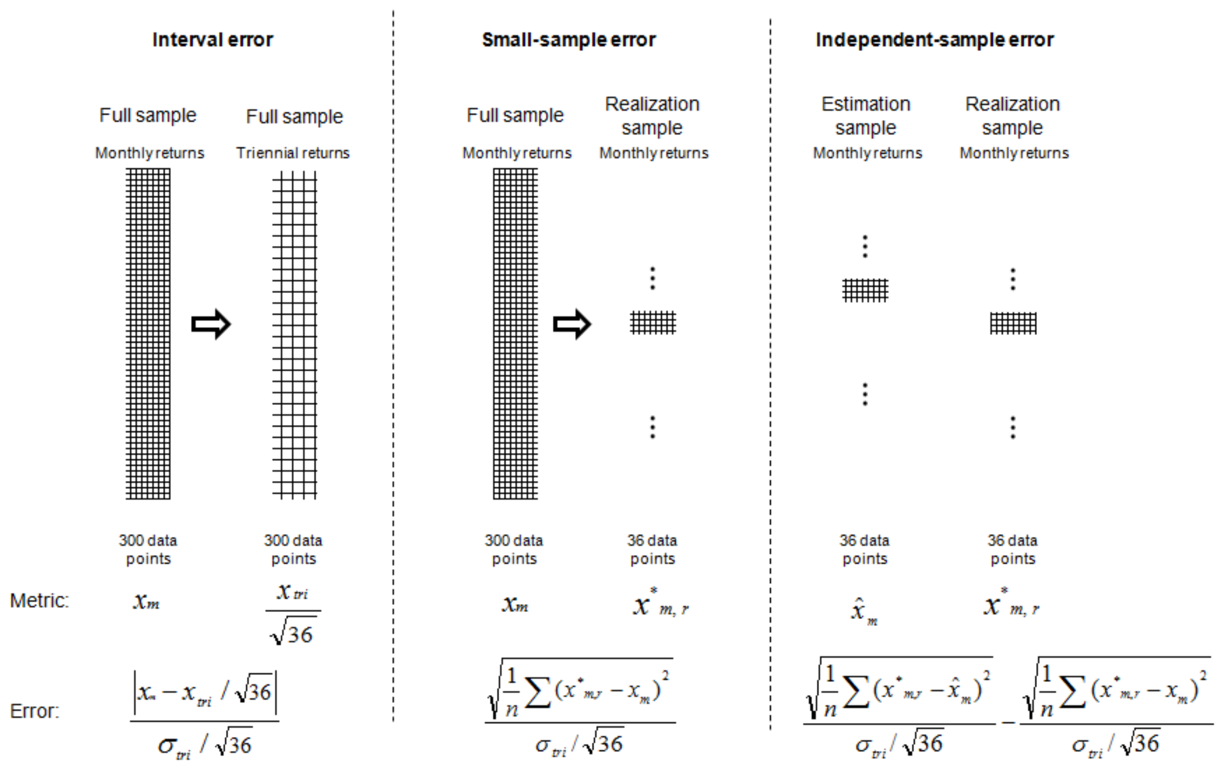

Assets, Factors, and Data

Until now we have used the term, "asset," to refer to macro asset classes such as stocks and bonds and industry groupings, and we have used the term, "factor," to refer to any grouping that is not an asset class or industry. Going forward we distinguish between two types of assets: asset classes and industry groupings. And we distinguish between three types of factors: fundamental factors, security attributes, and statistical factors derived from principal components analysis.

Our first set of experiments compares the stationarity of broad asset classes to the stationarity of fundamental factors and principal components. Exhibit 3 presents the specific groupings.

