

Turning Negative Into Nothing:

AN EXPLANATION OF “ADJUSTED FACTOR-BASED PERFORMANCE ATTRIBUTION”

Factor attribution sits at the heart of understanding the returns of a portfolio and assessing whether a manager has invested in a manner consistent with his value proposition. In this paper¹, we will step back and look at factor-based attribution from first principles, as well as describe a methodology that will help correct some of the underlying issues that may arise and produce misleading results.

One of the biggest problems that many users of attribution have noted is that factor contributions may be misclassified as asset specific or vice versa. This misclassification has mainly plagued users who employ factor-based investment approaches, and expect to show that their factor exposures have led to their investment results, but who don't always find that to be the case. This issue forms the basis of our analysis later in the paper.

What is Attribution, Anyway? We use factor-based attribution to try to explain a portfolio's return using returns on a set of representative portfolios, each of which is simple and well understood. For example, we may want to model a portfolio that is a combination of large-cap, mid-cap and small-cap U.S. stocks using three representative portfolios: the constituents of the S&P 500 Index, S&P 400 Index, and S&P 600 Index. We estimate how much each representative portfolio explains our portfolio by regressing the holdings of our portfolio against the holdings of the representative portfolios. That is, we regress

$$(\text{Our portfolio}) = (\text{S\&P 500 Index}) * \beta_1 + (\text{S\&P 400 Index}) * \beta_2 + (\text{S\&P 600 Index}) * \beta_3 + \epsilon,$$

wherein the β variables are the weights of each of the representative portfolios.

¹ For a much more detailed discussion of this topic, see Stubbs and Jeet, “Adjusted Factor-Based Performance Attribution” *Journal of Portfolio Management*, Vol. 42, No. 5: pp. 67–78.

After breaking down our portfolio into a weighted set of representative portfolios, we explain the returns as the weighted sum of returns on the representative portfolios plus a residual, which is the return unexplained by the factors. The residual may also be termed “stock-specific return.”

If our portfolio is an S&P 1500 Index tracking portfolio, then the previous example works perfectly because we can completely explain our portfolio’s return with those of the three representative portfolios². For every 1% change in the S&P 500 Index, we know that our portfolio will go up by β_1 percent since β_1 is the sensitivity of our portfolio to the S&P 500 Index. The holdings in this example represent *sensitivities* of the stocks in our portfolios to the returns of the indices.

In standard factor-based attribution, we do not explicitly compute sensitivities to the representative portfolios using a regression of holdings, as demonstrated above. Instead, we use *exposures* to the representative portfolios as an approximation³. (We use Factor Mimicking Portfolios, or FMPs, as our set of representative portfolios. These are the portfolios whose returns define factor returns.) Exposures of a representative portfolio to the factor portfolios are calculated as the weighted average of each stock’s exposure to the factor in question⁴, where the weights are the weights in the portfolio. *Exposures serve as a proxy for sensitivities but may contain errors and biases.*

FMPs. What is a factor return? Factor returns are essentially the weighted average return to a long-short minimum variance portfolio that has unit exposure to the factor in question and no exposure to any other factor in the risk model. At Axioma, we call the underlying portfolios “Factor-Mimicking Portfolios,” or simply FMPs.

Ideally, for factor-based investors, the factors would explain as much return as possible, and the residual portion would be small. Two of the underlying assumptions in classical attribution are that the asset-specific returns (over and above those of the factors) are uncorrelated with the factor returns, and that they are uncorrelated with each other. However, this is not always the case. Below, our example illustrates that the first assumption does not necessarily hold when factor attribution is run. Importantly, if the residual return is related to the factor return(s), then the exposure estimates must be wrong. Our solution seeks to adjust the *exposure*, so the residual and factor returns are uncorrelated.

Our Response to the Problem of Misspecified Exposures

Adjusted factor-based attribution gives us the chance to say, “Oops, our factor exposures were not correctly represented,” and fix our model after the realization of returns has occurred. Although this may seem fishy, it is justifiable. While we cannot make such adjustments to our model when estimating expected returns or risk, in attribution, we are trying to understand what has already happened. Therefore, after we have seen the realized returns, we can adjust our sensitivities (β_i) to better represent what occurred. We are simply adjusting the model to reflect reality.

² Technically, this will be true only if the portfolio is exactly the S&P 1500.

³ Under certain assumptions, the exposures used to approximate the sensitivities to representative portfolios are equivalent to the sensitivities computed in the regression.

