Risk-adjusted Performance Attribution

This article presents an approach for attributing the risk-adjusted performance, as given by the information ratio, to a set of investment decisions. The portfolio information ratio is shown to be the weighted average of the component information ratios. The relevant weights, however, are the risk weights and not the investment weights. The component information ratios are given by the stand-alone information ratios scaled by an inflation factor due to diversification.

Jose Menchero, Ph.D., CFA
is an Executive Director in the research department at MSCI Barra, where he focuses on factor modeling and portfolio analytics. Prior to joining MSCI Barra, Jose was Head of Quantitative Research at Thomson Financial, where he worked extensively on performance attribution analysis and factor risk modeling. Jose has several publications in these areas, including two that received Dietz Awards.

Before entering the field of finance, Jose was a Professor of Physics at the University of Rio de Janeiro, Brazil. His area of research was in the Quantum Theory of Solids, and he has several publications in this field. Jose serves on the Advisory Board of The Journal of Performance Measurement*. He holds a B.Sc. degree in Aerospace Engineering from the University of Colorado at Boulder and a Ph.D. in Theoretical Physics from the University of California at Berkeley.*

INTRODUCTION

Performance attribution analysis is widely used in the investment industry for evaluating the performance of active fund managers. Attribution effects are specified, which quantify the impact of investment decisions on fund performance. For such an analysis to be meaningful, of course, the attribution model must reflect the investment process. For example, if the manager follows a sector-based investment strategy, then a sector-based performance attribution model, such as developed by Brinson and Fachler (1985), would be appropriate. In this approach, active return (the difference between portfolio and benchmark return) is decomposed into an allocation effect and a selection effect. The former measures the impact of over-weighting or under-weighting individual sectors, whereas the latter measures the effect of over-weighting or under-weighting individual securities within sectors.

While the Brinson model gives clear insight into the sources of active return, it is silent on the question of risk. For active managers trying to outperform a benchmark, the single most pertinent measure of risk is the tracking error (also known as active risk), defined as the volatility of active returns. Mina (2002), however, showed that *ex-ante* (predicted) tracking error could be attributed to the same set of decision variables as used in the original Brinson model. Later, Gregoire (2005) showed how to attribute *ex-post* (realized) tracking error at the security level. Xiang (2005/2006) generalized this *ex-post* decomposition to the sector-level attribution effects. Later, Menchero and Hu (2006) treated the case of *ex-ante* and *ex-post* risk attribution at both the security level and the sector level.

A natural question is whether performance attribution and risk attribution can be treated simultaneously within a risk-adjusted performance attribution framework. For active managers, a widely used measure of risk-adjusted performance is the information ratio. Technically speaking, the information ratio is defined as the portfolio alpha divided by residual risk, as discussed by Grinold and Kahn (2000). However, as noted by Goodwin (1998), if the portfolio Beta is reasonably close to one, then the information ratio is simply active return divided by tracking error. In practice, even if the portfolio Beta does differ from one, the information ratio is still often defined as active return divided by tracking error, and this is the working definition that we employ in this article.

One approach to risk-adjusted attribution was developed by Ankrim (1992), in which he adjusted the original Brinson attribution effects to reflect systematic market risk, as measured by the asset Betas. Missing market risk, as measured by the asset Betas. Missing from Ankrim's approach, however, was any consideration of the information ratio. Later, Obeid (2005) reformulated Ankrim's approach and proposed certain modifications, such as replacing the market portfolio by the benchmark. Nonetheless, Obeid kept within the Ankrim framework of performing a Beta adjustment, and similarly did not consider the information ratio.

The purpose of this paper is to present a risk-adjusted performance attribution analysis based on the information ratio. We first provide an overview of how performance and risk can be attributed to a set of decision variables. We then present a concrete example of such analysis for an investment process consisting of asset allocation and security selection decisions. Next, we describe a methodology for decomposing the information ratio according to the same set of decision variables as used in the performance attribution and risk attribution analysis. Finally, we present an example which illustrates the concept of risk-adjusted performance attribution.

