WHITE PAPER

Performance Attribution: A Powerful Tool for Identifying the Sources of Investment Performance

A Basic Introduction for Asset Managers

Advent Software, Inc.

David Spaulding, The Spaulding Group

ADVENT®

TABLE OF CONTENTS

Introduction: A Powerful Analytic Tool

Performance is an investment manager's calling card, and it's often said that performance attribution has become the hottest area in performance measurement. Attribution has become a standard part of the portfolio manager's analytical tool kit, providing insights into what is working and what is not. It provides a way to communicate where the returns have come from when meeting with prospects as well as existing clients. And it has allowed the performance measurement team to work more closely with the investment group by providing important information on the results of their investment decisions.

This white paper provides an overview of what performance attribution is and how it works and contrasts two of the more popular models to calculate equity attribution. It is designed to be an introduction to the topic, not an exhaustive discussion, but it is hoped that readers will come away with a firm grasp of the fundamental principles of attribution as well as a practical understanding of how they can be used to measure performance. Our focus is solely on equities, since fixed income attribution is a complex topic beyond the scope of this paper.

A Familiar Concept

Attribution, by itself, isn't a new concept. People have engaged in attribution analysis for many, many years; probably centuries. Many individuals are trained in attribution analysis, although they may not know it.

For example, when police officers come upon an automobile accident, they will often conduct an analysis to determine the cause(s) of the accident. Similarly, firefighters typically conduct investigations after a fire to determine the cause. Sports teams regularly engage in a post-game analysis as a way to determine if their strategy was successful or not. And you're no doubt aware of many "Monday morning quarterbacks" who conduct their own analysis following the prior weekend's sporting events.

The basic concept of attribution is therefore something we see on a fairly regular basis. In the case of police and fire officials, they're analyzing something that occurred independently of their actions. For sports teams, they no doubt had some strategy that they developed in advance of the event and now want to review the game to determine what worked and what didn't. The process may be summed up in three words: strategy, execution, analysis.

When it comes to conducting analysis in the investment arena, our approach is quite similar to what a sports team does: the portfolio manager(s) usually starts with a strategy, then implements or executes the strategy, and then wants to conduct analysis to determine if the strategy worked. We call this analysis "attribution."

Relative versus Absolute Attribution

One of the basic rules of attribution is that the approach we use should be related to the 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. The decisions produce "effects," and consequently we want to analyze the effects that result from the decisions. We call these "allocation effects."

We can split attribution into many different categories. For the purpose of this paper, we will focus primarily on relative attribution. We will now briefly contrast this approach with absolute attribution, which is also referred to as "contribution."

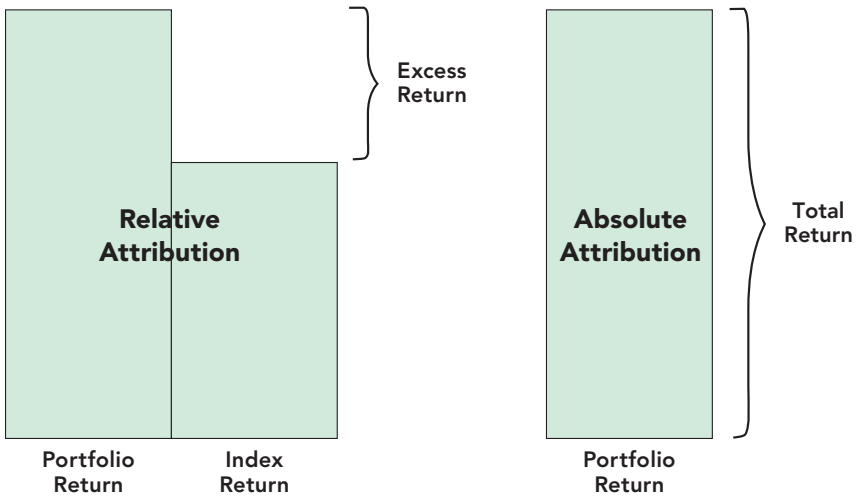

Figure 1. Relative versus Absolute Attribution

In both cases, we are attempting to reconcile to some number: for relative attribution, we're reconciling to the excess return (i.e., the portfolio minus the benchmark or index return); for absolute attribution, we're reconciling to the total return. In each case, we want to understand where that number came from, what contributed to it as well as what detracted from it.