Exhibit 3: Asset Classes, Fundamental Factors, and Principal Components

	Six dimensions	Three dimensions
Asset classes	U.S. large cap U.S. small cap	Equities
	U.S. government bonds U.S. corporate bonds	Fixed income
	Commodities Hedge funds	Alternatives
Fundamental factors	Inflation Growth	Macro factors
	Term premium Credit premium	Fixed income factors
	Size premium Value premium	Equity factors
Principal components	6 principal components	3 principal components

We create factor mimicking portfolios for the fundamental factors by regressing the asset class returns on the factor values. The specific asset class indexes and fundamental factor time series are reported in Exhibit 4. We use principal components analysis to form six portfolios representing statistical factors. Each portfolio corresponds to an eigenvector. We also evaluate groupings of three dimensions. We regress the six asset classes on three broader fundamental factor categories to derive three groupings of fundamental factors. And we use the top three eigenvectors instead of the top six. It is important to note that all of these groupings represent different stratifications of the same underlying investable units.

Exhibit 4: Indexes and Time Series

Name	Source	Transformation	Period
Asset classes			
Equities			
U.S. Large Cap	S&P 500 Composite	Log return	Jan 1990 - July 2014
U.S. Small Cap	Russell 2000	Log return	Jan 1990 - July 2014
Fixed Income			
U.S. Government Bonds	Barclays U.S. Aggregate Government	Log return	Jan 1990 - July 2014
U.S. Corporate Bonds	Barclays U.S. Aggregate Corporate	Log return	Jan 1990 - July 2014
Alternatives			
Commodities	S&P GSCI Commodity	Log return	Jan 1990 - July 2014
Hedge Funds	HFRI Fund of Funds Composite	Log return	Jan 1990 - July 2014
Fundamental factors			
Macro			
Inflation	Simple return of U.S. SA CPI	Difference	Jan 1990 - July 2014
Growth	One year ahead U.S. GDP growth forecast	Difference	Jan 1990 - July 2014
Fixed Income			
Term premium	10-Year Minus 2-Year Treasury	Difference	Jan 1990 - July 2014
Credit premium	Baa Corporate Yield to 10-Year Treasury	Difference	Jan 1990 - July 2014
Equity			
Small minus Big	Fama-French SMB factor	Log return	Jan 1990 - July 2014
High minus Low	Fama-French HML factor	Log return	Jan 1990 - July 2014

We also compare the stationarity of various industry classifications to security attribute classifications and principal components derived from the 222 individual equity securities that have been present in the MSCI U.S. Index from January 1989 through January 2015. Exhibit 5 shows how we stratify the U.S. equity market.

Exhibit 5: Industries, Attributes, and Principal Components

	49 dimensions	24 dimensions	10 dimensions
Industries	49 portfolios formed on GICS level III classifications.	24 portfolios formed on GICS level II classifications.	10 portfolios formed on GICS level I classifications.
Size	49 portfolios formed on market caps.	24 portfolios formed on market caps.	10 portfolios formed on market caps.
Value	49 portfolios formed on book-to-market ratios.	24 portfolios formed on book-to-market ratios.	10 portfolios formed on book-to-market ratios.
Momentum	49 portfolios formed on 1-year moving average returns.	24 portfolios formed on 1-year moving average returns.	10 portfolios formed on 1-year moving average returns.
Principal components	49 principal components	24 principal components	10 principal components

We stratify this universe into three industry groupings of 49, 24, and 10 dimensions corresponding to GICS classifications. We then stratify the same universe into 49, 24, and 10 quantiles of market capitalization, book-to-market value ratios, and trailing one-year returns to form size, value, and momentum portfolios. Finally, we construct groupings of 49, 24, and 10 statistical factors using principal components analysis.

Results

We first compare asset classes to fundamental factors and principal components. We aggregate the errors, first showing interval error, then small-sample error, and finally independent-sample error. Exhibits 6, 7, and 8 show the aggregation of errors for standard deviation, correlation, and return. Within each exhibit we compare errors associated with asset class stratification, fundamental factor stratification, and principal component stratification.