PERFORMANCE AND RISK ATTRIBUTION

Let *m* denote an individual investment decision, and let Q_{mt} be an attribution effect that measures the contribution that this decision makes to the active return for period *t*. Properly defined, these attribution effects fully explain the active return for the period. Hence,

$$R_{At} = \sum_m Q_{mt} \quad , \tag{1}$$

where R_{At} is the active return.

For concreteness, it may be convenient to think of Q_{mt} in terms of the classical attribution effects: allocation and selection. The allocation effect for sector *i* and period *t* is defined as

$$A_{it} = \left(w_{it}^P - w_{it}^B\right)\left(r_{it}^B - R_{Bt}\right) \quad , \tag{2}$$

where w_{it}^P is the portfolio sector weight, w_{it}^B is the benchmark sector weight, r_{it}^B is benchmark sector return, and R_{Bt} is the overall benchmark return for the period. The quantity $\left(w_{it}^P - w_{it}^B\right)$ is the active weight in the sec tor, and $\left(r_{it}^B - R_{Bt}\right)$ is the relative return of the sector. Thus, a manager earns positive allocation effect either by over-weighting outperforming sectors or by under-weighting underperforming ones.

The selection effect is given by

$$S_{it} = w_{it}^P\left(r_{it}^P - r_{it}^B\right) \quad , \tag{3}$$

where r_{it}^P is portfolio sector return. The quantity $\left(r_{it}^P - r_{it}^B\right)$ represents the active return within the sector. Hence, for a long-only portfolio, a manager earns positive selection effect within a sector by outperforming the benchmark within the sector.

The allocation effect and selection effect, summed over sectors,

$$R_{At} = \sum_i \left(A_{it} + S_{it}\right) \quad , \tag{4}$$

fully explain the active return for the period. Note that Equation 4 represents a specific example of Equation 1, in which the decisions consist of allocation and selection.

While Q_{mt} can be thought of concretely in terms of allocation and selection effect, it proves useful to consider Q_{mt} more generally as any component of active return. For instance, in a fixed income context, Q_{mt} might represent the contribution to active return from a steepening of the yield curve, or a tightening of a sector spread, and so forth. The only constraints we place on Q_{mt} are: (a) the attribution effects must be defined in a manner that reflects the investment process, and (b) the attribution effects must fully account for the active return.

Typically, attribution effects are computed over sub-periods (*e.g.,* daily) and then aggregated over the longer interval of time (say, six months) that is of interest to the performance analyst. This is accomplished by means of a multi-period linking algorithm,

$$Q_m = \sum_t \beta_t\, Q_{mt} \quad , \tag{5}$$

where Q_m represents the linked attribution effect, and is the linking coefficient for period *t*. These coefficients β_t are computed, for example, using the Optimized Linking Algorithm, as described by Menchero (2000, 2004). The sum of the linked attribution effects fully accounts for the multi-period active return; *i.e.,*

$$R_A = \sum_m Q_m \quad , \qquad (6)$$

where R_A is the difference between the compounded portfolio return and the compounded benchmark return. Equations 1, 5, and 6, taken together, form a general representation of performance attribution analysis, which quantifies how individual investment decisions contributed to the multi-period active return.

Return analysis, however, only tells part of the story. One must also ask how much risk was taken in order to achieve the realized return. For instance, a manager who outperforms a benchmark by 50 bps with two percent tracking error has certainly done better than another manager who outperforms by the same 50 bps, but while assuming six percent tracking error.

Tracking error can be computed either as a predicted value or as a realized value. Typically, tracking error forecasts are made using a multi-factor risk model. The realized tracking error, by contrast, is computed by the time series of active returns,

$$\sigma(R_A) = \sqrt{\frac{1}{T-1}\sum_t \left(R_{At} - \overline{R}_A\right)^2} \quad , \qquad (7)$$

where T is the number of periods.

Menchero and Hu (2006) presented an approach for decomposing both *ex-ante* and *ex-post* risk according to the decision variables used in the investment process. For the case of *ex-post* risk attribution, the realized tracking error can be decomposed as

$$\sigma(R_A) = \sum_m \sigma(Q_m)\,\rho(Q_m, R_A) \quad , \qquad (8)$$

where $\sigma(Q_m)$ is the realized volatility of attribution effect Q_m, and $\rho(Q_m, R_A)$ is the realized correlation between the attribution effect and the active return. The expression for the *ex-ante* decomposition of tracking error is actually the same as in Equation 8, except that the volatilities and correlations are based on risk model forecasts.

