

Why Have an Attribution Model to Break Out the Investment Decisions When the Answer is Explicit? Advocating a Decision-based Approach to Attribution

Attribution approaches are based on models and approaches that do not reflect the modern portfolio system or computing capability. Historically, all of the decisions made in a portfolio are consolidated into a total portfolio, with an attribution model then attempting to determine which decisions were made and which added positive returns. A whole industry has subsequently built up debating the resulting complexity and minutia: arithmetic vs. geometric; how to handle interaction; multi-period linking; etc. This article presents an alternative approach: don't aggregate and then attempt to disaggregate; just use the trades themselves, as every decision is executed in the trades. The objective of this article is to provide a detailed view on how this can provide a dynamically different and viable methodology.

Jem Tugwell

is a specialist Investment Management Consultant, focusing on institutional investment management strategy and analytics. Jem's analytical focus involves how to represent different asset classes inside portfolios, the separation of value and exposure, as well as pricing models and risk measurement across all asset classes and their associated derivatives. The founding director of Jem Tugwell Associates, Jem has more than 20 years' expertise in the investment management sector; across the buy-side and sell-side of the business.

INTRODUCTION

Attribution: A Definition

One definition of performance attribution is “a set of techniques that performance analysts use to explain why a portfolio's performance differed from the benchmark” (Wikipedia). This explanation of portfolio performance is important, as it is the embodiment of a fund management firm's collective skill and investment process. If the attribution analysis reflects the investment decisions, then it can have many valuable uses. The fund managers can use it to demonstrate their added value, clients can be kept accurately informed about their fund's performance, while sales and marketing teams can use the reports to help them win business. As Spaulding (2009) clearly states: “The basic rule of attribution is that the approach we use should be related to the investment strategy; otherwise we may be analyzing the wrong things.” We want to focus on the decisions that were made to determine if they were effective or not. If the attribution measures decisions that were not made or mixes different decisions up, then the analysis becomes devalued, risking being a mere “tick the box” report that is not actively used. This means that attribution can sometimes be seen as the “end” in its own right. However, it is also an extremely important part of the overall running of

an investment firm.

The Importance of Attribution

Attribution provides fund management firms with one of the key tools for winning new business and retaining existing clients. Holappa (2008) claims that fund sponsors continue to look for timely and more precise performance calculations. For an institutional investor to appoint an active manager, they must first believe that active fund management can outperform the benchmark. While there have been many studies into performance persistence, many suggest that active returns do not persist. Studies by Sharpe (1968) and Jensen (1968) found that superior performance does not persist over time, and Blake *et al.* (2002) showed that the performance of U.K. pension funds is close to the benchmark. If this is universally the case, then there is no need for performance attribution, as the active return will always be zero.

By looking at the actual published performance return figures, we can see that specific funds and fund managers do outperform their benchmarks over specific periods. Harmstone's study (2000) agreed, saying that active management is not a failure across the board and that actively managed portfolios successfully outperform their benchmarks in certain asset classes. For the

institutional investor, it is paramount to have belief in their chosen management firm as an “out-performer” before investing funds. If the attribution reports clearly and accurately describe the decisions made by a fund management firm, the skill of the firm and repeatability of the active return should be obvious to the investor, and the attribution reports will be a central part of their “beauty parade” of fund management firms.

Traditional Attribution Approaches

Most of the available attribution models are based on, or have their roots in, the seminal work of Brinson, Hood and Beebower (1986) and Brinson and Fachler (1985). These, along with many others, split a portfolio into “allocation decisions” and “selection decisions.” These models are applied to the aggregate holdings in the portfolio relative to the benchmark and introduce the concept of interim or notional portfolios into the analysis. This approach continues with newer models proposed by Colin (2007) amongst others.

The main problem with these models is that they only have a limited number of decision “types,” *e.g.*, allocation and stock selection, and that multiple decisions of the same type are not considered separately unless they are in different sectors. For example, the performance of two separate decisions to overweight equities will be explained by one net overweight decision. Adding different decision types requires a new model. The other problem of “notional” portfolios is that it describes the active returns in abstract terms, and this can then introduce other “decisions” that were not made; for example, interaction. No fund manager takes an interaction decision, and it is therefore misleading to use “decisions” that were not taken to describe the active return. Indeed, Bacon (2004) explains that interaction is not well understood, presumably because it is not intuitively part of the investment decision.

