PORTFOLIO MANAGEMENT RESEARCH

The Journal of Portfolio Management

JPM.pm-research.com

QUANTITATIVE STRATEGIES: FACTOR INVESTING

SPECIAL ISSUE 2017

Facts about Factors

Paula Cocoma, Megan Czasonis, Mark Kritzman, and David Turkington

Facts about Factors

PAULA COCOMA, MEGAN CZASONIS, MARK KRITZMAN, AND DAVID TURKINGTON

PAULA COCOMA
is a PhD student at INSEAD in Fontainebleau, France.
paula.cocoma@insead.edu

MEGAN CZASONIS
is a vice president at State Street Associates in Cambridge, MA.
mczasonis@statestreet.com

MARK KRITZMAN
is the CEO of Windham Capital Management in Boston, MA, and a senior lecturer at MIT Sloan School of Management in Cambridge, MA.
kritzman@mit.edu

DAVID TURKINGTON
is a senior vice president at State Street Associates in Cambridge, MA.
dturkington@statestreet.com

It is becoming increasingly common for investors to stratify portfolios into factors rather than assets. This approach may be advantageous for a variety of reasons. For example, certain factor structures may offer risk premiums or even alphas that are not otherwise available from conventional asset structures. Or investors may discover unwanted risk exposures by decomposing a portfolio into factors. Or factor stratification may help to explain a portfolio's performance. We acknowledge these important benefits to factor stratification, but our focus in this article is to evaluate the use of factors as building blocks for forming portfolios.

We argue that the motivation for allocating portfolios to factors rather than assets is often misguided. For example, some claim that factors are less correlated with each other than are assets; hence, they enable investors to achieve greater diversification. It is also argued that reducing the dimensionality of a large set of assets into a smaller set of factors reduces noise. And some believe that investors may be able to predict factor behavior more dependably than asset behavior. Finally, some argue that the risk properties of factors are more stable than the risk properties of assets. We address each of these assertions and conclude that investors should not replace assets with factors as the building blocks for forming portfolios.

We illustrate our arguments with factors that are fixed-weight linear combinations of underlying assets. Although many investors define factors this way and stratify their portfolios accordingly, we acknowledge that there are a variety of ways to define factors, including proprietary replication methods that may depend on dynamic rebalancing. We cannot address factors that depend on proprietary replication processes, but we see no reason why our general conclusions would not apply to these factor descriptions as well.

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. The argument that factors can offer superior diversification is specious if the factors represent regroupings of the assets, even if these regroupings are less correlated with each other than the component assets.[1] The factors would be less correlated only because they would include some short exposures to the assets. Assets will always deliver the same degree of efficiency as factors 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. This equality can be proven mathematically, as in Idzorek and Kowara [2013].

We understand that after careful consideration, most people will agree with this conclusion. However, given the amount of confusion around this subject, it may also be helpful to look at a simple empirical example. Although this equality applies to any collection of factors—including fundamental factors (such as inflation and GDP) and attributes (such as value, size, and momentum)—it is convenient to illustrate it for principal components because principal component factor portfolios are uncorrelated by construction and they span the same opportunity set as the underlying assets.[2] Exhibit 1 shows that efficient frontiers with and without leverage are identical regardless of whether they are formed from principal components or from the underlying assets that produced the principal components.[3]

Exhibit 1
Efficient Frontiers of Factors and Assets

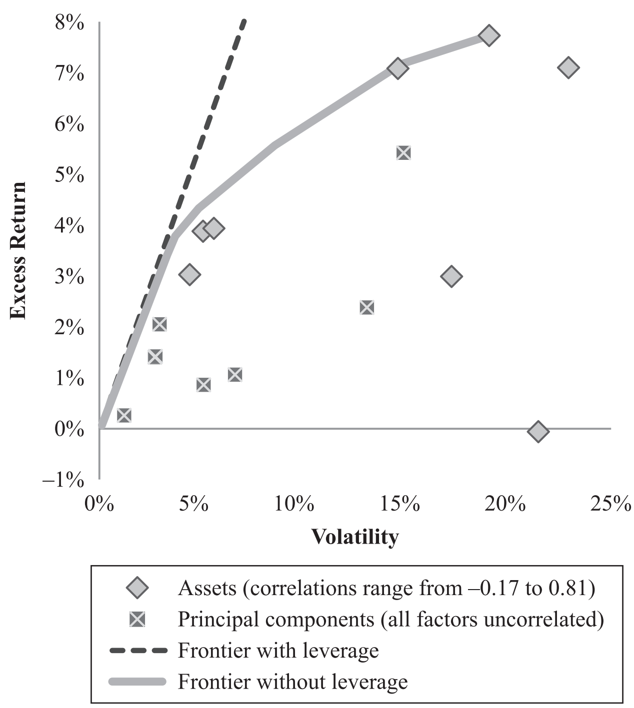

Notes: This analysis incorporates the following asset classes: U.S. large cap, U.S. small cap, EAFE equities, emerging equities, U.S. government bonds, U.S. corporate bonds, commodities, and hedge funds. It is based on monthly returns over the period January 1990 through December 2015. Excess returns represent the return over the risk-free rate. It is important to note that there are indeed eight distinct principal component factors represented in this exhibit. Two of them have very similar volatility and excess return: (2.73%, 1.40%) and (2.75%, 1.42%).