The risk decomposition specified by Equation 8 is very intuitive. It simply says that there are two drivers of tracking error within a portfolio. The first is the stand-alone volatility $\sigma(Q_m)$ of the attribution effects, and the second is the degree of correlation $\rho(Q_m, R_A)$ between these attribution effects and the active return. Thus, for positive correlation, the risk associated with an investment decision rises in proportion to the stand-alone volatility of the corresponding attribution effect. On the other hand, even if the stand-alone volatility of an attribution effect is high, it will contribute only a small amount to tracking error if its correlation with the active return is low. Some investment decisions actually contribute negatively to tracking error. If the correlation is negative, then when the active return is up, the attribution effect tends to be down, and vice versa. Either way, the result is to dampen the magnitude of the active return, and thus reduce tracking error, consistent with Equation 8.

Example 1: Performance Attribution

We now consider a concrete example using a global value portfolio versus a global growth benchmark. This example, though contrived, nonetheless clearly illustrates the relevant concepts. We construct our portfolio and benchmark using the MSCI Developed Markets Value and Growth Indices. The growth benchmark is segmented into Europe (ex U.K.), Japan, U.K., and USA., with weights of 25, 10, 15, and 50 percent,

Table 1: Performance Attribution Analysis for 6/25/2005 to 12/31/2005

Market	***Port. Weight***	***Bench. Weight***	***Port. Return***	***Bench. Return***	***Alloc. Effect***	***Select Effect***	***Total Effect***
Europe	***45.01***	***24.97***	***15.52***	***9.36***	***-0.17***	***2.78***	***2.61***
Japan	***20.15***	***10.16***	***26.82***	***36.69***	***2.44***	***-1.71***	***0.72***
U.K.	***4.96***	***14.97***	***4.56***	***6.65***	***0.35***	***-0.10***	***0.24***
U.S.A.	***29.88***	***49.90***	***5.46***	***6.94***	***0.66***	***-0.51***	***0.15***
Total	***100.00***	***100.00***	***14.06***	***10.32***	***3.28***	***0.46***	***3.74***

respectively. The value portfolio is also invested in the same markets, but with weights of 45, 20, 5, and 30 percent, respectively. Hence, the portfolio is over-weight Europe by 20 percent, over-weight Japan by 10 percent, under-weight U.K. by 10 percent, and under-weight U.S.A. by 20 percent. The portfolio and benchmark are rebalanced monthly to the specified weights.

In Table 1, we present the performance attribution analysis (using U.S. Dollar as the base currency) for the 27-week period from June 25, 2005 to December 31, 2005. The attribution effects were computed on a daily basis and linked using the Optimized Algorithm. The portfolio and benchmark weights reported in Table 1 represent the average of the beginning-of-day weights, and differ slightly from the stated nominal weights due to intra-month fluctuations. The portfolio and benchmark returns were 14.06 and 10.32 percent, respectively, for an active return of 374 bps. Table 1 immediately reveals that the allocation effect was the main contributor (328 bps) to the active return. Moreover, the majority of this (244 bps) was attributable to the single decision to over-weight Japan, which outperformed significantly over the period. While selection effect contributed relatively little (46 bps) to active return, Table 1 reveals that this was the result of a large cancellation between positive selection in Europe and negative selection in the other markets.

Example 2: Risk Attribution

In Table 2 we present the *ex-post* risk attribution analysis for the same portfolio/benchmark combination over the same 27-week period as considered in Table 1. The risk was computed as the standard deviation of weekly returns and was annualized by multiplication of these weekly volatilities by $\sqrt{52}$. The realized tracking error over this period was 3.56 percent. Of this, slightly more than half (193 bps) was attributable to selection decisions, and the rest (163 bps) was due to allocation decisions.

