Adjusted Factor-Based Performance Attribution

Robert A. Stubbs and Vishv Jeet

Robert A. Stubbs
is the vice president of strategic innovation at Axioma, Inc. in Atlanta, GA.
rstubbs@axioma.com

Vishv Jeet
is a senior associate of research at Axioma, Inc. in Atlanta, GA.
vjeet@axioma.com

Factor-based performance attribution is commonly used to explain the sources of realized return of a portfolio. The methodology relies on a factor model of asset returns to decompose a portfolio's return according to a set of factors. The portion of the portfolio return that can be explained by the model factors is called the *factor contribution*, and the remainder is called the *asset-specific contribution*, or *specific contribution* for short. (For a description of factor-based performance attribution, see Fischer and Wermers [2013], chapter 4.)

Unfortunately, the inferences from a standard attribution report can be misleading for several reasons—one of which is the misclassification of factor contributions as asset-specific contributions or vice-versa. Aside from missing factors, this misclassification can also be due to biased factor exposure estimates. As a result, inferences on the statistical significance of the contributions may also be incorrect. With the trend toward "smart beta" and factor investing in general, the ability to draw correct inferences about factor contributions from attribution reports is increasingly critical. In this article, we will address the causes of erroneous attribution analysis and propose a methodology that produces better representations of reality from which more accurate inferences can be drawn.

Before we jump into the causes of inaccurate inferences and our proposed solution, we will illustrate the types of inaccurate inferences that can be made from a standard attribution on a particular example. We constructed a portfolio that is rebalanced monthly from January 1995 to October 2013 according to the following strategy:

maximize Expected Return
subject to: Long Only and Fully Invested
        Active Risk Constraint 3%   (*Strategy*)
        Active Sector Bounds of ± 4%
        Active Asset Bounds of ± 3%

We used exposure to a growth factor as the expected return and the Russell 1000 Index as the benchmark. We then considered two different returns models to use in attributing returns of this portfolio. The first model, RM1, uses 10 sector factors and 4 style factors—market sensitivity, momentum, size, and value. The second model, RM2, adds the exact growth factor used to construct the portfolio to those factors present in RM1.[1] The active risk constraint used a factor risk model based on the RM2 returns model.[2]

Exhibit 1 summarizes the attribution results for the active portfolio using these two models. All contributions are annualized values computed using the linking methodology of Cariño [1999]. When we use RM1

Summary of Performance Attribution Results

Returns Model	RM1	RM2
Active Return	1.47%	1.47%
Specific Contribution (SC)	1.65%	−0.88%
Factor Contribution (FC)	−0.18%	2.35%
SC *T*-Stat	2.67	−1.58
FC *T*-Stat	−0.69	4.29
SC-FC Correlation	−0.09	−0.32

Notes: The values in the top section are the annualized active returns and their decompositions. In the bottom section, we give additional summary statistics for the factor contribution (FC) and specific contribution (SC).
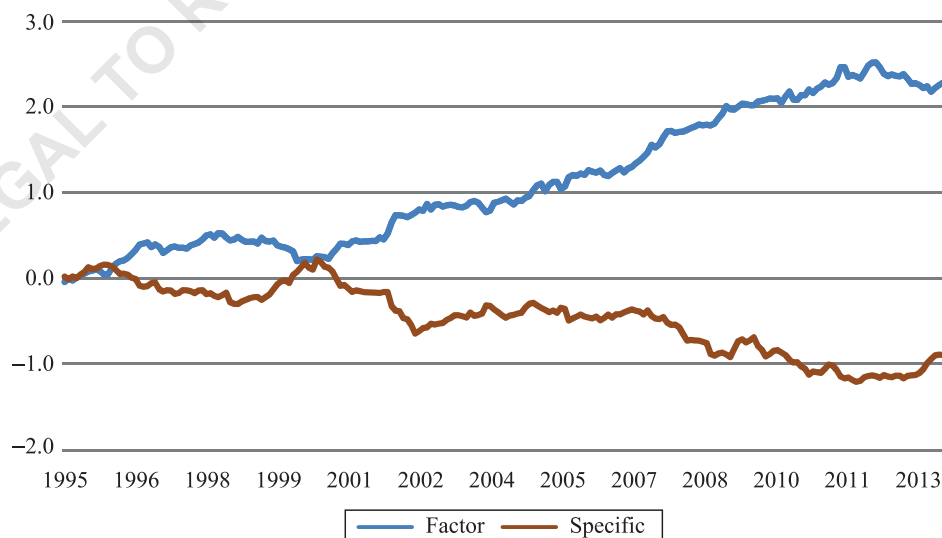
to perform the attribution, the specific contribution explains most of the active return; we expect this because we are betting on a factor that is not in the returns model. When we use RM2 to perform the attribution, the overall annualized factor contribution increases from −0.18% to 2.35%—which is primarily due to the contribution of the growth factor in RM2. Because our portfolio maximized exposure to this factor, this seems to be exactly what we want. However, note that the asset-specific contribution decreased commensurately from 1.65% to −0.88%. Moreover, the *t*-statistic on the specific contribution changed from being significantly positive (2.67) to almost significantly negative (−1.58).

What is the correct inference to draw from this analysis? Unfortunately, the truth probably lies somewhere between the two attributions. The attribution using RM1 misses the contribution from the growth factor, whereas the attribution using the RM2 model seems to overstate the contribution from factors. Exhibit 2 illustrates this overstatement: The cumulative factor and specific contributions are moving in opposite directions throughout the backtest period, suggesting that the contributions are negatively correlated. In fact, the correlation between the factor and specific contributions over the entire backtest is −0.32. Thus, for this particular portfolio, our "specific" contribution is actually related to our factors. As we will see later, an exaggerated exposure to the growth factor causes this dependence of the "specific" contribution on the factors, and this biased factor exposure estimate leads to incorrect inferences from the attribution results.