NOISE REDUCTION

Some proponents of factor stratification argue that consolidating a larger set of assets into a smaller set of factors reduces noise. This is only partly true, and even so, it may or may not be a good outcome. We know, a priori, that more granularity produces better results in sample than less granularity to the extent that 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 if 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?

Suppose our focus is to reduce out-of-sample noise. Is it really true that coarser data is less noisy? It is true for returns but not for covariances. We know that returns for consolidated groups of assets will be less noisy than the returns of the underlying assets because of diversification. But the effect of diversification should apply equally to asset classes and factors.

We have already shown that efficient frontiers composed of factors are identical to efficient frontiers composed of assets. It follows, therefore, that we cannot achieve greater reduction in dispersion around means by using factor groupings rather than asset groupings. Consider principal components as an example. The full set of principal component eigenvectors spans the full opportunity set of the assets, and the sum of principal component variances will always equal the sum of asset variances. Therefore, the amount of noise in returns around their means is the same in both cases. If we were to reduce dimensionality by retaining only the top few eigenvectors, we would be guaranteed to increase noise in returns because the top eigenvectors are the least stable by construction (they have the largest variance).

Although diversification reduces noise in returns around their means, it does not necessarily reduce noise in covariances. At the portfolio level, the only source of risk instability is the noise in the portfolio's variance. Just as factors cannot yield more efficient portfolios than assets in sample, they also cannot yield portfolios with less noise in variances because assets define the opportunity set. It is possible that different portfolios on the efficient frontier have different levels of noise in their variances. However, for an investor who seeks to maximize expected utility, there is no reason to believe that optimal portfolios will be biased toward assets with either more noise or less noise in covariances.

For a universe of assets or factors, risk instability pertains to the entire covariance matrix. If asset returns are normally distributed and serially independent, we should not expect the average noise in covariances across a consolidated group of assets or factors to be less than the average noise in covariances across the individual assets that form those groups. In fact, if correlations are closer to zero as a result of grouping, we would actually expect the noise in covariances to increase because correlations closer to zero are inherently less stable.[4]

For assets with non-normal distributions or serial dependence, different methods of grouping may result in different levels of noise. We see no clear reason to expect that factor groupings would consistently result in less noise than asset class groupings, but we can test this hypothesis empirically. In the data sets we test, we find no evidence that factor groupings reduce noise more effectively than asset class groupings. In fact, we find the opposite, as we will show in the results section later on.

PREDICTIVE SKILL

Some investors may have particular skill in predicting factors as opposed to predicting assets. It is also possible that many investors have more skill in predicting assets than factors. We cannot test this conjecture generically because it depends on individual investor skill. It is important to note, however, that investors who choose to predict factor performance face two additional hurdles that do not apply to assets. First, factors are not directly investable. Investors must identify a collection of assets that mimic the movement of factor values, which they typically do by regression analysis. However, the coefficients derived from historical data will differ from those that yield the best fit in the future. The difference in past and future coefficients is called *mapping error*. Second, even if the true asset-to-factor mappings were known in advance, investors would need to trade periodically to rebalance the factor portfolios, which would increase transaction costs. Both of these realities erode the value that can be achieved through superior factor prediction.

SOURCES OF INSTABILITY

Investors often rely on long samples of historical data to forecast covariances over shorter future periods such as a few years.[5] These forecasts are subject to four distinct types of error that contribute to instability. First, 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 subsample may be meaningfully different. Second, independent-sample error arises when known parameters from one sample are projected onto a separate independent sample. Third, investors also face mapping error when they use factors to build portfolios, because they must translate these factor values into mimicking combinations of assets. Fourth, interval error arises when the covariances of high-frequency returns, such as monthly returns, differ from the covariances of longer-period returns, such as three-year returns. We next describe how we measure each of these distinct sources of error, which are pictured in Exhibit 2.[6]

Small-Sample Error

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

To isolate small-sample error, we first estimate the covariances of two assets, A and B, from our full sample of monthly data.[7] Covariances are in squared units, which exaggerate outliers in error computations. To mitigate this outlier problem, we take the square root of the standard deviation product in covariances. We retain the correlation of the two assets and multiply it by the product of the square root of their standard deviations. We modify covariance estimates in this way across all our measures of error.