Exhibit 6: Errors in Standard Deviation

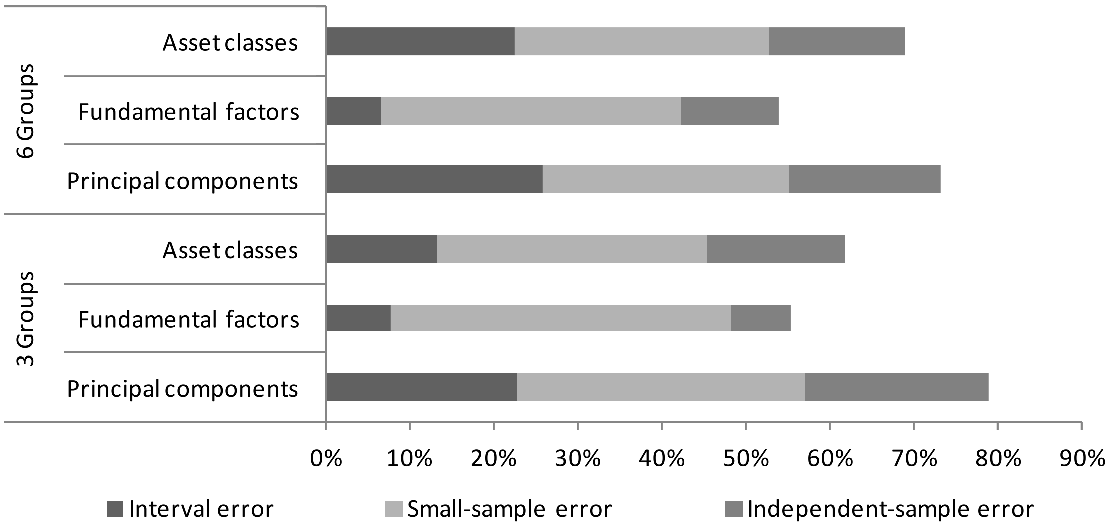

Exhibit 7: Errors in Correlation

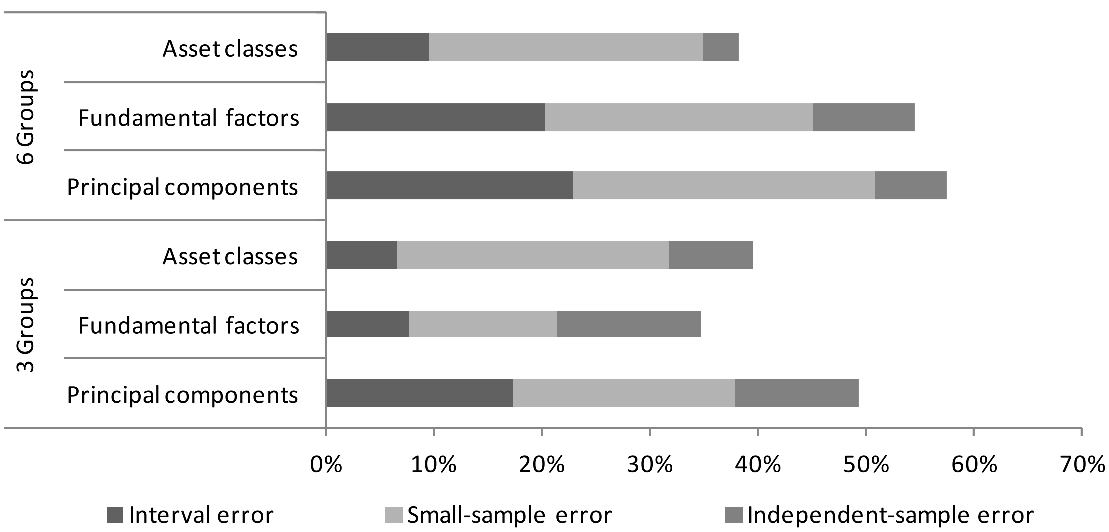

Exhibit 8: Errors in Return

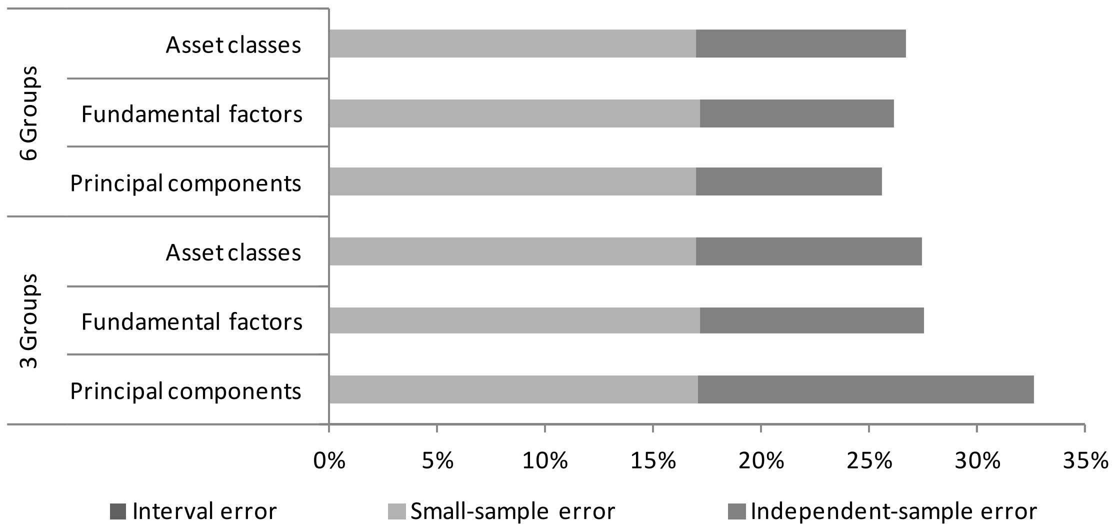

There are only two sources of errors for return because high- and low-frequency estimates of return are always the same. In some cases, assets classes yield the smallest errors, in other cases fundamental factors do, and in still other cases, though less often, principal components result in the smallest errors. It is difficult to generalize from these bar charts about the relative stationarity of asset classes, fundamental factors, and principal components. We therefore present the average errors across standard deviation, correlation, and return, and across only standard deviation and correlation. The reason we show the average of standard deviation and correlation with and without return is because most investors do not rely solely on historical means to estimate expected returns. They may instead estimate equilibrium returns from historical covariances or form views based on fundamental analysis. Or they may blend historical means with equilibrium returns or views. Exhibit 9 presents this information.

Exhibit 9: Average Errors for Asset Classes,

Fundamental Factors, and Principal Components

Classification	Standard Deviation, Correlation, and Return	Standard Deviation and Correlation
Asset Classes	0.44	0.52
Fundamental Factors	0.47	0.57
Principal Components	0.53	0.65

Exhibit 9 reveals that asset classes, on average, have more stationary statistical properties than either fundamental factors or principal components, both including and excluding return. One might argue that averages are misleading because errors in standard deviation are more important than errors in correlation, which of course, is true. But it is also true that beyond three assets or factors, there are more correlations than standard deviations.

We next turn to a comparison of industries, which we construe as assets,[11] portfolios based on security attributes, and principal components formed from the underlying securities. Exhibits 10, 11, and 12 show the aggregation of errors for standard deviation, correlation and return based on these equity market groupings.