On the selection side, most of the risk was concentrated in Europe (87 bps) and the U.S.A. (81 bps). The high risk associated with these decisions can be understood as the result of (a) large stand-alone volatilities, and (b) high correlations with the active returns (70 percent). By contrast, Japanese security selection contributed only a small amount (10 bps) to overall tracking error. From Table 2, we see that this was mostly due to the low correlation (0.14) between Japanese value-versus-growth and global value-versus-growth.

Greater insight into the sources of selection volatility can be found from the definition of selection effect, given by Equation 3. If the weights are taken as constant (which is nearly true in our example), then the selection volatility is simply the portfolio sector weight times the volatility of active returns within the sector. This relationship allows us to infer the tracking error within the sector. For instance, dividing the 70 bps selection volatility in Japan by the 20 percent portfolio weight leads to a realized tracking error of 3.5 percent for Japanese value-versus-growth.

On the allocation side, the greatest risk was due to the decision to over-weight Japan, which contributed 82 bps to tracking error. This is perhaps a surprising result, considering that the nominal size of the bet (10 percent over-weight) was less than, say, the over-weight in Europe (20 percent). Nonetheless, the result can be understood as the consequence of a large stand-alone

Table 2: *Ex-post* Risk Attribution Analysis (annualized) for 6/25/2005 to 12/31/2005

Market	***Alloc Vol.***	***Alloc Corr.***	***Alloc Risk Contrib***	***Select Vol.***	***Select Corr.***	***Select Risk Contrib***	***Active Risk Contrib***
Europe	***1.04***	***0.36***	***0.38***	***1.26***	***0.70***	***0.87***	***1.25***
Japan	***1.23***	***0.66***	***0.82***	***0.70***	***0.14***	***0.10***	***0.91***
U.K.	***0.60***	***-0.39***	***-0.23***	***0.28***	***0.53***	***0.15***	***-0.08***
U.S.A.	***0.88***	***0.75***	***0.66***	***1.17***	***0.70***	***0.81***	***1.48***
Total	***2.27***	***0.72***	***1.63***	***2.51***	***0.77***	***1.93***	***3.56***

allocation volatility (1.23 percent) and a high correlation (0.66) with the active return.

Greater insight into the sources of allocation volatility can be found from the definition of allocation effect, given by Equation 2. We see that the volatility of the allocation effect is driven by the size of the active weight as well as the volatility of the relative returns. Again, if we assume constant weights over the analysis period (an excellent approximation in our example), then we can estimate the volatility of relative return. For instance, dividing the allocation volatility in Japan (1.23 percent) by the active weight (10 percent) leads to a relative-return volatility of 12.3 percent. This volatility, which is quite large, represents the tracking error of the Japanese growth index versus the global growth benchmark. By contrast, similar computations for other markets give much smaller values. For example, European growth versus global growth gives merely 5.20 percent tracking error (1.04 percent divided by 20 percent).

Interestingly, the decision to under-weight the U.K. market actually reduced tracking error. The reason is that the U.K. relative return was positively correlated with the active return, but the U.K. allocation effect was negatively correlated with the active return due to the under-weight. This example highlights the fact that the sources of risk are not the same as might be suggested by merely considering the nominal sizes of the bets.

RISK-ADJUSTED PERFORMANCE ATTRIBUTION

A key measure of risk-adjusted performance for active managers is the information ratio, defined here as the active return divided by the tracking error,

$$IR = \frac{R_A}{\sigma(R_A)} \quad . \tag{9}$$

Substituting Equation 6 for the numerator, we have

$$IR = \sum_m \frac{Q_m}{\sigma(R_A)} \quad . \tag{10}$$

It now proves useful to multiply and divide by the contribution to tracking error from Q_m, specified by Equation 8. The result is:

$$IR = \sum_m \left(\frac{\sigma(Q_m)\,\rho(Q_m, R_A)}{\sigma(R_A)} \right) \left(\frac{Q_m}{\sigma(Q_m)\,\rho(Q_m, R_A)} \right) . \tag{11}$$