The problem of under- or overattributing portfolio returns to factors is not unique to this particular example; it can manifest itself in varying degrees in nearly all combinations of portfolios and returns models. As we see from this example, misattributing portfolio returns manifests itself in correlated factor and specific contributions and leads to inaccurate inferences. In the remainder of this article, we will discuss the cause of this misattribution and propose an adjusted

E X H I B I T 2

Cumulative Factor and Asset-Specific Contributions to Portfolio Return Using the RM2 Returns Model for Attribution

performance-attribution methodology that more accurately attributes portfolio returns to the factors under consideration. The resulting adjusted factor and specific contributions are uncorrelated, leading to more accurate inferences from the attribution report. Using the motivating example shown earlier, this article will introduce methodology that reduces the correlation between factor and specific contributions to near zero and drives the annualized specific contribution to near zero, while the contribution from the growth factor remains large and highly significant.

SOURCE OF THE PROBLEM

There are a couple of different ways to think about factor-based attribution. The most common approach is to start with a linear model of excess asset returns, r, in period t that is written as

$$r = Xf + \varepsilon \qquad (1)$$

where X is an $n \times k$ factor exposure matrix, n is the number of assets, k is the number of factors, f is a vector of factor returns, and ε represents the asset-level residual returns. In this article, we will assume that the returns model is cross-sectional, where X is given and the factor returns are estimated using a weighted least squares regression. The factor returns are given as $f = H^T r$, where

$$H = WX(X^T WX)^{-1} \qquad (2)$$

is a matrix whose columns are the pure factor-mimicking portfolios (FMPs) associated with X and regression weights, W. That is, the factor returns are the returns of the set of FMPs associated with the returns model. Taking a portfolio-weighted sum of the returns model, we can explain the returns of an (active) portfolio, h, according to

$$h^T r = \underbrace{h^T Xf}_{\text{factor contribution}} + \underbrace{h^T \varepsilon}_{\text{specific contribution}} \qquad (3)$$

The second approach explains h with a set of factor portfolios—specifically, the FMPs given by H. We can compute the exposures, λ, to these factor portfolios using the regression

$$h = H\lambda + u \qquad (4)$$

We can then take a return-weighted sum of Equation 4 to decompose the return of the portfolio as

$$r^T h = r^T H\lambda + r^T u \qquad (5)$$

It is straightforward to show that Equation 5 is equivalent to Equation 3 if $\lambda = X^T h$—that is, if the estimated factor exposures are equivalent to those classically computed in Equation 3. The advantage of this second way of thinking about attribution is that we can see that exposures are not exact: They are least-squares estimates of a linear regression. And as with all regressions, the estimates contain errors and may be biased if all underlying model assumptions are not satisfied.

This second approach is in the spirit of Grinold [2006, 2011], who explains the portfolio with a set of portfolios. However, in this article, we do not try to customize an attribution by mapping user factors to those in a particular model. Instead, we propose a method for adjusting attributions to correct for biases implicitly present in standard attribution approaches. The methodology works independently of the particular attribution methodology or factor model of returns.

Now, we will consider the conditions under which it makes sense to use $X^T h$ as the factor exposure estimates. To explain much of the portfolio return with our factors, we want to minimize the variance of the unexplained portfolio, u. We would ideally minimize $u^T Qu$ when estimating λ, where

$$Q = X\Omega X^T + \Delta^2$$

is the estimated asset covariance matrix, $\Omega = \text{cov}(f)$ is the factor covariance matrix, and Δ^2 is the diagonal matrix of asset-specific variances. If the FMPs in H are constructed with $W = \Delta^{-2}$ and we use Q as the generalized least squares (GLS) weights in Equation 4, then $\lambda = X^T h$. That is, with this particular choice of FMPs and regression weights, the factor exposure estimates are equal to the standard factor exposures.

However, FMPs are not always constructed (even implicitly) using weights equal to the inverse of specific variances. If regression weights equal to market caps or square-roots of market caps are used in the returns model, and thus in the FMPs, then $X^T h$ is the weighted least squares (WLS) solution to Equation 4 if the specific variance estimate in Q is the inverse of those regression weights.[3] Using the inverse of asset market capitalizations

as estimates of specific variances seems incorrect, but this is implicitly what we are doing when we take the factor exposures as a given when the returns model is estimated with market-cap weights.[4] Using such estimates of specific variances can produce factor and residual components of the portfolio whose returns are correlated ex ante, let alone ex post.

To see the extent of the correlations in our numerical example, we plot the estimated correlation between the fitted (factor) and residual components of the portfolio through time in Exhibit 3, where H is computed with the market-cap weights in RM2.

To see why this is a problem, consider a portfolio that has unit exposure to a single factor but has exactly half the return of the factor in each period. The return of the portfolio, r_p, can be decomposed as

$$ r_p = f + (-0.5 f) $$

where f is the factor contribution and $-0.5f$ is the specific contribution. The factor and specific contributions of the portfolio have a correlation of -1 that is caused by the portfolio having an exposure that is twice the correct value.

One of the key assumptions of linear regression is that the conditional mean of u_i, given H, is zero—that is, $E[u_i|H] = 0$ (see Greene [2003], chapter 2). This says that the expected value of each position in the unexplained portfolio should be independent of H. Violation of this assumption leads to biases in the estimates of λ, the factor exposures. If $\text{cov}(u, H) \neq 0$, then the conditional mean cannot be zero, thus violating this assumption.[5]

SOLUTION: ADJUSTED ATTRIBUTION

To eliminate the correlation between the fitted portfolio and residual portfolio, we model the residual portfolio as a function of the factor portfolios. Specifically, we assume that the residual portfolio is a linear function of the individual FMPs, rather than of the fitted portfolio, in order to provide more flexibility in our residual portfolio model. We consider the following model of the residual portfolio u:

$$ u = \sum_j H_j \beta_j + \tilde{u} \qquad (6) $$

Taking the inner product with asset returns, we can model the residual return of the portfolio in period t as

$$ r_t^T u_t = \sum_j r_t^T H_{tj} \beta_j + r_t^T \tilde{u}_t = \sum_j f_{tj} \beta_j + r_t^T \tilde{u}_t \qquad (7) $$

Estimating β from the cross-sectional regression in Equation 6 requires knowing H, which is not generally