Absolute attribution is the typical source for a manager's top ten or bottom ten investments, for example. Many managers want to highlight their winners and losers. Our focus in this paper will be on relative attribution: we will demonstrate a process to understand where the excess rate of return (ROR) came from.

The Effects to Analyze: Allocation and Selection

If we think about equity managers, one typical way to distinguish them is to compare top-down managers with bottom-up ones. In the case of top-down managers, the process often begins with a review of the economy; the managers may also take into consideration the political activities that may be occurring. Based on their projections on what the future may hold, they can make some predictions on how the various segments of the market might respond.

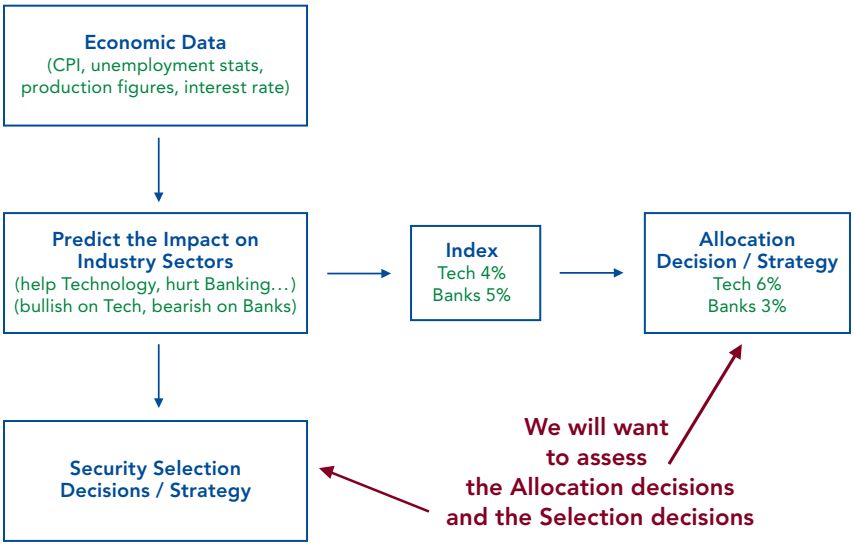

Figure 2. A Top-Down Approach to Investing

If their analysis suggests that certain segments should do well, they may want to overweight them relative to the index they're managing against; on the other hand, if they're projecting that other segments won't fare too well, they may underweight them. These over- and under-weighting decisions are the allocation decisions; thus they're ones we may want to review to determine if they worked or not.

The next step in the process would be to select the stocks in which they want to invest. This selection decision will also be evaluated.

In the case of bottom-up managers, the decision often begins with the identification of securities they will want to purchase as well as those they wish to sell. Thus, we want to analyze the manager's selection decision as we did in the case of the top-down manager. A controversial aspect of attribution deals with whether or not we should also look at the allocation decision. In some cases, bottom-up managers have been known to argue vehemently in opposition to using a top-down model. And, if you think about our rule, one might argue against its use, suggesting that we should only focus on the selection decision.

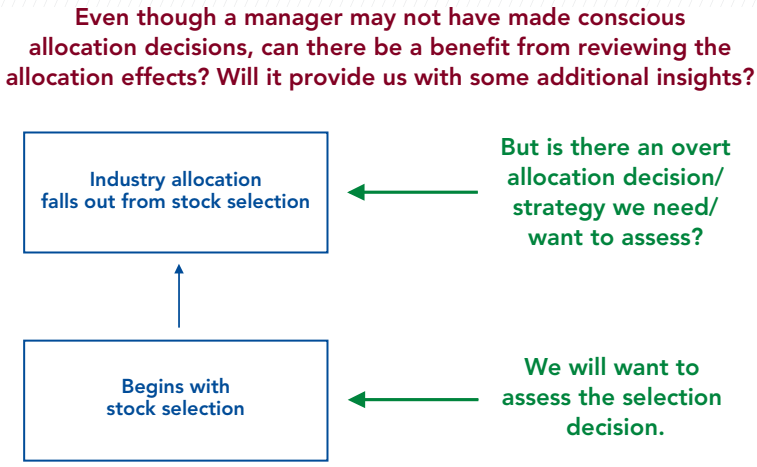

Figure 3. A Bottom-Up Approach to Investing

While it may be true that the manager didn't make a conscious allocation decision, there will no doubt be some differences in the way the securities are allocated in the portfolio relative to the index; in addition, these allocation differences can have an impact on the manager's return. Consequently, we generally recommend that even bottom-up managers take a look at the allocation effect.