Next, we reestimate the parameter from all overlapping 36-month realization subsamples within our data sample, which we denote using a subscript for each subsample j.[8] These subsamples are not independent of the full sample; hence, we are isolating the small-sample effect. For each realization subsample, we compute the difference between the modified covariances and the full-sample-modified covariances, and we divide by the square root of the product of both assets' full-sample standard deviations

Exhibit 2

Visualization of Sources of Estimation Error

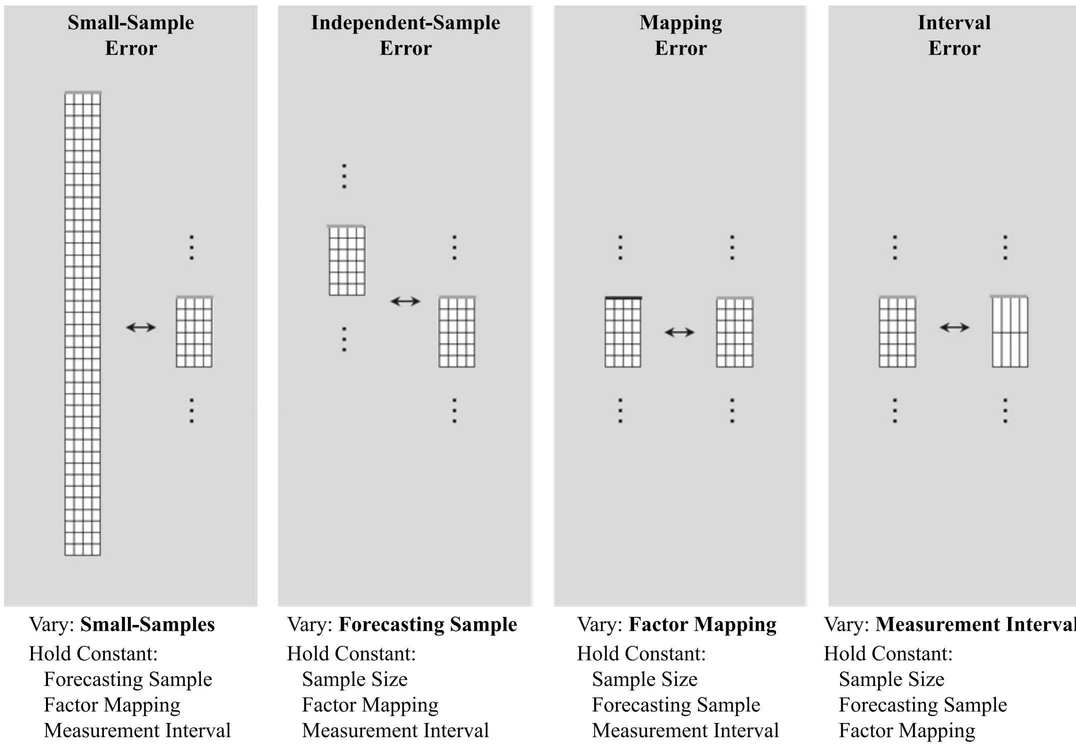

to normalize this quantity. We then compute the root-mean-squared error of these individual normalized errors, which we define as small-sample error. Note that when assets A and B are the same, this formula pertains to the standard deviation of a single asset, rather than to the relationship between two assets:

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

Independent-Sample Error

Investors also face independent-sample error because they project historical estimates of covariances onto a future period that is independent of the historical sample. We capture this error by first estimating statistics from all possible overlapping estimation subsamples of 36 months within our data sample, which we denote using the subscript $\hat{m}, j$. Next, we reestimate the same statistics from all independent realization samples of 36 months, each of which immediately follows one of the 36-month estimation subsamples and is denoted by the subscript m, j.

For each estimation and realization subsample pair, we compute the difference between the estimation sample statistic and the realization sample statistic, and normalize it by dividing by the square root of the full-sample product of the assets' standard deviations, as we did for small-sample error. We then compute the root-mean-squared error across all samples. Because both the estimation samples and the realization samples are subject to small-sample error, we subtract the previously estimated small-sample error in order to isolate the incremental impact of independent-sample error:

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

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 estimated from an independent historical sample. We call this mapping error, and it applies only to factors.[9]

The calculations we have just shown for small-sample error and independent-sample error do not incorporate mapping error. We isolate mapping error by comparing the covariances of factor-mimicking portfolios whose weights are optimized for the realization sample to those whose weights are optimized on an independent prior sample, holding the covariance evaluation period constant. Specifically, for each 36-month estimation subsample, we estimate factor-mimicking portfolios based on the data from that subsample, and we compute the realized covariances of those portfolios over the subsequent 36-month realization subsample. We denote these covariances using the subscripts $\hat{A},m,j$ and $\hat{B},m,j$. We then subtract these covariances from the realized covariances of factor-mimicking portfolios whose weights are derived from the same realization subsample on which they are evaluated, which we denote using the subscripts A,m,j and B,m,j.