E X H I B I T 3
Estimated Correlations of Fitted and Residual Components of the Portfolio Using a Factor Risk Model Based on the RM2 Returns Model
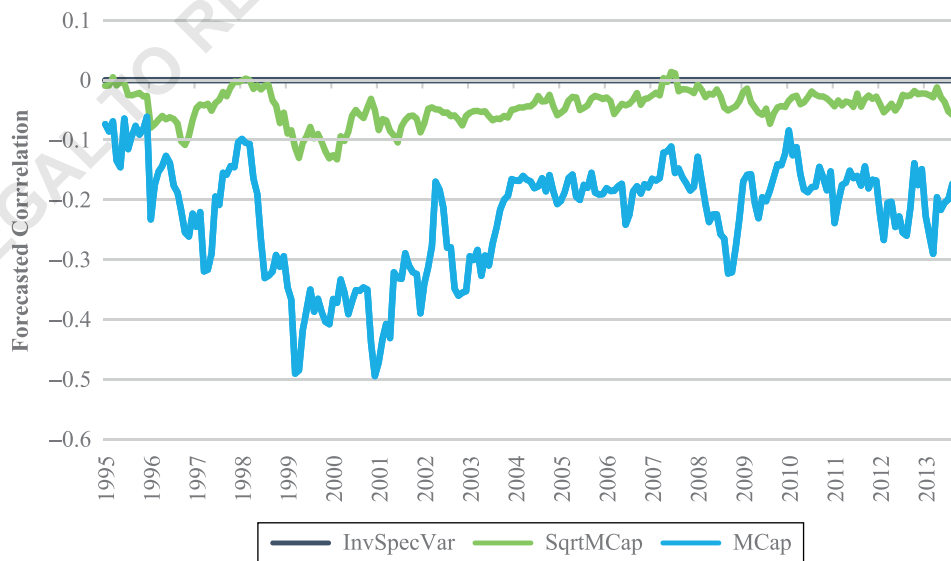
available and may suffer from the biases mentioned even if it were. Instead, we propose to estimate β using the time-series regression in Equation 7.[6] This also has the benefit of using a single correction through time, as opposed to modifying the factor exposures differently in each period. We refer to this adjustment to the factor exposures as an *absolute* adjustment. After estimating β via a time-series regression, dropping the time subscript, and substituting Equation 6 into Equation 5, we can decompose the period returns of the portfolio as

$$r^T h = \sum_j r^T H_j X_j^T h + \sum_j r^T H_j \beta_j + r^T \tilde{u}$$
$$= \sum_j f_j (X_j^T h + \beta_j) + r^T \tilde{u} \qquad (8)$$

In our experience, exposures are typically off by a relative amount, rather than an absolute amount. Therefore, we propose the following alternative to Equation 8 that we refer to as the relative adjustment:

$$r^T h = \sum_j f_j X_j^T h (1 + \beta_j) + r^T \tilde{u} \qquad (9)$$

The β values for this model are estimated from the following alternative to Equation 7, which uses factor contributions, as opposed to factor returns, as the independent variables:

$$r_t^T u_t = \sum_j f_{tj} X_{tj}^T h_t \beta_j + r_t^T \tilde{u}_t \qquad (10)$$

A relative adjustment can also be more appropriate if factor exposures are changing through time. For these reasons, we prefer the relative adjustment to the absolute adjustment and use it in all computational results. Nevertheless, an absolute adjustment may work better in some situations, and the remainder of this article is relevant to either form of adjusted attribution.

Up to this point, we have looked primarily at motivational examples and solutions. Before looking at results on realistic portfolios, we consider the issue of potentially overfitting the adjustment regression. Overfitting would allow factors to account for some of the noise in the portfolio returns, rather than accounting for only factor contributions in the portfolio's residual returns. To avoid this, we use a robust procedure to estimate the time-series model in Equation 10. Using

contributions as opposed to factor returns (relative adjustment versus absolute adjustment) has advantages with regard to this issue, in addition to previously mentioned benefits. Because we are making relative adjustments to the exposures, the adjustment procedure will not suddenly allow a factor to explain a large portion of returns when the unadjusted factor exposure is near zero. If the exposure was near zero prior to adjustment, it will remain near zero after the adjustment. In this sense, the proposed adjusted attribution methodology behaves like a Bayesian method with the standard exposures as the prior.

It might be tempting to simply estimate the factor exposures from the time series in Equation 8—similar to Sharpe's style analysis (see Sharpe [1988]). This works adequately with a small number of factors relative to the number of periods in the time series. However, the typical number of factors in an equity factor risk model is large—possibly larger than the number of periods—which leads to significant errors in the estimates. Factor-based attribution uses cross-sectional information from the holdings and factor exposures to compute contemporaneous factor exposures. However, as we illustrated, these estimates can be biased. The proposed approach is a hybrid approach that starts with cross-sectional estimates and then corrects them if there is a systematic bias through time.

To avoid the potential issues we described in the time-series regression, we use a variable selection scheme to select a reduced set of factors. Rather than use an OLS estimate of β, we use only a set of statistically significant β values to adjust the factor and specific contributions of the portfolio. The β values for all other factors are set to zero.

For all numerical results presented here, we use a heuristic variable selection scheme to select the independent variables (factor contributions) of Equation 10 based on their statistical significance, as measured by their p-values. We use an iterative regression scheme that starts with all variables present. After each iteration, we remove the variable with the greatest p-value if it is greater than the specified tolerance 0.02. If none of the p-values exceed the tolerance, we stop the iterative procedure of removing factors. Thereafter, we employ a reentry procedure in which we consider reentering rejected variables into the regression one at a time. A variable can reenter the regression only if its entry does not increase the p-value of any variable (including itself)

above the tolerance. After the reentry trials, we run a final regression with the selected variables to compute the final estimate of β. The resulting β is then used to adjust the factor contributions as prescribed by Equation 9.

EXAMPLES OF ADJUSTED ATTRIBUTION

We will now look at results for many different portfolios exhibiting varying degrees of correlated factor and specific contributions when using standard factor-based attribution. In all computational results shown here, we restrict potential adjustments to style factors only. For the strategies used in our tests, these are the only factors likely to have a large contribution to returns because of the constraints in the portfolio construction strategy.[7] We use the relative adjustment regression shown in Equation 10 in all numerical results.