Exhibit 10: Errors in Standard Deviation

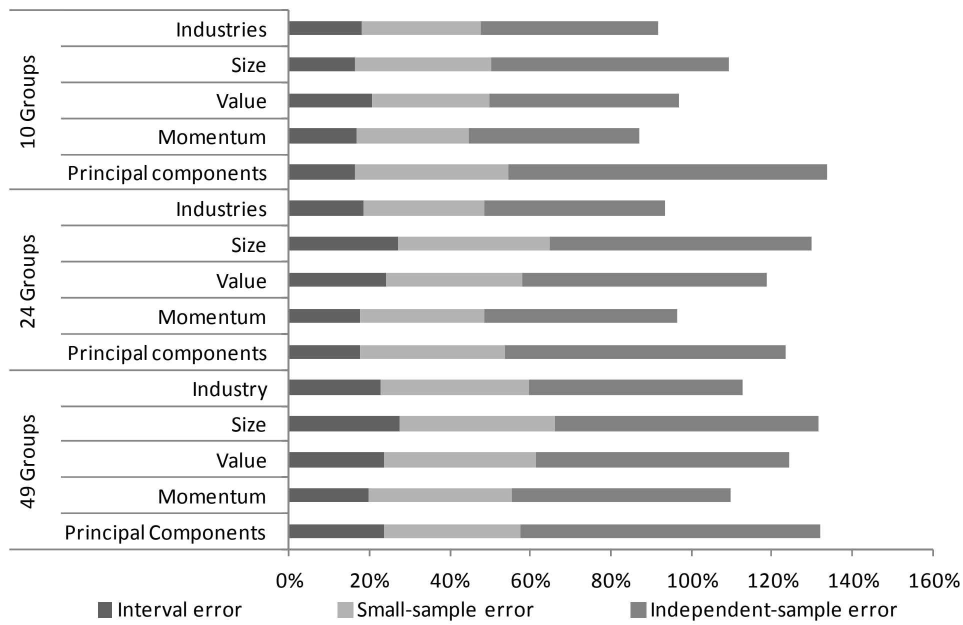

Exhibit 11: Errors in Correlations

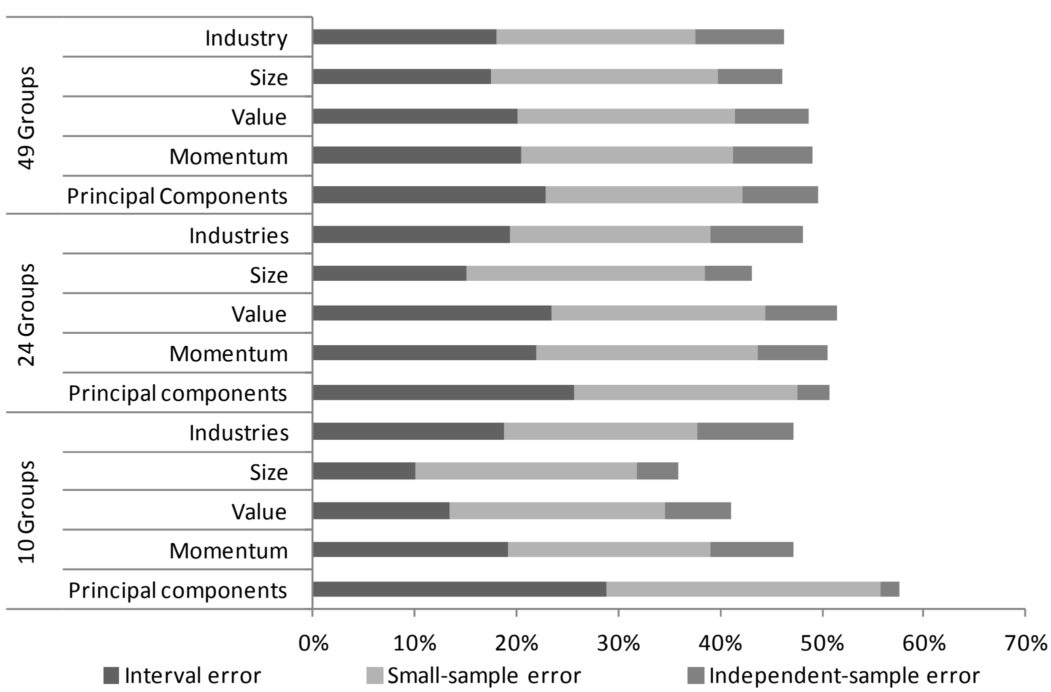

Exhibit 12: Errors in Return

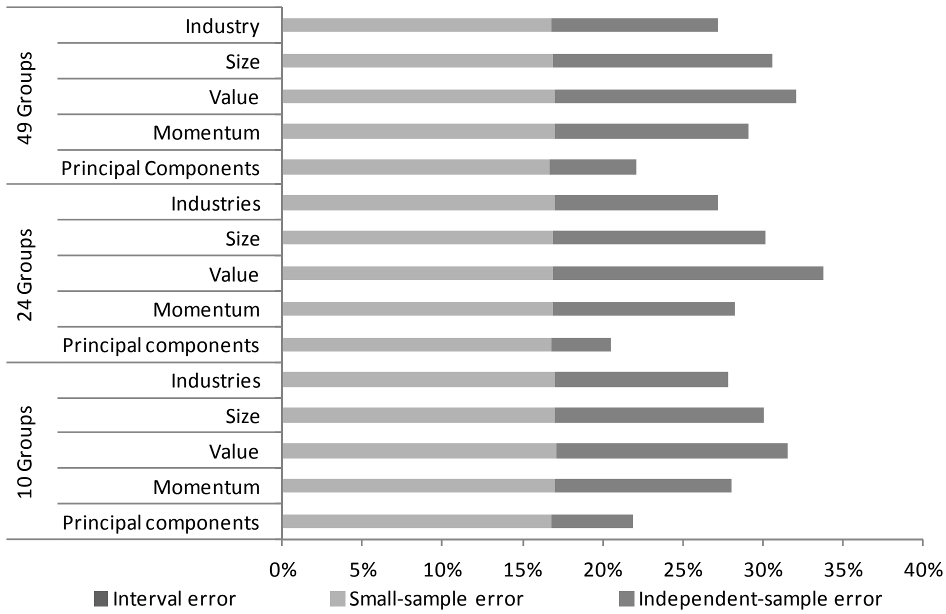

Again, it is difficult to generalize about the relative stationarity of industries, attributes, and principal components from these bar charts. Thus we again summarize these results by averaging the errors across standard deviation, correlation, and return, and across standard deviation and correlation. We present these averages in Exhibit 13.

Exhibit 13: Average Errors for Industries,

Attributes, and Principal Components

Classification	Standard Deviation, Correlation, and Return	Standard Deviation and Correlation
Industries	0.48	0.59
Attributes	0.52	0.63
Principal Components	0.56	0.74

Again we find that the statistical properties of assets, defined as industries, are more stationary than factors, defined as attributes or principal components.

Finally, we turn to the issue of noise reduction. As stated earlier, it may be the case that consolidating a larger set of assets into a smaller set of factors reduces noise more effectively than consolidating a larger set of assets into a smaller set of assets. Even if this is true, it may or may not be a good outcome. We know, a priori, that a more complete set of assets or factors will produce better results in sample than a less complete set to the extent the additional assets or factors are not purely redundant. This result occurs because greater granularity provides additional information. When we move out-of-sample, however, the more granular information may degrade more severely than the composite information because it is less stationary. We thus face the following trade-off. Should we take a more granular approach to portfolio construction in order to capture additional information, noisy though it may be, or should we approach portfolio construction in a more consolidated way, thereby