⁴ Asset style-factor exposures are calculated cross-sectionally, normalized and trimmed. See Axioma’s “Risk Model Handbook” for more details. Stocks in a given industry or country, or that are based in a given currency, are given an exposure of 1 to that industry, country or currency, while all other stocks have a zero exposure.

To illustrate the issues, we start with a simple optimized portfolio:

Maximize: Expected return (based on Axioma's Growth Factor)
Subject to: Long Only and Fully Invested
Active Risk Constraint 3%
Active Sector Bounds +/- 4%
Active Asset Bounds +/- 3%
Benchmark: Russell 1000

The portfolio was rebalanced monthly, and the tests covered the period 1995 to 2014.

We then ran attribution using a simplified version of our standard risk models. Factors included are shown in Table 1.

Table 1. Factors Included in Sample Attribution


10 GICS Sector Factors
Market Sensitivity
Momentum
Size
Value
Growth

In Table 2, we show some results from factor attribution using the risk model described above. The specific contribution is significantly negative, which also suggests the attribution is not telling the whole story. This becomes even more apparent when we look at the time series of factor and specific contributions, which indicate a strong negative correlation (see the chart). Since there should have been little specific risk in the portfolio (if it had been constructed to maximize factor exposure), this negative specific contribution seems to be too high in magnitude. Performance attribution will usually show specific return, and that return may be negative or positive. The real issue here, however, is the high negative correlation (-0.32) between factor and specific returns. A biased, exaggerated exposure to Growth led the attribution to provide an incorrect inference, namely that the specific return was so significantly negative. In this case, correcting the correlation brought the specific return close to zero, although that will not always occur.

Table 2. Selected Attribution Results

Active Return	1.47%
Specific Contribution (SC)	-0.88%
Factor Contribution (FC)	2.35%
SC T-Stat	-1.58
FC T-Stat	4.29
SC-FC Correlation	-0.32

Cumulative Factor and Specific Contributions



Of course, this could just be an artifact of this particular portfolio. However, we have tested this on many combinations of portfolios and returns models, and the results are consistent. When the exact factors used to construct a portfolio are present in the attribution model, we typically see that the time series of factor and specific returns are negatively correlated. In contrast, if a portfolio is constructed with factors that are not in the attribution model, or factors that are similar but not exact (for example, where the portfolio uses sales/price but the returns model uses earnings/price), then the correlation between factor and specific returns is likely to be positive.

What Causes the Misattribution? Estimates of a portfolio's sensitivity to a factor may be biased. To illustrate this, suppose that two active portfolios both bet on a single factor in the following way: the first goes long the first decile and short the 10th decile; the second goes long the third decile and short the eighth decile. Suppose that the two active portfolios both have an exposure of one to the factor: the two active portfolios are likely highly correlated, but they will not be identical. Further, we will assume we have a factor whose return is equal to the return of the first active portfolio, and we will use this factor to explain the returns of the two active portfolios. The first active portfolio will have a beta of 1 to the factor. So, our exposure of 1 is correct. The second active portfolio may have a beta of only 0.8 to the factor and yet still have an exposure of 1. This means that our exposure is a biased estimate of the sensitivity.

Our Proposal Recognizes and Corrects For the Negative Correlation, Eliminating This Source of Bias

Adjusted Performance Attribution uses the realized returns to correct for the bias mentioned above. Our goal is to eliminate the correlation between the specific return and the factors. The residual (specific) return is regressed on the factors' contributions⁵ to identify which factors can explain the residual return. It is likely that the factors ultimately considered for adjustment will be those factors for which the portfolio has a reasonably high exposure. The adjustment regression produces betas, or sensitivities, of the residual return to the factors

in the regression. The exposure of the portfolio is then adjusted by the beta⁶, and that exposure is then multiplied by the factor return to get the “new” contribution. This is done for each factor that had a statistically significant beta coefficient in the adjustment regression, but is generally assumed to apply to only a few of the factors. So, if Value’s beta from the regression was -0.5, and the initial exposure was 1, the new exposure would be 0.5, and that would be multiplied by the factor return to get the new contribution. The adjusted specific return is, like the original specific return, the residual after accounting for all of the factor contributions.

The resulting adjusted factor and specific contributions are uncorrelated, leading to more accurate inferences from the attribution report.