First, we will return to the motivating example shown in the introduction and look at the adjusted attribution results when RM2 is used to explain returns. The R-squared for the adjustment regression was 0.17, and the statistics for the β estimates are shown in Exhibit 4. Two factors ended up in the final list of statistically significant factors: value and growth. The β was −0.65 for value and −0.39 for growth, which means that our exposures to value and growth were too large and should be only 35% and 61% of their original values, respectively. The average active exposures before and after adjustment are shown in Exhibit 5. In Exhibits 5 and 6, the shaded cells indicate those that notably changed from the standard attribution.

In Exhibit 6, we compare the attribution results using adjusted performance attribution (PA) to those results originally shown in Exhibit 1 that used standard PA methodology. Comparing the last two columns of Exhibit 6, we see that the correlation between the adjusted factor and specific contributions changed from −0.32 to 0.09. Furthermore, the annualized factor contribution decreased from 2.35% to 1.50%, and the annualized specific contribution increased from −0.88% to

E X H I B I T 4
Beta Coefficients and Their Significance as Determined by the Time-Series Regression

	Coefficient	Standard Error	*T*-Stat	*P*-Value
Value	−0.65	0.21	−3.07	0.00
Growth	−0.39	0.06	−6.39	0.00

E X H I B I T 5
Active Portfolio Exposures to Style Factors Using Standard and Adjusted Attribution

	Active Exposures	
Attribution	**Standard**	**Adjusted**
Growth	1.010	0.616
Market Sensitivity	0.012	0.012
Momentum	0.017	0.017
Size	−0.114	−0.114
Value	−0.205	−0.072

E X H I B I T 6
Summary of Performance Attribution Results

Returns Model	RM1	RM2	RM2 (Adjusted PA)
Portfolio Return	10.72%	10.72%	10.72%
Benchmark Return	9.25%	9.25%	9.25%
Active Return	1.47%	1.47%	1.47%
Specific Contribution (SC)	1.65%	−0.88%	−0.03%
Factor Contribution (FC)	−0.18%	2.35%	1.50%
Industry	−0.14%	−0.15%	−0.15%
Styles	−0.04%	2.50%	1.65%
Growth	–	2.65%	1.60%
Market Sensitivity	−0.10%	−0.09%	−0.09%
Momentum	0.10%	0.09%	0.09%
Size	0.13%	0.15%	0.15%
Value	−0.17%	−0.29%	−0.10%
FC Volatility	1.13%	2.37%	1.53%
SC Volatility	2.66%	2.41%	2.21%
SC *T*–Stat	2.67	−1.58	−0.06
FC *T*–Stat	−0.69	4.29	4.22
SC–FC Correlation	−0.09	−0.32	0.09

Note: The values are the annualized returns attributable to each factor or to a group of factors.

−0.03%. In the proposed methodology, we consider the correlation between period factor and asset-specific contributions through time. It is possible to have a significant annualized specific contribution even if the period returns are Gaussian noise. Thus, although the adjusted attribution significantly reduced the annualized specific contribution in our example, it is not necessary for this to be the case in general. Our goal is to eliminate the relationship between factor and specific contributions.

To further validate the proposed adjusted PA methodology, we ran a variety of backtests and used different returns models with and without the adjustment in the attributions of the portfolios. We started with 12

backtests (combinations of three benchmarks and four expected returns) using U.S. market data. Our backtests used the Russell 1000, Russell 2000, and Russell 3000 indexes as benchmarks, as well as the investable universes. We used growth, value, momentum, and an equal-weighted combination of value and momentum as four different expected returns in our backtests. The optimization strategy was described in Equation 1. All backtests started with cash in January 1995 and were rebalanced monthly through October 2013. All attributions used the RM2 returns model that includes the growth factor. In our computational results, we show the results for two factor-based attributions: the standard factor-based attribution methodology labeled as *prior* and the adjusted factor-based attribution using the same model labeled as *adjusted*.

In this first set of tests, all 12 cases fall under the scenario in which the alpha factors used as expected returns in the construction of the strategy are also present in the returns model used to perform the attribution. As shown in Exhibit 7, all 12 cases have very large negative correlations between the factor and specific contributions. For each of these cases, the adjusted attribution reduced this correlation to near zero. If we had used an ordinary least squares (OLS) solution of Equation 10

with all factors, the adjusted correlation would be zero. However, using only the set of significant factors is more intuitive and nearly eliminated the issue.

The annualized contributions for the prior and adjusted cases using the Russell 1000 Index as the benchmark are shown in Exhibit 8. Notice that most of the factor contribution adjustments took place in the alpha factors and the specific contributions were adjusted accordingly. That is, if value was used as the expected return, then the factor contribution of value was adjusted downwards and the specific contribution was adjusted by the opposite amount. In Exhibit 8, Panel D, in which the expected return is a combination of the value and momentum exposures, we can see that the value contribution was adjusted downward significantly, but the momentum contribution was unchanged. In this case, the *t*-statistic for the value contribution was −10.54, while the *t*-stat for the momentum contribution was only −1.52. With a large *p*-value of 13%, the momentum contribution did not meet our significance threshold of 2%, thus the momentum contribution was not adjusted.

Exhibit 9 shows the split of total active returns between factor and specific contributions. In all cases, the annualized specific contribution increased when using adjusted attribution. When using standard attribution,

