The Decomposition Verses The Decision-Evaluation of Active Risk-Adjusted Returns

Andre Mirabelli, Ph.D., Opturo Inc.

Modeling is more than a mathematical exercise. Its proper intent is to provide insight into the actual structure of the system being addressed. This requires more care then is typically exercised in the construction of models evaluating the amount to be attributed to each decision that contributes to a fund's active return or active risk or risk-adjusted return.

Decomposition

Traditional 'Decision' Attribution

One can take the aim of decision (as opposed to market) attribution to be the decomposition of a fund's ex post active return, in terms of the ex post weights and returns of the components of the fund in the manner of Brinson-type Return Attribution (Brinson, 1986). Then modeling success can be rigorously gauged by whether the weights sum to one and, for arithmetic attribution, whether the attributes sum to the active return. On a less rigorous level, successful models are also required to have their attributes have a reasonable relationship to properties that can be interpreted as separate causes of the active return.

Active Return Decomposition

For some single period, consider the following performance attribution. Begin by defining some category, c, level basis, B_c and the portfolio level basis $B = \Sigma_c B_c$ and the category-level gain by $G_c = V_c + S_c + Div_c - P_c - V0_c$. Here, V_c is the closing value for the day held in category c, S_c is the amount of c that was sold during the day, Div_c is the generalized dividend earned by c during the day (possibly including as negative values any fees paid), P_c is the amount of c that was purchased during the day and Vo_c is the value of c at the open of the day. Then define the category weight by $W_c = B_c/B$ and category return, R_c, by the category gain, G_c per unit basis, $R_c = G_c/B_c$.

By defining the portfolio-level return as the portfolio level gain per unit basis, $R = G/B$, and noting the additivity of gains, $G = \Sigma_c G_c$, it follows that

$R = G/B = \Sigma_c G_c/B = \Sigma_c [(B_c/B)*(G_c/B_c)] = \Sigma_c W_c*R_c$.

(Though it will not be addressed herein, common approaches to specifying B_c, in the presence of trades made during the trading day that are settled after the close, are also legitimate targets of an analysis similar to that presented below.)

In the manner of Brinson-Fachler (Brinson, 1985), one can then decompose the portfolio level active return $S \equiv R^F - R^B$ (Fund minus Benchmark) into Allocation, Selection and Interaction components:

$All^{BF} = \Sigma_c (W^F{}_c - W^B{}_c)*(R^B{}_c - R^B)$,

$Sel^{BF} = \Sigma_c W^B{}_c*(R^F{}_c - R^B{}_c)$ and

$Int^{BF} = \Sigma_c (W^F{}_c - W^B{}_c)*(R^F{}_c - R^B{}_c)$.

(It will be seen that the critique by this paper of standard approaches is self-contained at the level of decisions. Drilling down further, to the component-level for each decision where one might define, for example, $All^{BF}{}_c$, only uncovers additional problems for standard arguments.)

Thus, it analytically follows that

$\Sigma_c W_c = 1$ (Justification 1)

for both the fund and for the benchmark and that

$S = R^F - R^B = All^{BF} + Sel^{BF} + Int^{BF}$ which can be stated in general as

$S = \Sigma_j X_j$, (Justification 2)

where the set $\{X_j\} = \{All^{BF}, Sel^{BF}, Int^{BF}\}$ spans the attributes that decompose the active return.

The Brinson-Fachler Allocation, All^{BF}, is seen to be an increasing function of the category level bet, $W^F{}_c - W^B{}_c$, when the category return, $R^B{}_c$, of the benchmark did better than the benchmark return, R^B, as a whole (Justification 3). This is consonant with the intuition that it should be helpful to the aim of outperforming the benchmark by means of allocation for one to have over-weighted a category that did better than the benchmark as a whole.

The Selection, Sel^{BF}, is seen, for positive definite benchmark weights, to be an increasing function of the amount to which the fund return in a category outperformed the benchmark return in that category (Justification 4). This is consonant with the intuition that it should be helpful (i.e. add to the active return) to have the fund outperform the benchmark in a long category.

The Interaction, Int^{BF}, is seen to be an increasing function of the product, of active category weight times active category return, that one obtains when one over weights a category in which the fund outperformed the benchmark (Justification 5). This is consonant with the intuition that it should be helpful to have more of what one does well in.