sacrificing information in favor of noise reduction? This trade-off is an empirical issue which depends on the specific data and portfolio construction algorithm. Nonetheless we report the noise reduction associated with consolidating the full set of 222 surviving securities in the MSCI U.S. Index into various sets of industry and attribute groupings. Exhibit 14 shows the average percentage reduction in non-standardized errors across standard deviation, correlation, and return and then across only standard deviation and correlation. In this analysis, we do not divide the errors by standard deviation because doing so would obscure the level of noise, as the average standard deviation of individual securities is greater than the average standard deviation of industries and attributes.[12]

Exhibit 14: Noise Reduction for Industries and Attributes

Standard Deviation, Correlation, and Return			
	10 Groups	24 Groups	49 Groups
Industries	20%	11%	4%
Attributes	22%	13%	12%

Standard Deviation and Correlation			
	10 Groups	24 Groups	49 Groups
Industries	19%	11%	3%
Attributes	22%	10%	8%

We find no compelling evidence that consolidating with factors reduces noise more effectively, on average, than consolidating with assets.

Summary

Investors have become infatuated with factors, often based on faulty logic or incomplete information. Some investors propose using factors instead of assets as the building blocks for forming portfolios, because factors appear to be less correlated with each other than assets and thus seem to afford greater potential for diversification. These ostensibly low correlations, however, reflect the fact that the asset combinations used to mimic factors include short positions in some of the assets and not superior diversification opportunities. We have shown that it is impossible to generate a superior in-sample portfolio, given the same constraints, by regrouping assets into factors if the investable units are assets from which the factors are formed.