The way in which factor-mimicking portfolios are obtained may differ across analyses, depending on the type of data and investment universe we consider. For broad asset allocation, we apply regression or principal components analysis on each subsample. For grouping individual securities, we use security attributes observed in each subsample.

$$ME\,(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}} - \rho_{\hat{A}\hat{B},m,j}\sqrt{\sigma_{\hat{A},m,j}\sigma_{\hat{B},m,j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2} \tag{3}$$

Interval Error

Investors typically estimate covariances from monthly or higher-frequency returns when determining the optimal composition of a portfolio intended to be held for horizons as long as several years. This practice implicitly assumes that the covariances estimated from monthly returns pertain as well to longer periodicities. Specifically, it assumes that standard deviations scale with the square root of time and correlations estimated from high-frequency returns are the same as longer-interval correlations. However, these two relationships hold only if asset returns are independently distributed across time, which means that autocorrelations and lagged cross correlations are zero. Evidence reveals that lagged correlations are significantly nonzero.[10]

Equation 4 shows how high-frequency standard deviations are related to low-frequency standard deviations. The left-hand side of Equation 4 is the standard deviation of the cumulative continuous returns of x over q periods, where σ_x is the standard deviation of x measured over single-period intervals.

$$\sigma(x_t + \cdots + x_{t+q-1}) = \sigma_x\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\,\rho_{x_t,x_{t+k}}} \tag{4}$$

Equation 5 shows how high-frequency correlations are related to low-frequency correlations. The left-hand side of Equation 5 is the correlation between the cumulative returns of x and the cumulative returns of y over q periods. The numerator reflects the covariance of the assets taking lagged correlations into account, whereas the denominator reflects the product of the assets' standard deviations as described by Equation 4. This equation allows us to assume values for the autocorrelations of x and y, as well as the lagged cross correlations between x and y, in order to compute the correlations and standard deviations that these parameters imply for longer horizons.

$$\rho(x_t + \cdots + x_{t+q-1}, y_t + \cdots + y_{t+q-1}) = \frac{q\rho_{x_t,y_t} + \sum_{k=1}^{q-1}(q-k)(\rho_{x_{t+k},x_t} + \rho_{x_t,y_{t+k}})}{\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t,x_{t+k}}}\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{y_t,y_{t+k}}}} \tag{5}$$

Equations 4 and 5 reveal that longer-interval covariances will differ from shorter-interval covariances to the extent lagged correlations differ from zero. If all the lagged correlations are zero, the standard deviation will scale with the square root of the number of periods, q, and the correlation will be invariant to the measurement interval.

We capture the error due to nonzero lagged correlations, which we term *interval error*, as follows. For each realization sample, we first estimate covariances using one-month returns, which we denote using the subscript *m*, *j*, as before. Next, we estimate the annual covariances implied by lagged correlations over the same sample, which we denote using the subscript *ann*, *j*. We perform these calculations within the same subsamples in order to isolate this effect from errors that arise from having different estimation and realization samples. For each subsample, we compute the difference between the statistics estimated from monthly returns and those estimated from annual returns divided by 12,[11] and we divide by the product of both assets' full-sample standard deviations to normalize this quantity. We then compute the root-mean-squared error across all subsamples.

$$IE(A,B) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left(\frac{\rho_{AB,ann,j}\sqrt{\sigma_{A,ann,j}\sigma_{B,ann,j}}/12 - \rho_{AB,m,j}\sqrt{\sigma_{A,m,j}\sigma_{B,m,j}}}{\sqrt{\sigma_{A,m}\sigma_{B,m}}}\right)^2} \tag{6}$$

Total Covariance Error

The error metrics defined above can be applied to the covariances of any collection of assets (or factors). For a given asset universe of *n* assets, we compute the *n*-by-*n* matrix of covariance errors for all asset pairs. We calculate a total covariance error by averaging the squared errors contained in this entire matrix. We weight each element in the matrix equally because they are comparable quantities and each is important to risk measurement and portfolio construction.[12] Total error can be calculated for any component of error.

$$TE = \sqrt{\frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}E(i,j)^2} \tag{7}$$

Computing the average of the errors across the elements of the covariance matrix implicitly recognizes the relative importance of variances and correlations:

1. Any one variance is probably more important than any one correlation. In our computation of total covariance error, errors in variances receive greater weight than errors in correlations, because the error of each covariance in the matrix reflects the error in the correlation of that asset pair as well as the error in each asset's variance. Errors in variance are also present in the diagonal terms in the matrix, and there is no correlation error present in those terms because they pertain to a single asset.
2. However, there are more correlations than variances. Even though an error in one correlation may be less important than an error in one variance, averaging across the entire matrix accounts for the fact that there are many more correlations than there are variances.