E X H I B I T 7

Correlation of Factor and Specific Contributions: Returns Model Contains Exactly the Same Alpha Factors Used in Expected Returns
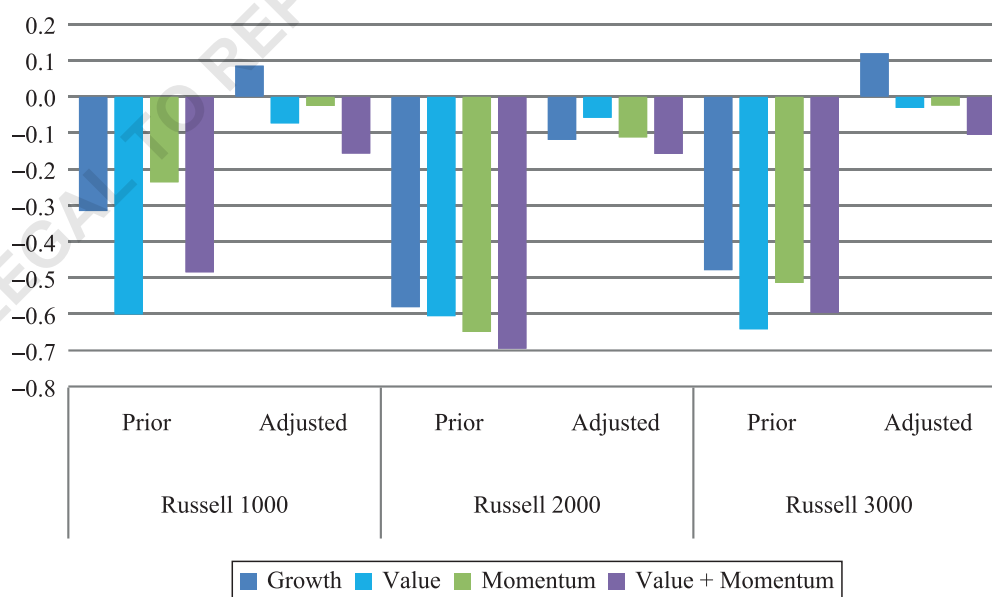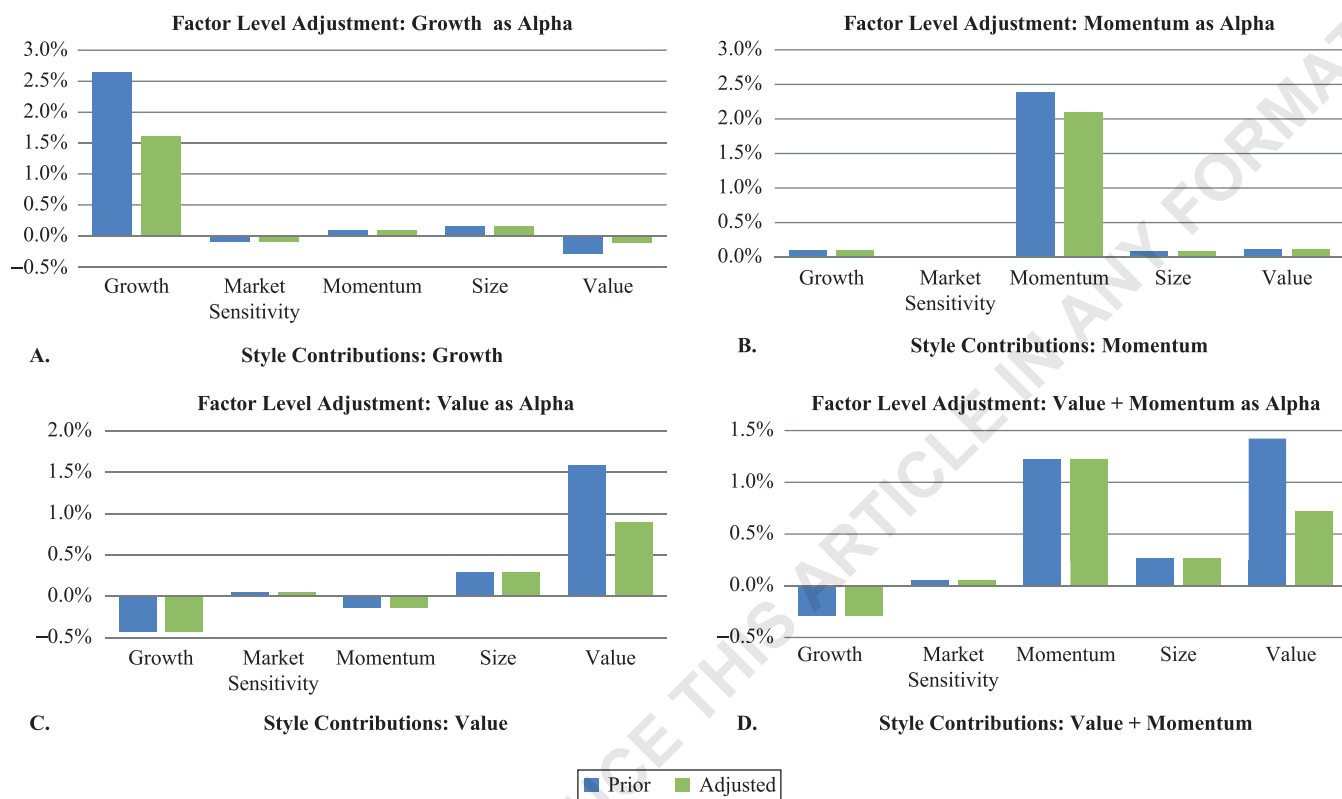
Style Contributions with and without Adjustment: Returns Model Contains Exactly the Same Alpha Factors Used in Expected Returns

Factor Level Adjustment: Growth as Alpha

A. Style Contributions: Growth

Factor Level Adjustment: Momentum as Alpha

B. Style Contributions: Momentum

Factor Level Adjustment: Value as Alpha

C. Style Contributions: Value

Factor Level Adjustment: Value + Momentum as Alpha

D. Style Contributions: Value + Momentum

Prior   Adjusted

Note: These results are for the case in which we used the Russell 1000 Index as the benchmark and investable universe.

the exposures to the alpha factors were overestimated, thus causing a downward-biased specific contribution.

Thus far, we have only discussed attribution, but the decomposition of realized risk will also change when using adjusted attribution; namely, the absolute volatility of the factor contributions may decrease because of a reduction in exposures. This can also have a significant effect on the statistical significance of the factor and/or specific contributions. In these 12 backtests, the average volatility of the factor contributions decreased from 3.0% to 1.8%, and the average volatility of the specific contributions decreased from 2.5% to 2.0%. We were able to reduce the volatility of each of the contributions while the overall volatility remained the same—precisely because we eliminated the correlation between the contributions.