This attempted explication of the active return of a fund relative to a benchmark in terms of the allocation, selection and interaction initiated by the portfolio construction process can seem satisfying till one notices that one can add, in turn, to All, Sel and Int any set of terms, AllK, SelK and IntK, respectively, such that AllK + SelK + IntK = 0. None of the justifications given above are necessarily thwarted by this amendment. Crucially, it is still the case that

$R^F - R^B = All^{BF} + Sel^{BF} + Int^{BF}$.

However, such an amendment clearly eviscerates any economic meaning that it was hoped this modeling process would elucidate. What now does

$All^{BF'} = All^{BF} + Allk = \Sigma_c (W^F{}_c - W^B{}_c)*(R^B{}_c - R^B) + AllK$

explicate if AllK can be any value at all, with no dependence upon the results achieved in the investment process? A process that can report any value at all, unconstrained by what actually transpired in the fund, for the amount that the allocation step in the investment process contributed to the portfolio's active return is worse than uninformative. It actually misinforms future investment decisions by misleading investment managers about the results of their past process.

One possible defense against such criticism is that the intent of All^{BF} requires that it be zero when the fund category weights match those of the benchmark and the intent of Sel^{BF} requires that it be zero when the returns of the fund's categories match those of the benchmark. Thus, it might be argued that this should be made a requirement for any model. However, even with the addition of this new requirement, it is still possible to impose the arbitrary terms in the presence of this added requirement if it is simply required that the arbitrary terms go to zero in just these circumstances. For example one could define them, as continuous functions, as:

$AllK = \{-1 + Exp\ \Pi_c\ [W^F{}_c - W^B{}_c]\}*AllK0,$

$SelK = \{-1 + Exp\ \Pi_c\ [R^F{}_c - R^B{}_c]\}*SelK0,$

$IntK = \{-1 + Exp\ \Pi_c\ [(W^F{}_c - W^B{}_c)*(R^F{}_c - R^B{}_c)]\}*IntK0,$

where AllK0, SelK0 and IntK0 can still be arbitrary as long as it is still the case that $AllK + SelK + IntK = 0$. Thus, AllK and SelK can, in general, still be any value. They just have to be zero in the special unlikely circumstances where the fund precisely reproduces the benchmark in some respect. The most negligible difference allows for the arbitrariness to exist in full.

Risk-Adjusted Active Return Decomposition

Relying on returns without considering the risk taken to obtain those returns is not a responsible approach to investment management. If risk is ignored, for all one knows, a good return value could just be a lucky landing on a wildly fluctuating period-to-date return time series. This is why risk adjusted returns are more important than simple returns by themselves. And that is why it is so important to supplement performance attribution with risk-adjusted performance attribution.

Risk-adjusted performance attribution aims to move beyond an evaluation of the sources of active return to an evaluation of the sources of risk-adjusted returns. (See "the literature": Menchero 2006, Menchero 2007 and Bertrand, 2009.)

The common approach to risk attribution can be summarized as follows.

Begin by recalling a few statistical definitions and their properties as applied to a time, t, series over a period T.

Average: $<x> \equiv \Sigma_{t\in T}[x_t/T]$.

Covariance of x and z: $cov(x, z) \equiv \Sigma_{t\in T}[(x_t - <x>)*(z_t - <z>)]$.

From these definitions it follows that

$$Cov(x + y, z) = \Sigma_{t\in T}[(x_t + y_t - <x + y>)*(z_t - <z>)]$$
$$= \Sigma_{t\in T}[(x_t - <x >)*(z_t - <z>)] + \Sigma_t[(y_t - <y >)*(z_t - <z>)]$$
$$= Cov(x, z) + Cov(y, z).$$

This can be generalized to: $Cov(\Sigma_j x_j, z) = \Sigma_j Cov(x_j, z)$.

Variance of x: $V(x) \equiv Cov(x, x)$. {Omit this: $Cov(x, z) = 0.5*[V(x) + V(z) - V(x-z)]$}

Standard deviation of x: $\sigma(x) \equiv [V(x)]^{1/2}$.

Correlation of x and z: $\rho(x, z) \equiv Cov(x, z)/[\sigma(x)*\sigma(z)]$.

We now note the following. Attribution focuses upon the creation of active values, that is, upon the difference between a fund value and the corresponding benchmark value.