Composite Instability Score

The four sources of error introduced above are independent from one another,[13] which means we can sum the squares of each error and then take the square root of this sum to compute a composite instability score.[14]

$$\text{Composite Instability Score (CIS)} = \sqrt{SSE^2 + ISE^2 + ME^2 + IE^2} \tag{8}$$

Errors in Means

The methodologies we have introduced for quantifying parameter instability pertain only to covariances. Investors also care about mean returns. We do not address errors in means for several reasons:

1. The small-sample error of mean returns is proportional to its full-sample variance. As we noted earlier, it is impossible for factors to produce portfolios with lower variance than assets for a given level of expected return.
2. Independent-sample error for mean returns differs across assets, but it is of limited interest because investors do not usually extrapolate past returns to estimate expected future returns.
3. Mean returns are not subject to interval error.[15]

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 categorizations 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 stability of broad asset classes to the stability of fundamental factors and principal components. The specific asset class indexes and fundamental factor time series are shown in Exhibit 3. This data set spans January 1990 through December 2015.[16] We create factor-mimicking portfolios for the fundamental factors by regressing the asset class returns on the factor values and rescaling the betas such that the sum of absolute weights equals one. We use principal components analysis to form six portfolios representing statistical factors. Each portfolio corresponds to an eigenvector. It is important to note that all of these groupings represent different stratifications of the same underlying investable units.

We also compare the stationarity of various industry classifications to security attribute classifications and principal components derived from the 288 individual equity securities that were present in the MSCI USA Index from January 1989 through December 2015 and for which full data history is available for price, market capitalization, and book-to-market valuation. Exhibit 4 shows how we stratify the U.S. equity market.

We stratify this universe into three industry groupings of 56, 24, and 10 dimensions corresponding to GICS classifications.[17] We then stratify the same universe into 56, 24, and 10 quantiles of market capitalization, book-to-market value ratios, and trailing one-year returns to form size, value, and momentum portfolios. Within each group, we weight securities by their market capitalization, and we use average values of attributes within each window for sorting. Finally, we identify the most important sets of 36, 24, and 10 statistical factors using principal components analysis applied to the individual stocks.[18]

For fundamental factors, attributes, and principal components, the object we want to study is a portfolio of assets whose composition will change depending on the data sample used to estimate it. When we measure small-sample error, we compute principal components or factor-mimicking weights using the 36-month subsample, and we evaluate the statistics of that same portfolio over the full sample and the 36-month subsample from which the weights were derived. For independent-sample error, we compute principal components or factor-mimicking weights using the 36-month realization subsample, and we evaluate the portfolio over both 36-month subsamples in the formula. We estimate the standard deviations in the denominator of the formula by applying the weights from the subsample to the full-sample asset returns. For interval error, we estimate principal components or factor-mimicking weights

EXHIBIT 3
Asset Class and Factor Data

Asset Classes		
Equity	U.S. Large Cap	S&P 500
	U.S. Small Cap	Russell 2000
Fixed Income	U.S. Government Bonds	Barclays U.S. Aggregate Government
	U.S. Corporate Bonds	Barclays U.S. Aggregate Corporate
Alternatives	Commodities	S&P GSCI Commodities
	Hedge Funds	HFRI Fund of Funds Composite
Fundamental Factors		
Macro	Inflation	U.S. Consumer Price Index, seasonally adjusted
	Growth	One-year ahead U.S. GDP growth forecast
Fixed Income	Term Premium	10-year minus 2-year U.S. Treasury yield
	Credit Premium	Baa U.S. corporate yield minus 10-year U.S. Treasury yield
Equity	Small-Cap Premium	Fama–French small-minus-big factor
	Value Premium	Fama–French high-minus-low factor

Exhibit 4

Industries, Attributes, and Equity Securities Data

	56 Dimensions	24 Dimensions	10 Dimensions
Industries	56 portfolios formed on GICS level III	24 portfolios formed on GICS level II	10 portfolios formed on GICS level I
Size	56 portfolios formed on capitalization	24 portfolios formed on capitalization	10 portfolios formed on capitalization
Value	56 portfolios formed on book-to-market	24 portfolios formed on book-to-market	10 portfolios formed on book-to-market
Momentum	56 portfolios formed on trailing 1y return	24 portfolios formed on trailing 1y return	10 portfolios formed on trailing 1y return
Principal Components	Top 36 principal components	Top 24 principal components	Top 10 principal components

Exhibit 5

Sources of Errors and the Composite Instability Score: ***Asset Classes, Fundamental Factors, and Principal Components***