Next, we consider the case in which the returns model does not contain the exact definition of alpha used to construct the portfolio, but it does contain a related factor. We ran three backtests under this scenario with estimated earnings-to-price ratio (E/P) as the alpha factor and the same three Russell indexes used in prior experiments as the benchmarks. The same strategy and risk model were used to construct the portfolios. In Exhibit 10, we can see that the correlation between factor and specific contributions prior to adjustment is positive in each of the three cases. This is the opposite of what we saw in the case in which the returns model used for attribution contained the same alpha factor. Although the adjustment regression is able to explain a portion of specific contributions, it is not able to explain the portion that could potentially be explained by estimated E/P, as opposed to book-to-price ratio (B/P) that is in the returns model. In Exhibit 11, we can see that in each case, the specific contribution was positive prior to adjustment. Again, this is the opposite of what we observed when the returns model contained the same factors used in expected returns.

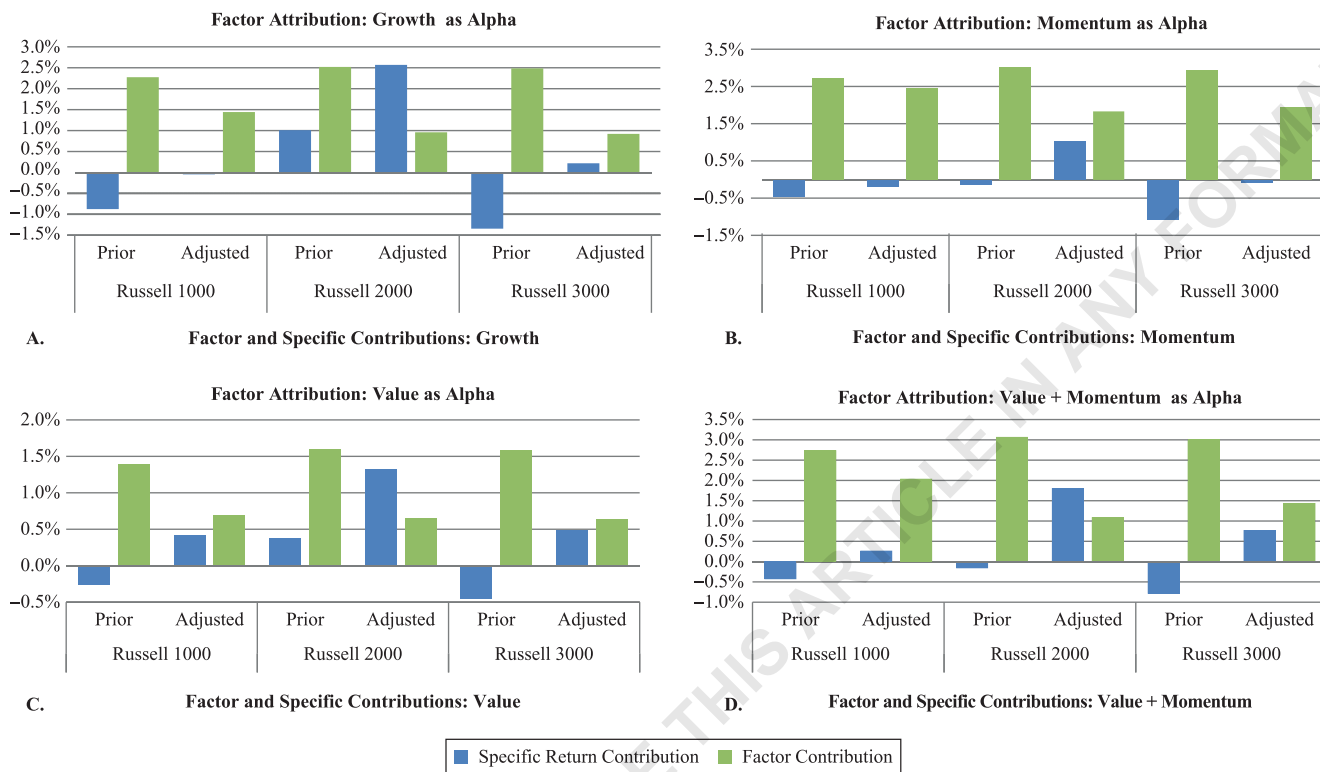**Performance Attribution: Returns Model Contains Exactly the Same Alpha Factors Used in Expected Returns**