There is a fundamental problem with this approach, in that there are a huge number of reasons why a fund manager makes a decision. It could be an allocation decision, a stock-picking decision, a value decision, a growth decision, a duration play, etc. The decisions also vary significantly by portfolio type and firm. Cash, equity, and fixed income funds will all have different decision types. Expecting an attribution model to be flexible enough accurately to break out all possible

decisions is an inherently flawed concept, as it would require a different model virtually per fund manager and fund.

It is also fair to say that the performance and attribution world is focusing considerable time and effort into problems and issues not directly related to explaining the investment decisions. The arithmetic versus geometric debate still continues; consideration of the linking of periods has resulted in many approaches, including Frongello (2001), Menchero (2000 and 2004), and Broberg (2006), among others. Interaction itself has warranted detailed discussion (Spaulding 2008), as have sectors with zero weights (Laker 2006). While this detailed research has added to the knowledge on attribution, it has not addressed the key issue of making sure that the attribution model accurately describes all of the investment decisions.

Status Quo Objection 1: Allocation Explains Everything

One reason to reject the model proposed in this paper is the status quo view that asset allocation (a portfolio-level activity) drives the vast majority of relative performance. This is supported by Brinson, *et al* (1991), finding that investment policy explained, on average, 91.5% of the variation in quarterly total plan returns. Other studies, including Sharpe (1992), Brinson, Singer, and Beebower (1991), and Hensel, Ezra and Ilkiw (1990), are also quoted to support these findings. However, as highlighted by Ibbotson and Kaplan (2000), the original question answered by Brinson *et al.* was about variation in returns, not about which decisions added value.

Consider a fund that for simplicity only invests in three listed funds (equity, cash and fixed income), with the following initial policy and adjustments.

Sector	Q1	Q2	Q3	Q4
Cash	10.0	12.0	12.0	12.0
Equity	45.0	44.0	45.0	46.0
Fixed Income	45.0	44.0	43.0	42.0
Total	100.0	100.0	100.0	100.0

When we look at the whole year, we would expect that allocation will explain all of the return variation to the benchmark. If this is the question being asked, then the

status quo answer is all that is needed. However, if the questions are more analytical and in-depth - for example: which decision added/lost the most?; which fund managers make decisions that more consistently add value?; did investing a large inflow mid-Q2 to the policy weights make sense in a falling market? - it therefore follows that asset allocation will not, nor was it originally proposed to, answer these questions. The applicability of Allocation also varies by fund type: Money Market fund managers do not make allocation decisions; fixed income funds also require different factors as demonstrated by the numerous fixed income attribution models available.

Fundamentally, a client looking to appoint a fund management firm, or a CIO looking to decide who the best fund managers are, need answers to much more detailed questions about each and every decision that was made on the fund. To do this, we therefore need a fundamentally different approach to attribution, which breaks out each actual decision that was made, works across all fund types, and fits the different investment processes.

A DECISION-BASED APPROACH

Approach Overview

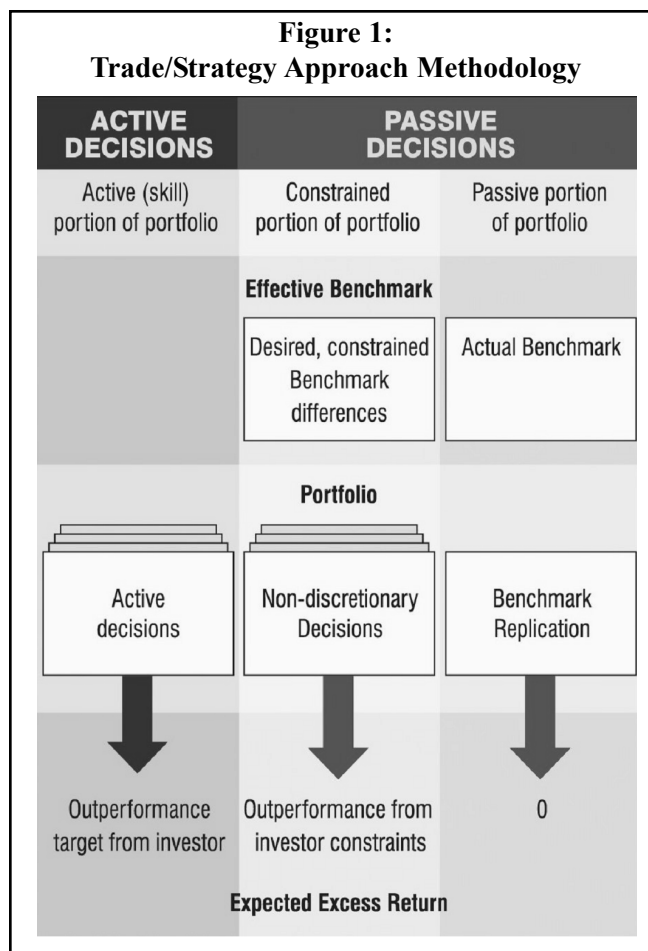
The first step is to recognize that we need to separate each investment decision and its reason. This means that we must measure each active decision, as well as any nondiscretionary decisions and any benchmark replication errors. We cannot therefore use aggregate holdings.