How do we analyze these effects? We employ models. Models are mathematical formulas which take into consideration various characteristics of the portfolio and its related index to both reconcile to the excess return as well as to apportion the excess return across two or more of our effects. Our experience and research have shown that the two most popular models for equity attribution were developed by Gary Brinson, along with various co-authors.[1] While these models are quite similar, there is a key distinction between them.

The Brinson Hood Beebower Model

The Brinson Hood Beebower (BHB) model was published in 1986 and has been quite popular (BHB (1986)). This model has three effects, two of which we've already discussed: allocation and selection. We will discuss the third effect shortly.

The allocation effect measures the impact of the manager's over- and under-weighting decisions. The formula is:

$$Allocation_{BHB} = (w_P - w_B) \times (r_B)$$

where:

w_P = the portfolio's weight in a sector
w_B = the benchmark's weight in the same sector
r_B = the benchmark's returns for the given sector

1. See TSG (2007), page 23.

The difference in weights represents the allocation decision. For example, if we were bullish on a particular segment, we would presumably want to overweight it in anticipation of it having good performance; on the other hand, if we were bearish we would probably want to reduce our exposure to it relative to the index, and therefore underweight it.

Technology = 2% ROR	
	BHB
Overweight	Positive
Underweight	Negative
Technology = –2% ROR	
Overweight	Negative
Underweight	Positive

Figure 4. Assessing the Allocation Decision

For example, in Figure 4 we see in one case that technology had a positive return. If we were to overweight it, our portfolio would have more invested in the sector than the index. This would result in a positive weighting factor and multiplying this by the positive return would yield a positive result. Here, we would basically be saying that we're being rewarded for overweighting a sector that had good, positive performance.

On the other hand, if we had chosen to underweight the sector, our portfolio's weight would be below the index. This would mean the difference would be negative and multiplying this difference by our positive return would yield a negative result. This can be interpreted as a negative assessment of our decision to underweight a positively performing sector (i.e., we failed to take advantage of a sector that had a positive return).

In the second case, our return is negative. If we had overweighted it, we would have a positive weighting (our weight being greater than the index's) being multiplied by a negative return. This will yield a negative result, suggesting that our overweighting of a poorly performing sector wasn't a good idea. If we underweighted it, we'd have a negative weight difference being multiplied by a negative return. This results in a positive effect, thus rewarding us for avoiding having too much of our investment in a poorly performing sector.

Our second effect is selection. The formula we use for selection is:

$$Selection_{BHB} = (r_P - r_B) \times (w_B)$$

where:

r_P = the portfolio's return in the sector

The first factor represents our skill at picking better securities than the benchmark. If our return is higher, then our resulting effect (after multiplying this value by the benchmark's weight) will yield a positive number. If our selection wasn't as good as the benchmark, then the difference will be negative, resulting in a poor score for our selection.

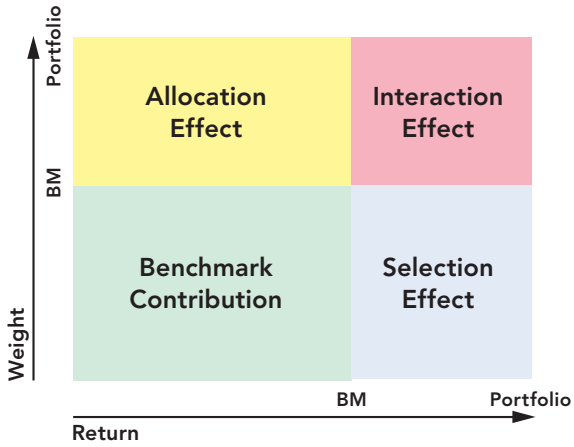

Figure 5. A Graphical View of the Effects

Figure 5 shows an attempt to portray these effects visually. The outside box represents the portfolio's total return; we show within this box the contribution that comes from the benchmark. Recall that with relative attribution we wish to account for the difference between the portfolio and the index. You can see how the allocation and selection effects contribute to this difference; however, there remains a bit that isn't accounted for by either of these effects. While the authors of the BHB model referred to this effect as simply "other," the industry has generally settled on the term "interaction." The formula for interaction can easily be derived by looking at Figure 5 and is:

$$Interaction_{BHB} = (w_P - w_B) \times (r_P - r_B)$$

If you refer back to our other two formulas, you can see that this formula includes the allocation factor (the difference in weights) as well as the selection factor (difference in returns); thus, it's the combination of these two decisions. While this may seem like a reasonable explanation, the reality is that the interaction effect is quite controversial (see, for example, Laker (2000), Spaulding (2003/2004), Campisi (2004), and Spaulding (2008)). Some contend that the effect should be included with selection (or perhaps even allocation), while others argue that it should be shown separately. Space isn't adequate to give this topic enough attention, so if you're curious about the arguments, we suggest you consult one of the aforementioned references.