Performance attribution focuses upon the active return of a portfolio, $S = R^P - R^B$. Risk attribution focuses upon values like the difference between the volatility of a fund and the volatility of its benchmark, $\sigma(R^P) - \sigma(R^B)$. Risk adjusted performance attribution focuses upon the active performance per risk, like the difference between the information ratio of a fund and the information ratio of a benchmark. However, since an information ratio is defined as the active return per unit of risk and since the active return for the benchmark is $R^B - R^B = 0$, it follows that the information ratio for the benchmark is zero and, thus, that the active information ratio is equal to the information ratio of the fund. Thus, it is the information ratio of a fund that risk-adjusted performance attribution addresses and which risk-adjusted active return decomposition decomposes.

For information ratio, it is standard to define risk as tracking error:

$TE \equiv \sigma(S)$.

Thus, recalling that X_j is the component of S,

Information Ratio: $IR \equiv <S>/\sigma(S) = \Sigma_j <X_j> / \sigma(S)$.

Following a procedure similar to the one above that decomposed the portfolio return into its weighted components we multiply and divide each addend in the IR equation by

$[\rho(X_j, S)*\sigma(X_j)]$ to get

$IR = \Sigma_j (\{ [\rho(X_j, S)*\sigma(X_j)]/ \sigma(S) \}*\{ <X_j> / [\rho(X_j, S)*\sigma(X_j)]\})$.

The literature now defines a "risk-weight"

$WR^{MHB}{}_j \equiv [\rho(X_j, S)*\sigma(X_j)]/ \sigma(S)$

and a "component information ratio" for the performance attribute X_j,

$IR^{MHB}(X_j) = \langle X_j \rangle / [\rho(X_j, S)*\sigma(X_j)]$.

The justification for this model is that

$\Sigma_j[WR^{MHB}{}_j] = \Sigma_j[\rho(X_j, S)*\sigma(X_j)]/ \sigma(S)]$

$= \Sigma_j[Cov(X_j, S)/\sigma(S)]/ \sigma(S)$

$= \Sigma_j [Cov(X_j, S)/V(S)]$

$= Cov(\Sigma_j X_j, S)/V(S)$

$= Cov(S, S)/V(S) =1$,

showing that the risk-weights are normalized, and that

$IR = \Sigma_j WR^{MHB}{}_j*IR^{MHB}(X_j)$,

showing that the total information ratio is the risk-weighted sum of the component information ratio for the risk-adjusted performance attribute. Just as in the case of performance attribution, these justifications are not at the level of the individual attributes but at the level of their sum.

For situations in which an attribute, X_j, is positively correlated, $\rho(X_{j,}, S) > 0$, with the active return, S, X_j does not hedge the portfolio. In such cases, the risk-weight, WR_j, for decision j is monotone increasing with the standard deviation of that decision's expected value of the attribute that contributes to the active return and the component information ratio for the performance attribute X_j is monotone with the attribute that contributes to the active return while being monotone decreasing with the standard deviation of that attribute. This is consonant with the intuitions that, in the absence of the above-mentioned hedging, more variable attribute contributions should be weighted more toward the risk-adjusted return and that attributes with a larger information ratio, $\langle X_j \rangle / [\rho(X_j, S)*\sigma(X_j)]$, should also contribute more to the active risk-adjusted returns.

However, just as is the case with performance attribution, it is possible to insert a degree of freedom.

For any arbitrary set $\{\rho\sigma K_j\}$, spanning the j decision attributes, such that

$\Sigma_j[\rho\sigma K_j] = 0$

it is the case that

$\Sigma_j (\{[\rho(X_j, S)*\sigma(X_j) + \rho\sigma K_j]/ \sigma(S)\}*\{ \langle X_j \rangle/ [\rho(X_j, S)*\sigma(X_j) + \rho\sigma K_j]\}) = \Sigma_j \langle X_j \rangle/\sigma(S) = IR$

and

$\Sigma_j[\rho(X_j, S)*\sigma(X_j) + \rho\sigma K_j]/ \sigma(S)]$

$= \Sigma_j[\rho(X_j, S)*\sigma(X_j)]/ \sigma(S)] + \Sigma_j[\rho\sigma K_j]/ \sigma(S)]$

$= \Sigma_j[\rho(X_j, S)*\sigma(X_j)]/ \sigma(S)] + 0 = 1$.