Status Quo Objection 2: Individual Trades Don't Add Up to the Whole.

One possible objection when looking at the individual trades is that it assumes that the sum of the parts equals the whole – something that diversification and modern portfolio theory tells us is not the case. However, the return of all of the trades will equal the return on the fund over any period. Another objection is that there will be many tactical trades or “housekeeping” trades that would be unclassified and/or classified as “miscellaneous.” It is important to note that the proposed approach attempts to measure each decision, as opposed to each trade. A decision could be implement-

ed as one trade (e.g., tactical purchase of an equity), but is more likely to involve multiple trades (e.g., invest inflow to the policy, sell out of one sector into another, buy bond and hedge interest rate risk with a second). If the fund really does have a significant number of miscellaneous decisions, then the use of the proposed approach is valuable. This is because it allows the client to question the CIO as to why there are so many miscellaneous decisions when the firm’s sales material claims how well understood and defined the investment process is. A significant number of tactical decisions would also be separated out by the proposed approach, allowing debate as to whether they add value and which fund managers are good at them, etc. Again, the proposed approach is valuable, as it provides the necessary granularity to answer these types of questions.

The decisions can be split as shown below:



The performance of the funds’ benchmark replication should be zero. Any differences would generate unexpected relative returns that are not a result of active

decisions. The measurement of this requires that all positions held for benchmark replication are explicitly tagged. Any unexpected relative return could be due to the manager's lack of skill in indexing, or that the benchmark violates one of the cardinal rules of benchmark selection; namely that the benchmark return must be exactly replicable. Large, broad benchmarks, such as Lehman Global Aggregate, are extremely difficult for a fund to replicate in every exposure and risk dimension and, therefore, some undesired relative return is inevitable.

From the diagram, it can be seen that there are two sources of expected excess return: from the active decisions and from any investor constraints. Any constraint decisions will need to be implemented where the benchmark chosen conflicts with the investor constraints. For example, an investor may require no tobacco stocks, but the benchmark may contain them, or the investor may want to be 50% currency hedged, or constrained on their exposure to banks, etc. The unwanted exposure will need to be "sold" and this will reduce the effectiveness of the "effective benchmark" replication. It is important that the performance of these constraint decisions is separated from the manager's decisions, as they are really part of the passive benchmark replication return.

The active decisions represent the implementation of the management firm's skill.

A WORKED EXAMPLE

For our analysis, we consider a benchmark of 60% iBoxx GBP Corporate Bond index and 40% FTSE All Gilt. The iBoxx index had 1,100 holdings in February 2009 and the FTSE held 32 U.K. Government bonds (Gilts). The iBoxx index is a broad U.K. index, covering 386 issuers across over 40 industry groups.

Let us consider three separate funds running against this benchmark.

- Fund 1: benchmark replication fund unable to match the benchmark exactly due to fund size constraints.
- Fund 2: benchmark replication fund with no fund size constraints, but a 10% constraint on banking stocks.

- Fund 3: As fund 2, but actively managed.

Considering Fund 2, the blended benchmark contains just over 19.3% in banking stocks.¹ Fund 2 is constrained to 10% of the fund in banks. It has "sold" the 9.3% of banks into cash to fit its constraint.

To illustrate the difference between a "traditional" Brinson style total portfolio view and our decision-based approach, consider Table 1, which shows traditional asset allocation weights.