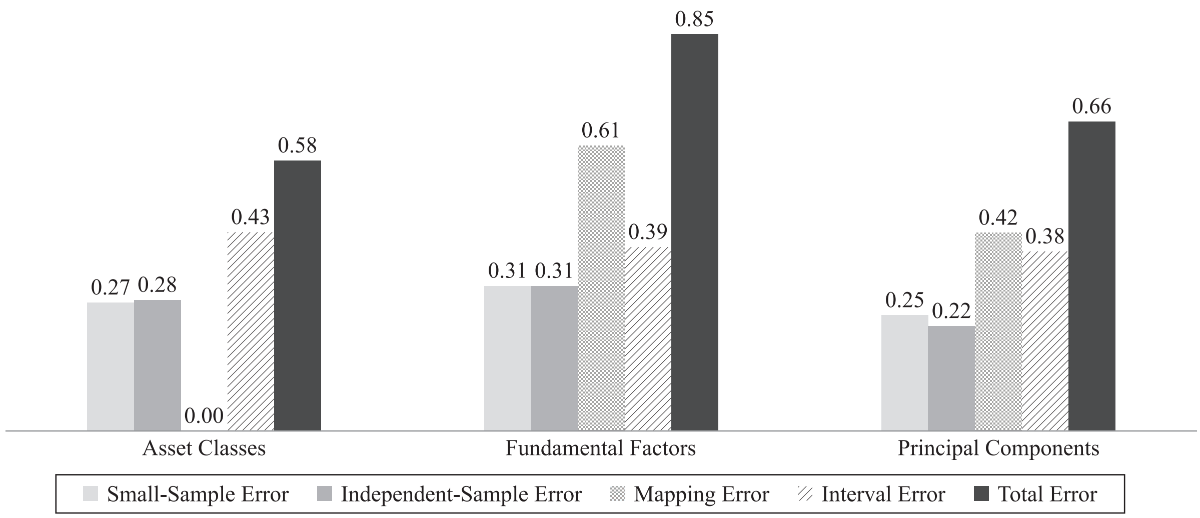

using the relevant 36-month period, and we estimate the standard deviations in the denominator by applying those weights to the full-sample asset returns.

RESULTS

We first compare asset classes to fundamental factors and principal components. Exhibit 5 compares small-sample error, independent-sample error, mapping error, and interval error for these three approaches to stratification, as well the composite instability score.

We next turn to a comparison of industries, which we construe as assets,[19] portfolios based on security attributes, and principal components formed from the underlying securities. Exhibit 6 shows the sources of error and the composite instability score for each approach to stratification in the case of 10 groups. To evaluate the significance of these differences in errors, we constructed 1,000 random groupings of stocks to form a distribution of errors due to random chance. The composite instability score of 0.57 for industries is less than the lowest score for any of the 1,000 random groupings, which is 0.62. This result suggests that there is some useful information in industries that contributes stability to the groups. In contrast, the composite instability scores for attributes (0.67) and principal components (0.87) are both larger than the maximum instability observed in any of the 1,000 random groupings, which is 0.64.

Exhibit 7 shows the composite instability score and each source of error for various levels of dimensionality. Consistent with the intuition discussed earlier, we find that grouping stocks by industry classification has very little impact on the instability of covariances. Reducing

EXHIBIT 6
Sources of Errors and the Composite Instability Score: ***Industries, Attributes, and Principal Components***