All the previous arguments for the justification of the decomposition still go through and we have, thus, that the risk weights can be defined by

$WR^{MHB'}{}_j \equiv [\rho(X_j, S)*\sigma(X_j) + \rho\sigma K_j]/ \sigma(S)$

and the information ratio of the performance attribute can be defined as

$IR^{MHB'}(X_j) \equiv \{ \langle X_j \rangle / [\rho(X_j, S)*\sigma(X_j) + \rho\sigma K_j]\}$,
where $\rho\sigma K_j$ is an arbitrary value.

As before, such an arbitrary value can be formulated to be zero in desired circumstances while still allowing it to be arbitrarily large in all other cases.

Thus, again we have a process that can report any value at all, unconstrained by what actually transpired in the fund, for the amount that the allocation step in the investment process contributed to the portfolio's active risk-adjusted return. Again, these reported values actually misinform future investment decisions by misleading investment managers about the results of their past fund-construction process.

Hats

Imagine that I have 2 hats and he has 4 hats and she has 9 hats.
And that some group purports to 'explain' the situation by the following model:

15 hats = 3 hats + 5 hats + 7 hats.

When it is pointed out that 3 hats does not correspond to the number of hats I have they respond that their reported values have the name 'hats' in them and that their total adds up to the correct value of 15. In addition they point out that their model preserves the intuitive property that I have less hats them him who, in turn, has less hats then her.

The correct response to such 'justifications' of their model is that global properties about the total number of hats and the mirroring of certain intuitions about the situation are not sufficient. An adequate model must assign values to the components that are structured in the way that the relevant evaluations of the components are structured and that this can only be achieved if the values assigned to the components meaningfully correspond to the properties of the components that we intend to model. This requires that we begin with a clear understanding of the component level properties themselves and model them, letting the global properties follow from the fact that we captured the component-level properties correctly.

In the case we are concerned with here, this means that if we are modeling the number of hats that each person has then we must assign that number to this property in the model. Then, at the global level, it will simply follow that the total number of hats assigned by the model is equal to the total number of hats had by the three hat holders and similarly for the relationships between these values.

Decision Evaluation

Modeling at the appropriate level

We clearly need an alternative approach to performance attribution and risk-adjusted performance attribution. Below, it will be seen that there is an alternative that ends up justifying the single period Brinson performance attribution model at the decision level with which we began but that leads to a totally different model for risk-adjusted performance attribution. In contrast to that offered by the literature, the alternative, decision evaluation approach justifies our terms at the level at which we want the terms to make economic sense. If we want the risk-adjusted attribute $\Delta IR(X_j)$ to make sense for each decision-type, j, then we must justify its definition at the level of each decision-type. In order to justify decision level values it is not enough to impose requirements at the portfolio level. To justify decision-level values we must impose their definitional-level values at the decision level. Decision-level values cannot be justified by imposing requirements on their roll-ups to the portfolio level. Because decision evaluation is an approach based on answering explicit economic questions, it does not just decompose $\Delta = \{\Delta R, IR, \Delta\sigma \ldots\}$. Instead it ascertains the effect of each individuated investment decision on Δ.

The amount of a parameter that is attributed to a decision is defined as the change in the parameter due to the decision. Thus, the economic question about a parameter addressed by decision evaluation is: By how much did the parameter change from right before to right after the implementation of the decision?

The Decision Process

Consider a single period investment process that, if no decision were made, would leave the fund a replica of the benchmark. Thus, if no decision is made the single period fund return would equal the benchmark return and the active return would be zero:

$R^F{}_0 = R^B = \Sigma_c W^B{}_c * R^B{}_c$ & $S_0 = R^F{}_0 - R^B = 0.$

If instead, only the allocation decision were implemented for that period, then the fund return would change so that the weights of each category are changed from the benchmark weights to the actual fund weights, changing the active return:

$R^F{}_1 = \Sigma_c W^F{}_c * R^B{}_c$ & $S_1 = R^F{}_1 - R^B = \Sigma_c (W^F{}_c - W^B{}_c) * R^B{}_c.$

If the investment process went even further and subsequently made changes to the fund within each category to create its actual holding, then the fund return and corresponding active return would become:

$R^F = \Sigma_c W^F{}_c * R^F{}_c$ & $S = R^F - R^B = \Sigma_c (W^F{}_c * R^F{}_c - W^B{}_c * R^B{}_c).$

Evaluating a Decision's Contribution to the Active Return

Following the approach of decision evaluation, define the allocation value to be the change in the active return brought about by implementing the allocation decision. Thus, the arithmetic allocation value is the arithmetic difference between what the active return was immediately before the implementation of the allocation decision to what it was immediately after the implementation of the allocation decision:

$All = S_1 – S_0 = \Sigma_c (W^F{}_c – W^B{}_c)*R^B{}_c = \Sigma_c (W^F{}_c – W^B{}_c)*(R^B{}_c – R^B)$.