Table 1

Sector	Benchmark	Fund 1	Fund 2	Fund 3
Cash			9.3	9.3
Corp	60.0	59.5	50.7	52.7
Govt	40.0	40.5	40.0	38.0
Total	100.0	100.0	100.0	100.0

Assuming no stock selection differences, just allocation, this gives us the following returns:

Table 2

Sector	Return	Benchmark	Fund 1	Fund 2	Fund 3
Cash	1%	0.00%	0.00%	0.09%	0.09%
Corp	6%	3.60%	3.57%	3.04%	3.16%
Govt	2%	0.80%	0.81%	0.80%	0.77%
Total		4.40%	4.38%	3.93%	4.02%
		Active Return	-0.02%	-0.47%	-0.38%

and attribution:

Table 3

	Fund 1	Fund 2	Fund 3
Cash	0.00%	-0.32%	-0.32%
Corp	-0.01%	-0.15%	-0.11%
Govt	-0.01%	0.00%	0.05%
Total	-0.02%	-0.47%	-0.38%

From these results, we would conclude that none of the funds has added value, as they all have negative active returns. Fund 1 had the best (least bad) decisions, as it only underperformed by 2bp and the following interpretation of the attribution.

Note that in none of the funds does the traditional attribution view match the actual bets placed (see Table 4).

In our example, Fund 3 has two active decisions: a deci-

Table 4

Fund	'Traditional' Interpretation	Actual Bets
Fund 1	Appears to have bets (underweight Corp and Overweight Govt)	No active bets, just a benchmark replication error due to the fund's size.
Fund 2	Underweight Corp, overweight cash	No active bets underweight corporate is due to a client constraint
Fund 3	Underweight Corp (-11bp) and Govt (+5bp), overweight cash (-32bp)	Mixture of active bets and client constraint

sion to buy 1% of the fund in Technology and a decision to buy 1% in Industrials. Being a bond fund, buying the corporate bonds necessary to implement the decision has a duration impact on the fund. To offset this, the duration impact of the corporate bond purchase is hedged using government bonds. The active decisions, therefore, each have two separate elements that must stay together to measure the decisions accurately. As Table 4 shows, the traditional approach actually gives a positive contribution to the underweight Govt "decision" and suggests that it is a good one. However, from the fund manager's perspective, there is no decision on Govt sector at all, only a need for an interest rate hedge as an intrinsic part of each of the actual corporate decisions.

Consider a Trade/Strategy representation using one level of decision "tags" on the trades used to implement

the decisions (see Table 5).

The difference in the funds, their constraints, and also the active investment strategy is completely transparent. This is particularly true of Fund 3, where the traditional approach implied a decision to underweight Govt bonds, but here the use of the government bonds as two separate hedges is clear.

Table 6 shows the return contributions (to two decimal places) in each fund.

With this approach, we have not performed any traditional attribution or used any model to separate the decisions. Despite this, it is clear that:

Fund 1 has no active decisions, other than that the fund is unable to accurately replicate the benchmark. This

Table 5

Trades/Strategy	Sector	Benchmark	Fund 1	Fund 2	Fund 3
Benchmark	Corp	60.0	59.5	60.0	60.0
Replication	Govt	40.0	40.5	40.0	40.0
Total Replication		100.0	100.0	100.0	100.0
Banking Constraint	Cash				9.3
	Corp				-9.3
Total Constraints					
Total Passive		100.0	100.0	100.0	100.0
Buy Technology	Corp				1.0
	Govt				-1.0
Total Buy Technology					0.0
Buy Industrials	Corp				1.0
	Govt				-1.0
Total Buy Industrials					0.0
Total Active					
Total		100.0	100.0	100.0	100.0

Table 6

Strategy	Sector	Return	Benchmark	Fund 1	Fund 2	Fund 3
Benchmark Replication	Corp	6%	3.60	3.57	3.60	3.60
	Govt	2%	0.80	0.81	0.80	0.80
Total Replication			4.40	4.38	4.40	4.40
Banking Constraint	Cash	1%			0.09	0.09
	Corp	6%			-0.56	-0.56
Total Constraints					-0.47	-0.47
Total Passive			4.40	4.38	3.94	3.94
Buy Technology	Corp	6%				0.06
	Govt	2%				-0.02
Total Buy Technology						0.04
Buy Industrials	Corp	6%				0.06
	Govt	2%				-0.02
Total Buy Industrials						0.04
Total Active						0.08
Total			4.40	4.38	3.94	4.02

enables the fund manager to have a meaningful discussion with their client and their advisors about the practical issues arising from the benchmark selection.