This decomposition of information ratio was also discussed by Xiang (2005/2006). The first term is interpreted as the *risk weight* of decision *m*,

$$u_m \equiv \frac{\sigma(Q_m)\,\rho(Q_m, R_A)}{\sigma(R_A)} . \tag{12}$$

Note that these weights sum to unity, *i.e.*, $\sum_m u_m = 1$. We define the *component information ratio* associated with decision *m* as

$$IR_m = \frac{1}{\rho(Q_m, R_A)} \left(\frac{Q_m}{\sigma(Q_m)} \right) . \tag{13}$$

Note that the quantity in parenthesis, $[Q_m / \sigma(Q_m)]$, represents the stand-alone information ratio associated with decision *m*. However, in the context of the portfolio, this is amplified by the factor $[1/\rho(Q_m, R_A)]$, which may be regarded as a diversification benefit. Substituting Equations 12 and 13 into Equation 11, we find

$$IR = \sum_m u_m\, IR_m \quad . \tag{14}$$

This says that the portfolio information ratio is the weighted average of the component information ratios. The relevant weights, however, are not the investment weights but rather the risk weights.

A nice feature of expressing the information ratio as the weighted average of component information ratios is that these quantities can then be aggregated in precisely the same way that security returns are aggregated to the sector level. For instance, let *M* denote a decision group consisting of individual decisions, each labeled by *m*. Then we can write the information ratio as the weighted average over decision groups,

$$IR = \sum_M U_M\, IR_M , \tag{15}$$

where U_M is the risk weight associated with decision group *M*,

$$U_M = \sum_{m \in M} u_m \quad , \tag{16}$$

and the decision-group information ratio is

$$IR_M = \frac{1}{U_M} \sum_{m \in M} u_m \, IR_m \, . \qquad (17)$$

To be specific, IR_M might be the component information ratio for allocation effect in aggregate, whereas IR_m would represent the component information ratio for a single allocation decision.

Example 3: Risk-Adjusted Performance Attribution

We measure the risk-adjusted performance by the information ratio. Note, however, that information ratios, by convention, are typically stated on an annualized basis. We have seen from Table 2 that the annualized risk was 3.56 percent. The active return in Table 1, however, has not been annualized. The annualized active return for the example in Table 1 is

$$\left[(1.1406)^{52/27} - (1.1032)^{52/27}\right] \quad ,$$

or 8.02 percent. Hence, to annualize the attribution effects in Table 1, we must multiply by (8.02/3.74). The annualized information ratio is therefore given by (8.02/3.56), or 2.25.

Table 3 reports the risk-adjusted performance attribution analysis for the same portfolio/benchmark combination considered in Tables 1 and 2. The risk weight, given by Equation 12, tells what fraction of the risk budget was spent on a particular investment decision. For instance, 23 percent of the risk budget was used on over-weighting Japan, and 25 percent on European security selection. The risk weight for the U.K. allocation effect was -7 percent, reflecting the fact that this decision actually reduced tracking error.

The component information ratio for each attribution effect was computed using Equation 13. On a risk-adjusted basis, the best decision was European security selection, with a component information ratio of 6.83, followed closely by Japanese allocation (IR_m = 6.42). Since these decisions also had large risk weights, they explain most of the realized information ratio. The worst performing decision was Japanese selection, with a component information ratio of -37.92. This number may seem astonishingly large, but the reader must remember that it is not a stand-alone information ratio, which would typically be much smaller. To obtain the stand-alone information ratio, we must multiply by the correlation, given in Table 2 as 0.14. Therefore, the stand-alone information ratio for this decision was -5.31, which, though still large, is much smaller than previously.

This illustrates an intrinsic difficulty with information ratio analysis: Low correlations, which appear in the denominator of Equation 13, can greatly inflate the component information ratios. This difficulty is confounded by the fact that the estimation error for correlation coefficients can be quite high. Assuming normally distributed disturbance terms, the standard error is approximately $\text{se}(\rho) \approx \sqrt{(1-\rho^2)/T}$. Therefore, for Japanese selection effect $(\rho = 0.14)$, over T=27 weeks, the standard error is roughly 0.19, which is larger than the correlation itself.