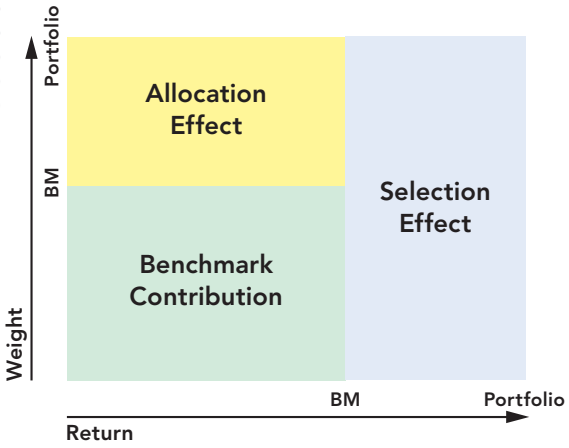

Figure 6. A Graphical View of the Effects without Interaction

If you choose to, you can eliminate the interaction effect by replacing the benchmark's weight with the portfolio weight in the selection formula:

$$Selection_{BHB} = (r_P - r_B) \times (w_P)$$

As can be seen in Figure 6, if we elect to do this we are essentially combining interaction with selection. Again, we won't discuss the pros or cons of this approach, but simply acknowledge that it's an option.

Applying the BHB Model

Now let's apply the BHB model to a portfolio and its index. Figure 7 shows the weights and returns for our benchmark and index as well as the ten sectors that comprise both. As you can see, the portfolio had a return of 0.29% while the index's return is 0.19%, meaning we have an excess return of 0.10% or 10 basis points. We want to reconcile to this value.

For illustration purposes, we'll begin by looking at the Basic Materials sector. We can see that the manager underweighted a positively performing sector; recall that this should yield a negative effect (refer to Figure 4). Our portfolio outperformed the index, so our selection effect should be positive. And because we have an underweighting factor from the allocation being multiplied by a positive selection factor, our interaction effect will be negative.

	ROR		Weight	
	Portfolio	**Index**	**Portfolio**	**Index**
Basic Materials	0.25%	0.15%	10%	11%
Industrials	0.50%	0.51%	11%	9%
Consumer Cyclical	1.00%	1.01%	8%	7%
Utilities	–0.80%	–0.75%	12%	13%
Energy	2.00%	1.95%	7%	5%
Financials	–0.30%	–0.31%	6%	8%
Healthcare	0.80%	0.79%	15%	13%
Technology	0.60%	0.70%	9%	10%
Telecommunications	–0.20%	–0.21%	13%	10%
Consumer, Non-Cyclical	–0.50%	–0.52%	9%	14%
Portfolio	**0.29%**	**0.19%**	**100%**	**100%**

Figure 7. The Scenario We'll Analyze

Our results:

$$AllocEffect = r_B \times (w_P - w_B) = 0.0015 \times (0.10 - 0.11) = -0.002\%$$

$$SelEffect = w_B \times (r_P - r_B) = 0.11 \times (0.0025 - 0.0015) = 0.011\%$$

$$InteractionEffect = (w_P - w_B) \times (r_P - r_B) = (0.10 - 0.11) \times (0.0025 - 0.0015) = -0.001\%$$

Hopefully these results make intuitive sense. Figure 8 shows the model applied to the entire scenario. As you can see, we successfully reconciled to our excess return.

We mentioned that if you'd prefer not to see interaction, we could substitute the portfolio weight for the benchmark's weight in the selection effect. If we do this, our allocation effect will be identical to what we see in Figure 8; the only difference will be the combining of the interaction and selection effects, as we see in Figure 9.