Appropriately, the method of calculating the effect of the allocation decision is arrived at by requiring that its formulation instantiates its definition at the level at which we intend it to be meaningful. Thus, the value assigned to Allocation is the answer to question: By what value did the active return change due to the implementation of the allocation decision? Or: By what value did the active return change from immediately before the implementation of the allocation decision to immediately after the implementation of the allocation decision? Thus, this method of calculating the value assigned to allocation is not justified by only imposing requirements at the portfolio level.

Again following the approach of decision evaluation, define the selection value to be the change in the active return brought about by deciding to implement the selection decision. Thus, the arithmetic selection value is the arithmetic difference between what the active return was immediately before the implementation of the selection decision to what it was immediately after the implementation of the selection decision:

$Sel = S – S_1 = (R^F – R^B) – (R^F{}_1 – R^B) = (R^F – R^F{}_1) = \Sigma_c W^F{}_c*(R^F{}_c – R^B{}_c)$.

It is noted that at the decision level at which All and Sel are defined, there is no freedom to add a free parameter since their values are completely specified by their definition. It is also noted that, while All comes out identical to All^{BF}, there is no Int. This is because decision evaluation only leads to effects of actual decisions and “Interaction” is not a decision. (This clarifies the situation for the formally created Interaction term that has caused so much angst in the field.) Sel does end up equaling the sum of Sel^{BF} and Int^{BF}. This makes it coincide with a common approach in standard performance attribution where Sel^{BF} and Int^{BF} are combined but for reasons that are not those of the argument given here.

The conclusion so far from the analysis of this simple two-step investment process is that decision evaluation provides an improved justification for the performance attribution results commonly reported. However, we will next see that decision evaluation leads to very different results for risk-adjusted performance attribution than is standardly proposed.

Evaluating a Decision’s Contribution to the Active Risk-Adjusted Return

Rather than introduce a new method for the decision evaluation for the active risk-adjusted return, we employ the exact same method of decision evaluation that was applied for active return.

For each period, before the first investment decision is implemented, the fund return for that period is identical to the benchmark return for that period, $R^F{}_0 = R^B$. Thus, at this point in the decision process for each day the active return is equal to zero, $S_0 = 0$. Therefore, at this point in the decision process for the time series of days the information ratio is equal to zero, $IR_0 = 0$.

Immediately after the implementation of the allocation decision for each day in the time series the fund return is $R^F{}_1$ and the active return is S_1. Thus, the information ratio for the time series that would have been achieved if the investment process had ended each day immediately after the allocation decision was implemented is $IR_1 \equiv \langle S_1 \rangle / \sigma(S_1)$.

Similarly, immediately after the implementation of the selection decision for each day in the time series the information ratio $IR \equiv \langle S \rangle / \sigma(S)$.

In a manner exactly analogous to the analysis for performance attribution, define the risk-adjusted allocation value to be the change in the active risk-adjusted return brought about by deciding to implement the allocation decision. Thus,

$IR(All) \equiv IR_1 - IR_0 = \langle S_1 \rangle / \sigma(S_1)$.

Similarly, define the risk-adjusted selection value to be the change in the active risk-adjusted return brought about by deciding to implement the selection decision. Thus,

$IR(Sel) \equiv IR - IR_1 = \langle S \rangle / \sigma(S) - \langle S_1 \rangle / \sigma(S_1)$.

It still follows that

$IR = IR(S) = IR(All) + IR(Sel) = \Sigma_j IR(Xj)$.

Only now, all X_j and $IR(X_j)$ are themselves meaningfully defined at the level of the actual investment decisions that apply to them. There is no opening in the analysis that allows for the imposition of arbitrary terms.