Note that the component information ratio for U.K. allocation is negative. In cases of negative risk weights, however, this is desirable; we see that the U.K. allocation decision contributed positively to information

Table 3: Risk-adjusted Performance Attribution Analysis (annualized) for 6/25/2005 to 12/31/2005

Market	*Alloc. Risk Weight*	*Alloc. Info. Ratio*	*Alloc. IR Contrib.*	*Select Risk Weight*	*Select Info. Ratio*	*Select IR Contrib.*	*Total IR Contrib.*
Europe	***0.11***	***-0.97***	***-0.10***	***0.25***	***6.83***	***1.67***	***1.57***
Japan	***0.23***	***6.42***	***1.47***	***0.03***	***-37.92***	***-1.04***	***0.43***
U.K.	***-0.07***	***-3.17***	***0.21***	***0.04***	***-1.47***	***-0.07***	***0.14***
U.S.A.	***0.19***	***2.13***	***0.40***	***0.22***	***-1.34***	***-0.30***	***0.10***
Total	***0.46***	***4.30***	***1.98***	***0.54***	***0.50***	***0.27***	***2.25***

ratio. This is akin to the idea of earning positive contribution to portfolio return by shorting stocks with negative returns.

Finally, we use Equation 17 to aggregate information ratios for the allocation and selection decision groups. The component information ratio for the allocation decision was 4.30, much greater than the 0.50 attained from selection decisions. This shows that the allocation decisions were the main drivers of risk-adjusted performance. Unfortunately, the manager used less than half of the risk budget on allocation decisions. In retrospect, the manager would have achieved an even higher information ratio by concentrating more risk on allocation decisions instead of selection.

SUMMARY

The information ratio is perhaps the most pertinent measure of risk-adjusted performance for active portfolio managers. In this article, we present a general methodology for attributing the realized information ratio to a set of investment decisions. We have shown that the portfolio information ratio is the weighted average of the component information ratios. The weights, however, are not the investment weights but rather the risk weights. We also studied an example of a global value portfolio relative to a global growth benchmark. We considered this example first from a performance attribution perspective, second from a risk attribution perspective, and finally from a risk-adjusted performance attribution point of view. The performance attribution analysis (Table 1) explains clearly where the active return came from but is silent on the sources of risk. The risk attribution analysis, presented in Table 2, explains the sources of risk. The risk-adjusted attribution analysis, in Table 3, combines the two by showing the impact of investment decisions on the information ratio.

REFERENCES

Ankrim, Ernest M., "Risk-Adjusted Performance Attribution," *Financial Analysts Journal*, March-April 1992, pp. 75-82.

Brinson, Gary, and Nimrod Fachler, "Measuring Non-U.S. Equity Portfolio Performance," *Journal of Portfolio Management*, Spring 1985, pp. 73-76.

Goodwin, Thomas, "The Information Ratio," *Financial Analysts Journal*, July/August 1998, pp. 34-43.

Grégoire, Philippe, "*Risk Attribution,*" Working Paper, 2005.

Grinold, Richard, and Ron Kahn, Active Portfolio Management, New York, NY: McGraw Hill, 2000.

Menchero, Jose, "An Optimized Approach to Linking Attribution Effects Over Time," *The Journal of Performance Measurement*, Fall 2000, Vol. 5 No. 1, pp. 36-42.

Menchero, Jose, "Multi-period Arithmetic Attribution," *Financial Analysts Journal*, July/August 2004, pp. 76-91.

Menchero, Jose, and Junmin Hu, "Portfolio Risk Attribution," *The Journal of Performance Measurement*, Spring 2006, Vol. 10 No. 3, pp. 22-33.

Mina, Jorge, "Risk Attribution for Asset Managers," *RiskMetrics Journal*, 2002, Vol. 3, No. 2, pp. 33-55.

Obeid, Alexander, "Reformulating Ankrim's Risk-adjusted Performance Attribution," *The Journal of Performance Measurement*, Spring 2005, Vol. 9 No. 3, pp. 8-25.

Xiang, George, "Risk Decomposition and its Use in Portfolio Analysis," *The Journal of Performance Measurement*, Winter 2005/2006, Vol. 9 No. 2, pp. 26-32.