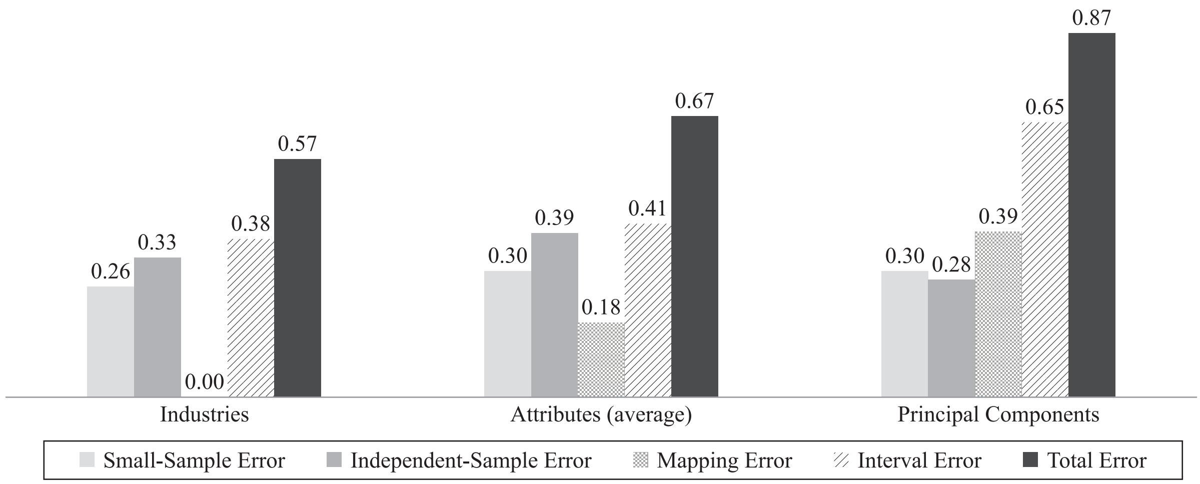

EXHIBIT 7
Noise Reduction for Industries, Attributes, and Principal Components

		56 Groups	24 Groups	10 Groups
Composite Instability Score	**Industries**	**0.52**	**0.56**	**0.57**
(288 stocks = 0.49)	**Attributes (average)**	**0.63**	**0.65**	**0.67**
	Principal Components	**0.67**	**0.72**	**0.87**
Small-Sample Error	Industries	0.23	0.26	0.26
(288 stocks = 0.22)	Attributes (average)	0.26	0.28	0.30
	Principal Components	0.14	0.20	0.30
Independent-Sample Error	Industries	0.28	0.31	0.33
(288 stocks = 0.25)	Attributes (average)	0.32	0.36	0.39
	Principal Components	0.20	0.22	0.28
Mapping Error	Industries	0.00	0.00	0.00
(288 stocks = 0.00)	Attributes (average)	0.27	0.23	0.18
	Principal Components	0.25	0.29	0.39
Interval Error	Industries	0.37	0.38	0.38
(288 stocks = 0.36)	Attributes (average)	0.38	0.40	0.41
	Principal Components	0.55	0.59	0.65

dimensionality through attribute-based groupings or principal components actually increases instability, largely because of the mapping error of factors.[20]

CONCLUSION

Investors are increasingly using factors rather than assets to search for higher returns, to control risk, to gain insights about performance, and to construct portfolios. We acknowledge that factors may offer risk premiums or even alphas as well as valuable insights about risk exposures and performance. But we do not agree that investors should use factors instead of assets as the building blocks for forming portfolios. We have shown that given the same constraints, it is impossible to generate a superior in-sample portfolio by regrouping assets into factors if the investable units are assets from which the factors are formed.

We also argued from first principles that reducing the dimensionality of a larger set of assets to a smaller set of factors is no more effective in decreasing noise than is reducing the dimensionality to a smaller set of assets.

We then 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. We did note, though, that investors who favor predicting factors face the additional challenge of mapping these factor predictions onto asset predictions, and they must also incur incremental trading costs to the extent that factor-mimicking portfolios change over time.

Finally, we identified the types of errors that contribute to covariance instability, and we quantified these errors for various sets of assets and factors. We found that factors are less stable than assets mainly because, unlike assets, they are subject to mapping error.[21] In our view, the case is yet to be made that investors should use factors rather than assets as building blocks for forming portfolios.

ENDNOTES

We thank Ken Froot, Robin Greenwood, Roy Henrikkson, Will Kinlaw, Paul Mende, Anya Obizhaeva, and participants at the London Quantitative Group Summer 2015 seminar and the JOIM/MIT Center for Finance and Policy Fall 2015 seminar for helpful comments.

The material presented is for informational purposes only. The views expressed in this material are the views of the authors and are subject to change based on market and other conditions and factors; moreover, they do not necessarily represent the official views of MIT, Windham Capital Management, State Street Global Exchange, or State Street Corporation and its affiliates.

[1]Of course, if factors are replicated dynamically they would expand the opportunity set. However, one could rebalance assets dynamically for purposes other than to replicate factors, such as to produce option-like payoffs, which would also expand the opportunity set. It is the dynamic rebalancing that changes the opportunity set, not the replication of factors.

[2]Another justification for using principal components for this illustration is that principal component factors will pick up on the information contained in explicit factors that have the ability to explain returns, as shown by Brown [1989].

[3]For a mathematical proof of this equality, please refer to Idzorek and Kowara [2013].

[4]As a simple and intuitive example, consider small-sample error for a universe of assets that is normally distributed. It can be shown that the expected normalized small-sample error of asset standard deviations (from Equation 1) will converge to $1/\sqrt{2k}$ for large samples. Portfolios of normally distributed assets are themselves normally distributed, which means that every portfolio will have the same expected instability per unit of volatility. The normalized errors in covariances between two assets will be distributed similarly to errors in correlation statistics, because much of the volatility contained in the numerator of Equation 1 will be cancelled out by volatility contained in the denominator. For normally distributed assets, it can be shown that expected small-sample error in correlation will converge to $\sqrt{1-\rho^2}/\sqrt{k}$ for large samples. Interestingly, this implies that assets with zero correlations will experience the greatest amount of small-sample error, while those with high correlation will experience the lowest amount of small-sample error.

[5]Investors typically estimate expected returns based on considerations other than purely historical precedent.