In a nutshell, we adjust our model by regressing the time-series of residual return contributions (returns not explained by the factor exposures) against the contributions from the factor portfolios. Any portion of the residual returns that can be explained by the factor portfolios is then pushed back into the portfolio contributions by adjusting the exposures, thereby getting a better representation of the actual exposures, and more accurate and intuitive attribution results.

Example of Adjusted Attribution

To illustrate the adjusted attribution methodology, we will expand upon the motivating example presented earlier. In running our regressions, two factors ended up on the final list of statistically significant factors, Value and Growth. The beta for Value was -0.65 and for Growth was -0.39. This means our exposures to Value and Growth were overstated, and should be only 35% and 61% of their original values, respectively.

Table 3 illustrates the exposures to all the factors in the model.

Table 3. Original and Adjusted Exposures

Factor	Active Exposure	
	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

In Table 4, we compare, and expand upon, the attribution results using adjusted attribution to the original results in Table 2. We can see that correlation between the adjusted factor and specific contributions changed from a significant -0.32 to just 0.09. Annualized factor contribution decreased from 2.35% to 1.50%. In addition, the factor return came mainly from the Growth factor, which is what we would have expected given how we constructed the portfolio. Specific contribution fell to almost zero, as compared with being far more negative under the initial attribution. Factor contribution, in total, maintained its significance, whereas the significance of the specific contribution was roughly zero. Industry contribution was unchanged.

⁵ Note that we use contributions rather than factor returns. This avoids the problem of allowing a factor to explain a large portion of return when the unadjusted factor exposure is near zero. It may also be more appropriate for cases where factor exposures change through time.

⁶ $(1 + \beta) \times \text{original exposure}$

Table 4. Full Attribution Results

	Initial	Adjusted
Portfolio Return	10.72%	10.72%
Benchmark Return	9.25%	9.25%
Active Return	1.47%	1.47%
Factor Contributions	2.35%	1.50%
Industry	-0.15%	1.45%
Styles	2.50%	0.05%
Growth	2.65%	1.60%
Market Sensitivity	-0.09%	-0.09%
Momentum	0.09%	0.09%
Size	0.15%	0.15%
Value	-0.29%	-0.10%
Specific Contribution	-0.88%	-0.03%
FC Volatility	2.37%	1.53%
SC Volatility	2.41%	2.21%
SC T-Stat	-1.58	-0.06
FC T-Stat	4.29	4.22
SC FC Corr	-0.32	0.09

The goal of this paper is to provide a peek into Axioma's Adjusted Performance Attribution theory and practice. As part of this study we ran a number of tests using single and multiple factors as alphas in our optimizations, as well as various risk models. Readers can find the results of these tests and much more detailed theory and analysis in Stubbs and Jeet (2016), which was referenced at the beginning of the paper⁷.

Conclusion

Factor attribution is a widely used methodology to determine the sources of return in a portfolio. However, even if the all factors used in portfolio construction are included in the attribution, the impact of some may be overstated. This may lead to a highly negative correlation between the factor returns and the residual return, which violates one of the major principles of factor attribution. The overstatement is in the form of factor exposures appearing higher than they should be.

Axioma has implemented a methodology derived by Stubbs and Jeet that corrects for the tendency of factor and residual returns to be negatively correlated and does a more accurate job of explaining portfolio returns. This adjusted performance attribution, found and easily run in Axioma's Portfolio Analysis product, uses a variable-selection heuristic to choose which variables are adjusted. To ensure that we do not overfit the data,

⁷ For more details on Axioma's view of attribution, see Stubbs and Jeet, "Adjusted Factor-Based Attribution," Axioma Research Paper No. 59, and Canova, Brown and Steffanus, "Enhancing the Investment Process with a Custom Risk Model," Axioma Research Paper No. 42

we adjust only those factors that are found to be statistically significant in the time series regression, as described above.

Although it does not correct for every potential bias that drives negative specific return, this methodology does eliminate the correlations between factor and specific contributions that we set out to accomplish. While there are multiple reasons for negative contributions, adjusted attribution eliminates one of them and goes a long way to helping managers, especially factor-based managers, explain their returns.



Contact us to learn more about how Axioma can help you better manage risk and enhance your investment process.

United States and Canada: +1-212-991-4500

Europe: +44-207-856-2424

Asia: +852-8203-2790

Sales: sales@axioma.com

Client Support: support@axioma.com

Careers: careers@axioma.com