Factor Attribution: Growth as Alpha

A.   **Factor and Specific Contributions: Growth**

Factor Attribution: Momentum as Alpha

B.   **Factor and Specific Contributions: Momentum**

Factor Attribution: Value as Alpha

C.   **Factor and Specific Contributions: Value**

Factor Attribution: Value + Momentum as Alpha

D.   **Factor and Specific Contributions: Value + Momentum**

■ Specific Return Contribution   ■ Factor Contribution

E X H I B I T   1 0

Correlation of Factor and Specific Contributions: Returns Model Contains Factors Related to the Alpha Factors
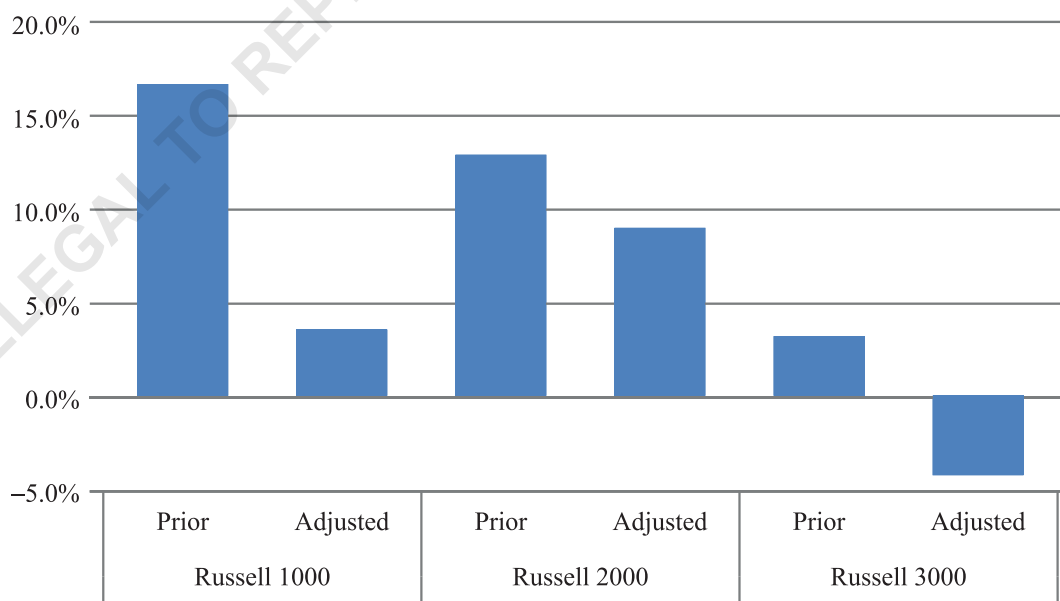
EXHIBIT 11
Factor and Specific Contributions: Returns Model Contains Factors Related to the Alpha Factors
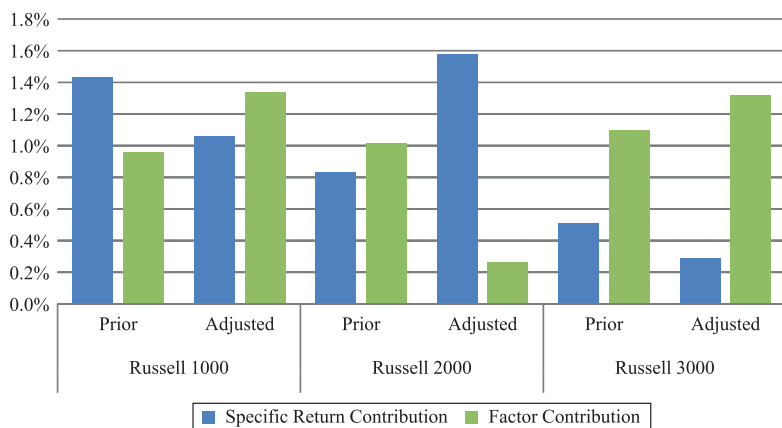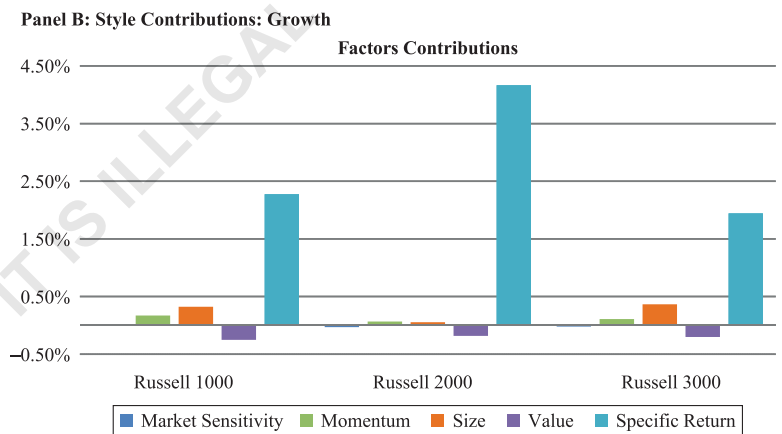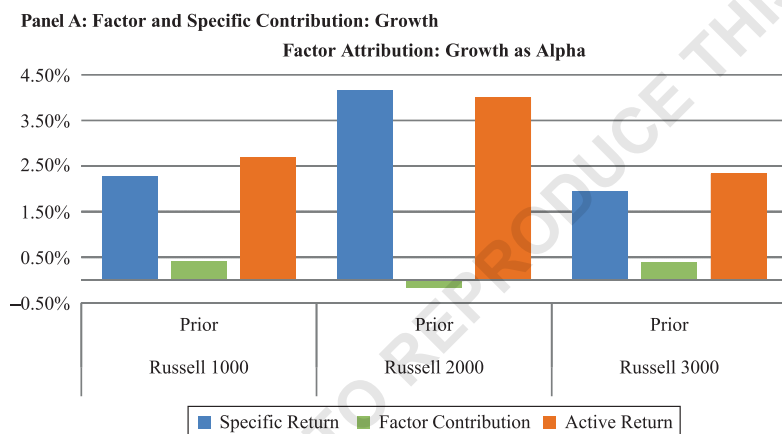

EXHIBIT 12
Performance Attribution: Returns Model Does Not Contain Exactly the Same or Similar Factors to Those Used in Expected Returns

Panel A: Factor and Specific Contribution: Growth

Panel B: Style Contributions: Growth

As we have shown, if we use factor-based attribution with a returns model containing factors that are either exactly the same as or similar to the factors used to construct the portfolio, adjusted attribution is needed to correct for the correlation between factor and specific contributions. One might wonder whether omitting all factors that are similar to the alpha factor(s) would eliminate this problem and produce better attribution results. In this last set of tests, we consider such a case in which the alpha factor and similar factors are not in the returns model.

We ran three more backtests in which we used the growth factor as our expected returns and a risk model based on the RM1 returns model (that did not contain the growth factor). This is the same scenario used in our motivating example in the introduction. The results are shown in Exhibit 12, Panel A. Because no factor related to the alpha factor is present in the returns model, most of the active return is attributed to asset-specific bets. In Exhibit 12, Panel B, contributions from the individual styles are plotted alongside the specific contribution to further show that factor contributions to each factor were small. So, although the issues related to correlated factor and specific contributions may not be present, the attribution does not show the true source of skill. And as we saw in the introduction, the specific contribution in such a scenario may not be statistically significant—whereas a factor contribution would be.

To illustrate this point, we look at one more case in which the alpha factor is completely missing from the returns model; Exhibit 13 summarizes the results. We used estimated E/P as the expected returns, and the returns model had only four styles: growth, market sensitivity, momentum, and size. In the "Prior" column, one can see that the contributions of the alpha factor are attributed to specific contributions. No adjustment regression is run in this case.