Fund 2 again has no active decisions. It is able to replicate the full benchmark correctly, with the fact that the banking constraint is on the fund and not the benchmark being the cause of the performance difference. Again, a meaningful discussion about benchmark selection can be had.

Fund 3 allows the same discussion about benchmark selection, but clearly shows that the fund manager has added value through their decision, with both the decisions adding 4 bp.

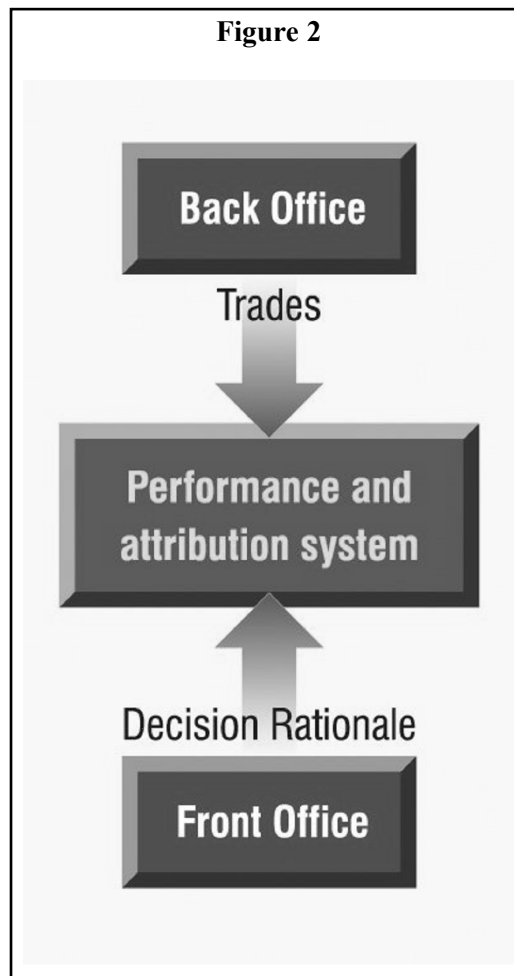
The proposed trade/strategy approach only requires that we know the purpose of the trades (why would there be a trade with no purpose in an active fund?) and that the benchmark replication holdings are identified. The performance is split into the actual trades. The approach also allows for any number of different decision types to be identified: *e.g.*, duration, value, growth, etc., by the trades being identified with the appropriate decision reason. More than one tag could be applied to each trade if necessary – from our example above, the Govt

bonds could be tagged as “hedge,” so allowing an even more granular analysis. The attribution analysis can then be applied to the decisions, enabling a discussion about why the decisions made money, as opposed to which decisions made money.

Getting the Data

A traditional environment, where the source of the data for the performance and attribution is the back office, will have holding and trade data. The rationale for the trades will not normally be stored, as this is regarded as front office information. Considerable work may be involved in ensuring that the decision for each trade is fed from the decision-support system, through the OMS and onto the back office, and then out again to the performance and attribution system. It will also probably involve a considerable upgrade to the back office system, with the associated costs and risks. Perhaps this is one reason why traditional approaches are still the norm, even with their failings in measuring the actual decisions. Indeed, inertia certainly represents an easy rationale behind some fund management firms’ decision making.

However, we now work under the constraints of a post



credit-crunch environment, where fund management firms are increasingly facing more demands:

- on the regulatory front;
- from clients requiring an explanation of exactly how their investment decisions work; precise information on why the fund performed as it did; and why the performance is repeatable.

The essential role of attribution is that it should accurately and clearly demonstrate how each decision added value - something that some fund management firms will capitalize on and achieve a clear competitive advantage as a direct result.

The data issue is therefore something that has to be addressed. However, rather than trying to get the data to and from the back office systems, a more pragmatic and deliverable approach is to recognize the need for a

front office data feed to the performance and attribution system, at least until a version of the back office is available that fully supports the concept of strategies.

In this approach, the trade rationale is fed from the originating system in the front office, together with the trade data as it is executed. The final “books and records” version of the trades is then sent through from the back office along with any corrections, corporate actions, money in and out, etc. If there are differences between the trades sent from the front office and the trades from the back office, then either a manual or automated reconciliation can be performed. The key is that the trades are kept with their original rationale. In this way, the attribution of the investment decisions can be made clearly. The trade/strategy attribution approach is a fully extensible model; when different decision rationales are required, it only requires the trades to be tagged with the appropriate rationale. This can be input by the fund manager at the same time as the trade or

orders are raised.