	ROR		Weight		Effects			
	Portfolio	Index	Portfolio	Index	Allocation	Stk Sel	Interaction	Total
Basic Materials	0.25%	0.15%	10%	11%	–0.002%	0.011%	–0.001%	0.009%
Industrials	0.50%	0.51%	11%	9%	0.010%	–0.001%	0.000%	0.009%
Consumer Cyclical	1.00%	1.01%	8%	7%	0.010%	–0.001%	0.000%	0.009%
Utilities	–0.80%	–0.75%	12%	13%	0.008%	–0.007%	0.001%	0.001%
Energy	2.00%	1.95%	7%	5%	0.039%	0.003%	0.001%	0.043%
Financials	–0.30%	–0.31%	6%	8%	0.006%	0.001%	0.000%	0.007%
Healthcare	0.80%	0.79%	15%	13%	0.016%	0.001%	0.000%	0.017%
Technology	0.60%	0.70%	9%	10%	–0.007%	–0.010%	0.001%	–0.016%
Telecommunications	–0.20%	–0.21%	13%	10%	–0.006%	0.001%	0.000%	–0.005%
Consumer, Non-Cyclical	–0.50%	–0.52%	9%	14%	0.026%	0.003%	–0.001%	0.028%
Portfolio	**0.29%**	**0.19%**	**100%**	**100%**	**0.100%**	**0.001%**	**0.000%**	**0.102%**

Figure 8. Applying the BHB Model to the Scenario

	Alternative		BHB		
	Stk Sel	Alloc'n	Alloc'n	Stk Sel	Interaction
Basic Materials	0.010%	–0.002%	–0.002%	0.011%	–0.001%
Industrials	–0.001%	0.010%	0.010%	–0.001%	0.000%
Consumer Cyclical	–0.001%	0.010%	0.010%	–0.001%	0.000%
Utilities	–0.006%	0.008%	0.008%	–0.007%	0.001%
Energy	0.004%	0.039%	0.039%	0.003%	0.001%
Financials	0.001%	0.006%	0.006%	0.001%	0.000%
Healthcare	0.001%	0.016%	0.016%	0.001%	0.000%
Technology	–0.009%	–0.007%	–0.007%	–0.010%	0.001%
Telecommunications	0.001%	–0.006%	–0.006%	0.001%	0.000%
Consumer, Non-Cyclical	0.002%	0.026%	0.026%	0.003%	–0.001%
Portfolio	**0.002%**	**0.100%**	**0.100%**	**0.001%**	**0.000%**

Figure 9. BHB with and without Interaction

The Brinson Fachler Model

The year before the BHB article was published, *The Journal of Portfolio Management®* published an article by Gary Brinson and Nimrod Fachler (Brinson Fachler (1985)). Although this model is older, many seem to favor it over the BHB.[2] The Brinson Fachler (BF) model is quite similar to the BHB, with one significant difference: how the allocation effect is calculated:

$$Allocation_{BF} = (w_P - w_B) \times (r_B - R_B)$$

where:

R_B = the overall benchmark return

In the case of the BHB approach, we multiplied our weighting differences by the benchmark's sector return. Here, we multiply the weighting difference by the benchmark sector's relative return, where the sector's return is relative to the overall benchmark return. The differences can be significant. Let's consider the following input values:

Portfolio weight: 8%
Benchmark weight: 2%
Sector return: –1.5%
Overall benchmark return: –3.5%.

As you can see, we decided to overweight this sector relative to the benchmark; however, the sector's performance is negative. As we previously saw, the BHB model doesn't like the idea of overweighting a negatively performing sector. But this sector actually outperformed the overall benchmark. So, what are the results? The BHB will calculate a –0.09% allocation effect while the BF will give us a positive 0.12% effect, meaning a 21 basis point swing, from negative to positive! We invite you to try the math yourself to see if you can come up with the same results.

Because the BF views allocation decisions from a relative basis, it is rewarding overweighted sectors that outperform the overall benchmark return, while penalizing when we overweight sectors that underperform.

Technology = 2% ROR	**Index = 4% ROR**	
	BHB	**BF**
Overweight	Positive	Negative
Underweight	Negative	Positive
Technology = –2% ROR	**Index = –4% ROR**	
Overweight	Negative	Positive
Underweight	Positive	Negative

Figure 10. Comparing How the BF and BHB Models Respond to Allocation Decisions

2. See TSG (2007), page 23.

Comparing the Two Models

Figure 10 compares how the two models respond to the allocation decisions. The BHB views allocation from an absolute positive or negative perspective (being above or below zero) while BF takes the benchmark's return into consideration. If the manager's job is to outperform the benchmark, wouldn't an effective allocation strategy be to overweight those sectors that are the greatest contributors to the benchmark's overall return? That's how the BF model views it.