Decision evaluation, thus, always obtains its evaluations by explicitly defining the economic meaning of the most basic terms it employs and then rolls them up to other meaningful concepts. This is starkly different than defining terms at higher levels and them formally decomposing them in ways whose justification allows for indeterminate assessments and creates formally consistent but financially meaningless results.

Since decision evaluation analysis does not allow for the imposition of any arbitrary values, All and Sel are determinate, as are IR(All) and IR(Sel). This is in vivid contrast to All^{BF}, Sel^{BF} and Int^{BF} and $IR^{MHB}(All)$, $IR^{MHB}(Sel)$ and $IR^{MHB}(Int)$ as standardly justified that all allow the inclusion of arbitrary values in their assessment.

If attribution is only a decomposition process, than any decomposition of active return can be paired with any decomposition of the active information ratio. (See Mirabelli, 2000 for an overly long but still incomplete list of the many meaningless ways in which multi-period active return has been decomposed.) However, if attribution is taken to be a financially meaningful method of analysis, by following the approach of decision evaluation, then the very same method can and should be uniformly applied to both the active return and the active information ratio. When this is done, performance attributes and risk attributes can be meaningfully pared in an investment optimization process.

An Example Decision Evaluation Report

To show what the results of such a proper analysis would look like in a more complete and more realist context, below is a partial screen shot of a report by Opturo employing full decision evaluation analysis of performance and two of the many risk and risk adjusted performance measures it addresses for a deeply structured investment decision tree. When interpreting this table, recall that the non-linearity of the standard deviation formula implies that Active Standard Deviation is not Tracking Error,

$\sigma(R^P) - \sigma(R^B) \neq \sigma(R^P - R^B) = TE.$

Decision Tree Attributes	Arithmetic Attributes of: Active Return	Active Standard Deviation	Active Information Ratio
Cumulative Effect	4.01%	0.26%	0.07%
--Invest Effect	0.01%	--	--
--Country Allocation Effect	0.15%	-0.04%	0.00%
--Asset Class Allocation Effect	-0.15%	0.18%	-0.00%
--Asset Class Forking Effect	--	0.09%	0.04%
--EQ GIC Sector Allocation Effect	-0.74%	0.02%	-0.01%
--EQ GIC Industry Allocation Effect	0.62%	0.01%	0.01%
--EQ Market Capitalization Selection Effect	1.21%	0.02%	0.02%
--EQ Issue Selection Effect	1.11%	-0.05%	-0.02%
--FI Duration Selection Effect	0.02%	0.05%	0.02%
--FI Convexity Selection Effect	0.09%	0.00%	0.00%
--FI Issue Selection Effect	0.18%	-0.05%	-0.02%
--Trade Effect	1.51%	0.04%	0.03%

Conclusion

Attribution is not simply the decomposition of a portfolio-level property into individual terms that are each ‘associated’ with individual decisions. Rather, at the fund’s decision level, it appropriately is the determination of the amount by which each decision-type affects a portfolio-level property. This process could and should be done both arithmetically and geometrically in order to get the full view of decision structures that are at all more complicated than the simplest one addressed in the body of this article. Also, it is important to apply decision evaluation at levels below the portfolio level, such as, in the decision structure in the example discussed above, for each independent parallel selection decision done within each allocation category, where there is no ordering between the selections within different categories.

One way of summarizing this paper is that portfolio level restrictions are sometimes insufficient for the specification of component level analysis. A stronger way of phasing the conclusion is that the introduction of a term into a dependable model needs to be accompanied by a meaningful question to which the term in isolation is a precise answer.

References

Bertrand, Philippe, "Risk-adjusted Performance Attribution and Portfolio Optimizations under Tracking-Error Constraints," *Journal of Asset Management* (2009) **10,** 75–88.

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

Brinson, Gary, P., L. Randolph Hood and Gilbert L. Beebower, "Determinants of Portfolio Performance," Financial Analysis Journal, July-August 1986, 39-44.

Menchero, Jose G., "Risk-adjusted Performance Attribution," *Journal of Performance Measurement* (2007) **11**, 22-28.

Menchero, Jose G. and J. Hu, "Portfolio Risk Attribution," *Journal of Performance Measurement* (2006) **10**, 22-33.

Mirabelli, Andre, "The Structure and Visualization of Performance Attribution," *Journal of Performance Measurement* (Winter 2000/2001) **5**, 55-80.