Even though alpha factor contributions are mixed with genuine specific contributions, there is no detectable skill in the specific contributions because the *t*-stat of specific contribution is insignificant. In the "Adjusted" column, the

Exhibit 13
Adjustment Regression with Missing Factor in the Returns Model

	Prior	Adjusted	Directional Change
Factor Contributions (FC)	0.53%	1.32%	Up
Specific Contributions (SC)	1.09%	0.29%	Down
Total Active Return	1.61%	1.61%	Same
T-stat (SC)	1.30	0.38	Stays Insignificant
T-stat (FC)	1.41	2.43	Becomes Significant

attribution was run with a returns model that contained the alpha factor and used the adjusted PA methodology. After the alpha factor's contributions are identified, the specific contributions are reduced appropriately and their *t*-stats stay insignificant. However, the *t*-stat for the factor contributions improved significantly.

CONCLUSIONS

Explaining portfolio returns with the factors driving the portfolio is generally beneficial in supporting the story of any portfolio based on factor investing. Perhaps more importantly, it can show that the main drivers of return are statistically significant—providing confidence that the strategy will outperform in the future. However, many active portfolios are not that similar to the pure factors used in attribution to explain them. They contain "noise" due to constraints and other frictions that will show up as a specific contribution (see Clarke, de Silva, and Thorley [2005]). If a would-be contribution from a factor missing from the returns model is mixed with the specific contribution, the statistical significance of the mixed contribution will generally be less than that of the would-be factor contribution. In other words, attributing the return to a factor that is truly driving the returns of the portfolio allows us to separate the signal from the noise. However, care must be taken when estimating factor exposures, particularly for those factors that the portfolio is intentionally tilting toward. Biased estimates can lead to incorrect inferences from an attribution analysis.

Having correlated factor and asset-specific portfolio contributions violates one of the basic assumptions of a factor model of asset returns and thus introduces error into factor-based performance attribution results. This error can be found to some extent when performing factor-based attribution with nearly every real portfolio and any returns model. We proposed a solution of adjusting the factor-based performance attribution methodology to account for the correlation between the factor and asset-specific contributions of return that are computed with standard factor-based PA. In essence, the proposed attribution is a combination of factor-based attribution and style analysis (see Sharpe [1988]). Here, the "styles" are the most significant factor contributions. And, because the same factor contributions are already present in the attribution, we distribute the asset-specific contributions that can be explained by the factors back into contributions to the factors, rather than separately accounting for the explanation of specific contributions.

In our experience, the classical bias/variance trade-off seems to exist in standard attribution results in which variance is the volatility of the unexplained portfolio, and bias is the over- or underestimation of factor contributions. The closer the alpha factors are to the factors in the returns model used for attribution, the lower the variance—but the greater the potential for bias. If the independent variables (FMPs in this case) explain much of the portfolio, then we will see a lower variance. However, if the assumptions of classical regression are not satisfied, the factor exposure estimates (and thus factor contributions) may be biased. If the independent variables do not explain much of the portfolio, the variance will be large and the exposure estimates may have large standard errors, but they are unlikely to be biased. In this article, we proposed a method that provides the best of both worlds: We can start with a low-variance model and correct for bias in the factor exposures by using the time-series regression we refer to as the adjustment regression.

ENDNOTES

[1] The five styles were growth, momentum, market sensitivity, size, and value. Our style factor exposures are defined as follows: Growth is the plowback ratio times return-on-equity, momentum is the cumulative return over the past year excluding the most recent month, market sensitivity is the beta of time-series regression of an asset's return against the market return using six-month daily data, size is the natural logarithm of market capitalization, and finally, value is book-to-price ratio (B/P). The style factors were normalized over the estimation universe by subtracting the market-cap-weighted mean from the raw descriptor and dividing by the equal-weighted deviation from zero. The ten sectors are defined by GICS. The factor returns were estimated using daily cross-sectional regressions weighted using stocks' market capitalizations.

[2]All results associated with this example are almost identical to those based on a portfolio optimized with a risk model based on the RM1 returns model. We include only one set of results here for brevity.

[3]Either $Q = X\Omega X^T + W^{-1}$ or $Q = W^{-1}$ will give $X^T h$ as the GLS solution to the portfolio decomposition regression.

[4]We used a model with market-cap weights only to make certain points more extreme. Such a model is not commonly used.

[5]The converse of this statement is not true. We may have a nonzero conditional mean of the residual portfolio even if the covariance between the residual and fitted portfolio is zero. These violations of the conditional mean assumption can arise for other reasons—such as a missing factor in the returns model or the holdings in the portfolio being a nonlinear function of the FMPs—than the linear relationship we are imposing. These are more structural issues that may or may not be addressed with the proposed methodology.

[6]We do not add a constant term to the adjusted attribution regression. If we did, then the constant term would still go into the adjusted specific contribution. By not including a constant term, we ensure that the factor explains the constant if possible (if it has a nonzero mean). However, we ensure that the correlation between the adjusted factor and specific contributions is zero only if we include a constant term in the regression or if the mean of the adjusted specific contributions is zero.

[7]In order for a factor to have any real impact on the adjusted attribution when using the relative adjustment, the exposure to the factor should be relatively large. The factors that are likely to have large exposures are those that the investment strategy is intentionally tilting. We often find that we can improve the intuition in the ultimate adjusted attribution by pre-selecting a subset of factors to be candidates for adjustment.

REFERENCES

Cariño, D. "Combining Attribution Effects over Time." *Journal of Performance Measurement*, Vol. 3, No. 4 (1999), pp. 5–14.

Clarke, R., H. de Silva, and S. Thorley. "Performance Attribution and the Fundamental Law." *Financial Analysts Journal*, Vol. 61, No. 5 (2005), pp. 70–83.

Fischer, B.R., and R. Wermers. *Performance Evaluation and Attribution of Security Portfolios*, 1st edition. Waltham, MA: Academic Press, 2013.

Greene, W.H. *Econometric Analysis*, 5th edition. Upper Saddle River, NJ: Prentice Hall, 2003.

Grinold, R. "Attribution." *The Journal of Portfolio Management*, Vol. 32, No. 2 (2006), pp. 9–22.

——. "The Description of Portfolios." *The Journal of Portfolio Management*, Vol. 37, No. 2 (2011), pp. 5–30.

Sharpe, W.F. "Determining a Fund's Effective Asset Mix." *Investment Management Review*, November/December 1988, pp. 59–69.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.