We also considered the possibility that investors are more skilled at relating current information to future factor behavior than to future asset behavior. However, we conceded that we could not test this conjecture generically, because skill is investor-specific.

However, we did test whether the statistical properties of factors are more stationary than the statistical properties of assets. We found no evidence that factors produce more stable results across varying frequencies, nor from large samples to small samples, nor across independent samples. On the contrary, we found evidence of the opposite, on average.

Finally, we tested whether reducing the dimensionality of a larger set of assets to a smaller set of factors reduces noise more effectively than reducing the dimensionality to a smaller set of assets. Again, we found no evidence that factors are meaningfully more effective than assets at noise reduction.

Where does this leave us? In our view, the case is yet to be made that investors should use factors as building blocks for forming portfolios rather than assets. We do believe, however, that investors may be able to gather useful insights about the performance of their portfolios by attributing performance to factor exposures in addition to asset exposures. And we believe as well that investors may be able to manage risk more effectively by considering a portfolio's factor exposures in combination with its asset exposures. Beyond these applications, investors must decide for themselves whether they have greater conviction about future factor behavior or future asset behavior.[13]

References

Idzorek, T., M. Kowara. “Factor-Based Asset Allocation vs. Asset-Class-Based Asset Allocation.” *Financial Analysts Journal*. May/June 2013, Vol. 69, No. 3.

Kinlaw, W., M. Kritzman, and D. Turkington. “The Divergence of High- and Low-Frequency Estimation: Causes and Consequences.” *The Journal of Portfolio Management*. Special 40th Anniversary Issue. 2014.

Kinlaw, W., M. Kritzman, and D. Turkington. “The Divergence of High- and Low-Frequency Estimation: Implications for Performance Measurement.” *The Journal of Portfolio Management*. Spring 2015.

[1] This result assumes that portfolio outputs are measured at the same return frequency as inputs and that both portfolio construction processes are subject to the same constraints.
[2] Within this context, high-frequency refers to intervals such as a month rather than fractions of a second.
[3] For a mathematical proof of equality, please refer to Idzorek and Kowara [2013].
[4] For more information, please refer to Kinlaw, Kritzman, and Turkington [2014] and Kinlaw, Kritzman, and Turkington [2015].
[5] For more detail about the mathematics of the relationship between high- and low-frequency estimation of these risk parameters, see Kinlaw, Kritzman, and Turkington [2014].
[6] We do not include return because it is unaffected by non-zero lagged correlations.
[7] We divide by the square root of 36 for standard deviation, but not for correlation.
[8] We only standardize standard deviation by three-year return standard deviation, not correlation.
[9] More specifically, the error between contiguous small samples may arise from three sources: (1) Using a small sample to estimate the true parameter (2) Using one small sample to forecast a realization over a different, random small sample (3) Using one small sample to forecast a realization over a different, contiguous small sample. This last source of error captures the impact of serial correlation in the parameter (standard deviation, correlation). This is different than interval error, which captures the impact of serial correlation in returns. It is important to clarify that by subtracting small-sample error from contiguous small sample error, we isolate the incremental impact of independent-sample error. We focus on the incremental impact throughout this paper.
[10] It is important to clarify that each of these represents a distinct source of error.

[11] Philosophically, it is difficult to label industries as a form of asset stratification and security attributes as a form of factor stratification. We do so only because investors have historically allocated across industries more typically than across security attributes.

[12]We do not report noise reduction for principal components because there is a wide variation in standard deviation across principal components, which translates into a wide variation of errors in return and standard deviation. It is therefore meaningless to aggregate non-standardized results from principal component groupings. It only makes sense to aggregate these results once they are standardized.

[13] We thank Roy Henrikkson, Paul Mende, and Will Kinlaw for helpful comments. Views expressed in this article are those of the authors, and may not represent the views of State Street Corporation or its affiliates.