[6]It is important to clarify that each of these represents a distinct source of error.

[7]We refer to the components of these formulas as assets, because even when we apply them to factors, the factors are represented by combinations of assets.

[8]Our choice of a 36-month performance horizon aligns with the common practice among many investors of reevaluating their strategic asset allocation every three years. We find that this assumption has little impact on results. We found that our results and conclusions remain intact for 60-month estimation and realization intervals. The only material difference is that, as expected, the degree of small-sample error decreases as the window size grows longer. Mapping error remained almost identical across assets classes, fundamental factors, and principal components.

[9]Mapping error may apply to assets if, for example, the constituents of a stock index change, or an individual security is reclassified within an index. These changes tend to happen rarely and only affect a small portion of the value of any given index, so mapping error for assets is usually close to zero.

[10]For more information, please refer to Kinlaw, Kritzman, and Turkington [2014] and Kinlaw, Kritzman, and Turkington [2015].

[11]Using Equations 4 and 5, we imply annual volatilities and correlations using monthly returns and 11 lagged auto- and cross correlations (where relevant).

[12]It is conceivable that the errors in covariances are correlated across multiple asset pairs. In particular, assets that themselves are correlated may have correlated sources of estimation error. For the data set used in our empirical study, the correlations among errors were consistently close to zero

across the covariance matrixes we analyzed, with only a small number of exceptions. For this reason, as well as the fact that an investor's perception of correlated errors could be either good or bad depending on the portfolio allocation, we use a simple equal weighting of errors across the covariance matrix.

[13]The sources of error we compute are conceptually distinct and will be uncorrelated when computed for a sufficiently long time series of simulated normally distributed data. For the data set used in our empirical study, the correlations among different sources of error were consistently close to zero.

[14]It is tempting to try to measure total estimation error directly by comparing one set of estimates that includes all errors to another set of true realizations. Unfortunately, this is not possible because doing so obscures the mapping error of factors that must be captured. As an example, consider principal components analysis. Suppose our investment objective is to allocate across the principal component factors that will prevail in the subsequent three-year realization sample. We do not know the future covariances, nor do we know the weights that will best track the desired factor exposures over that future period. If we estimate covariances and weights based on historical data, both will have error in tracking the desired exposures, but the off-diagonal covariances will be zero by definition (because they are orthogonal principal components) for both the fitted estimation sample and the re-fitted realization sample, and these correlations would then appear to have zero error. That conclusion is not correct because it ignores the fact that the estimation weights represent different factors than the desired future factor weights.

[15]The cumulative compounded return over any period will always equal the geometric average returns of the subperiods multiplied by the number of subperiods.

[16]We use log returns throughout our analysis to control for the effect of compounding on interval error. At the monthly frequency, log returns are nearly identical to holding period returns and the distinction would not affect our analysis.

[17]The Global Industry Classification Standard (GICS) was developed by and is the exclusive property of MSCI Inc. and Standard & Poor's. GICS is a service mark of MSCI and S&P and has been licensed for use by State Street.

[18]Because we run the principal components analysis on data subsamples of 36-months, the rank of the covariance matrix is only 36, so we include as many eigenvectors as possible in the case of 56 groupings.

[19]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 historically investors have typically allocated across industries rather than across security attributes.

[20]We also tested reducing dimensionality in the asset class universe. In this case, we consolidated assets into three major groups: equities, fixed income, and alternatives. We mapped the original six asset classes to three fundamental factors representing composite macro, fixed income, and equity factors, and we also mapped the six asset classes to the top three principal components. As with individual stocks, we found that dimensionality reduction in the asset class universe did not alter our conclusions. The composite instability of the three fundamental factors and three principal components was larger than that of the three asset classes, largely because of mapping error of the factors.

[21]Our analysis pertains to factors that are fixed-weight linear combinations of the underlying assets. It could be the case that dynamically replicated factors might be more stable, especially if they are rebalanced very frequently. But frequent rebalancing would drive up trading costs and likely offset the benefit of stabilizing factor covariances.

REFERENCES

Brown, S.J. "The Number of Factors in Security Returns." *The Journal of Finance*, Vol. 44, No. 5 (1989), pp. 1247-1262.

Idzorek, T., and M. Kowara. "Factor-Based Asset Allocation vs. Asset-Class-Based Asset Allocation." *Financial Analysts Journal*, Vol. 69, No. 3 (2013), pp. 19-29.

Kinlaw, W., M. Kritzman, and D. Turkington. "The Divergence of High- and Low-Frequency Estimation: Causes and Consequences." *The Journal of Portfolio Management*, Vol. 40, No. 5 (2014), pp. 156-168.

——. "The Divergence of High- and Low-Frequency Estimation: Implications for Performance Measurement." *The Journal of Portfolio Management*, Vol. 41, No. 3 (2015), pp. 14-21.