CONCLUSION

Traditional approaches to attribution, involving taking aggregate holdings and then attempting to disaggregate the separate decisions, are missing the fundamental point of attribution: the attribution analysis reflects the investment decisions. Existing models define a set number of investment decision types *e.g.*, allocation and selection and then attempt to take the aggregate data and fit the data to the predefined reasons. This approach is extremely limiting as it does not accommodate the different approaches necessary for the management of fixed income, equity, and cash funds; it does not separate different decisions of the same type; or allow for complete flexibility in the decision types described. It does not provide the necessary granularity to question fully the value of each decision and the fund managers who made them. This leads to attribution being undervalued or, more precisely, it leads to attribution systems and the reports generated undervaluing the true potential of attribution.

The alternative is to recognize that the attribution should be based on the actual reasons that the investment decisions were made. The decisions made are clear, as the decisions are explicitly described in the trades. In order to realize this approach, you certainly do need the trade rationale data to be made available to the attribution models, and this involves using the front office as a primary data source. The choice is really very simple. Firms can continue with the status quo view, which, as this article has demonstrated, represents an intrinsically flawed concept. Alternatively, they can change to a decision-based approach and, as a result, produce performance and attribution that is truly accurate and useful.

REFERENCES

Bacon, Carl, Practical Portfolio Performance Measurement and Attribution, John Wiley: New York, 2004.

Blake, David, Lehmann, Bruce N. and Allan Timmermann, "Performance Clustering and Incentives in the U.K. Pension Fund Industry," *Journal of Asset Management*, 3(2), 2002, pp. 173–194.

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

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

Brinson, Gary P., Brian D Singer, and Gilbert L. Beebower, "Determinants of Portfolio Performance II: An Update," *Financial Analysts Journal*, May/June 1991.

Broberg, Mikael, "A General approach for linking arithmetic attribution results over time: Another linking algorithm," *The Journal of Performance Measurement*, Fall 2006, pp. 48-54.

Colin, Andrew, "A Brinson model alternative: an equity attribution model with orthogonal risk contributions," *The Journal of Performance Measurement*, Fall 2007, pp. 59-63.

Frongello, Andrew Scott Bay, "Linking Single Period Attribution Results?," *The Journal of Performance Measurement*, Summer 2001, pp. 45-57.

Harmstone, Andrew R., "Alpha transfer: Optimising the benefit of active management," *Journal of Asset Management*, 2000, 196-206.

Hensel, Chris R., D. Don Ezra and John H. Ilkiw, "The Value of Asset Allocation Decisions," *Russell Research Commentary*, March 1990.

Holappa, Guy M., "Private Investments and performance implications from a fund sponsor's perspective," *The Journal of Performance Measurement*, Winter 2008-2009, pp. 36-41.

Ibbotson, Roger G. and Paul D. Kaplan, "Does Asset Allocation Explain 40, 90 or 100 Percent of Performance?," *Financial Analysts Journal*, January/February 2000.

Jensen, Michael, C., "The Performance of Mutual Funds in the Period 1945–1964," *Journal of Finance*, 23(2), 389–416.

Laker, Damien, "Performance Attribution With Zero-Weighted Sectors," *The Journal of Performance Measurement*, Summer 2006, pp. 33-46.

Menchero, Jose, "A Fully Geometric Approach to Performance Attribution," *The Journal of Performance Measurement*, Winter 2000/2001, pp. 22-30.

Menchero, Jose, "Multiperiod Arithmetic Attribution," *Financial Analysts Journal*, 2004, Vol. 60, No. 4, pp. 76-91.

Sharpe, William. F., "Mutual Fund Performance and the Theory of Capital Asset Pricing: Reply," *Journal of Business*, 41(2), 1968, 235-236.

Sharpe, William F., "Asset Allocation: Management Style and Performance Measurement," *Journal of Portfolio Management*, Winter 1992.

Spaulding, David, "Performance Attribution: An Introduction," *The Journal of Performance Measurement*, Spring 2009, pp. 33-41.

Spaulding, David, "Should the Interaction Effect be Allocated? A 'Black Box' Approach to Interaction," *The Journal of Performance Measurement*, Spring 2008, pp. 23-30.

ENDNOTES

¹ Using Bloomberg Industry Group classification scheme.

This article has been reprinted with permission of *The Journal of Performance Measurement*. All rights reserved. Spring 2010.