Figure 11 contrasts these two approaches. On the left we have a panel that ranks the sectors from the highest to the lowest performing; as you can see, the benchmark's return is –3%. Here, the BHB will want the manager to overweight when the return is greater than zero, while the BF wants the manager to overweight any sector that outperforms the benchmark—even Consumer Cyclical (Cons Cyclical) which shows a negative return (but greater than the benchmark).

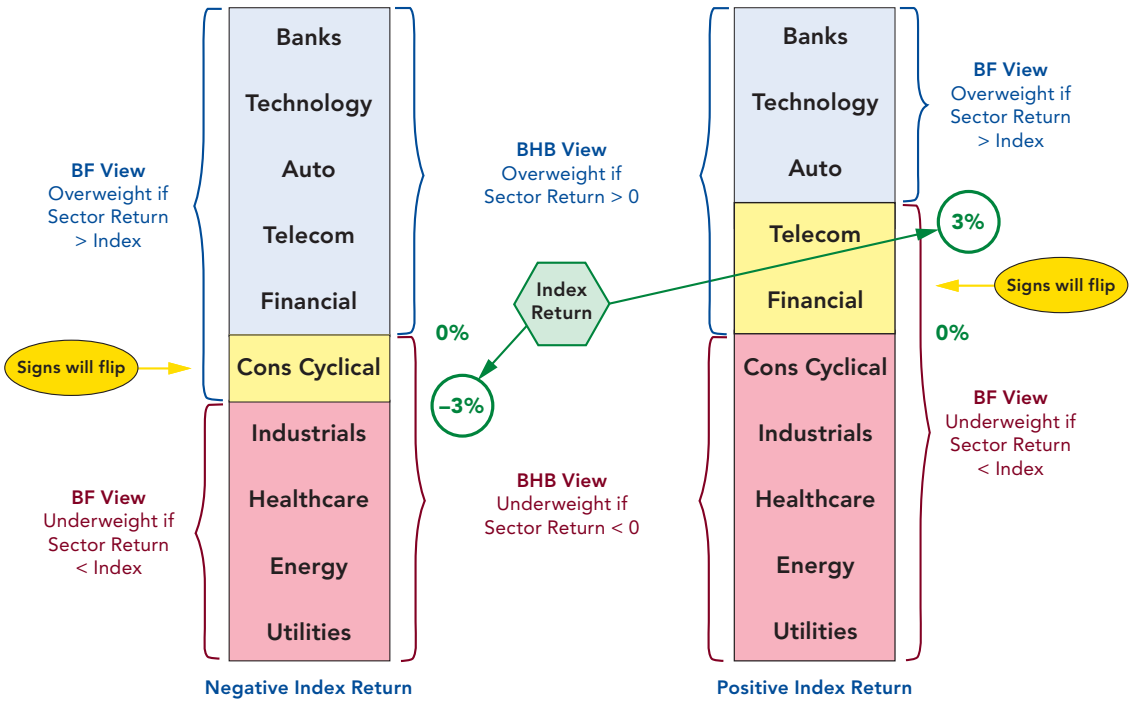

Figure 11. Comparing the BHB and BF Approaches to Allocation

The right-hand panel shows a scenario where we have a positively performing benchmark (3%). The BF model only wants the manager to overweight those sectors that are above this value. Thus, the BF would not want the manager to overweigh Telecommunications or Financials, even though they have positive returns, because they're below the benchmark.

If we return to our portfolio we can calculate the effects using the BF model. Since the selection and interaction effects will be the same as

what we calculated for the BHB model, we will only look at the allocation effects (see Figure 12). While the totals match (which we would expect), the individual sector effects are different, and in the case of Basic Materials, we see the signs flip. This is because the sector's return is positive 15 basis points but falls below the benchmark's 19 basis points.

The decision as to whether to use the BF or BHB models should be based on your view of investing. We've found that today most firms prefer the BF model.

	BF Allocation	BHB Allocation
Basic Materials	0.0004%	–0.002%
Industrials	0.0065%	0.010%
Consumer Cyclical	0.0082%	0.010%
Utilities	0.0094%	0.008%
Energy	0.0353%	0.039%
Financials	0.0099%	0.006%
Healthcare	0.0121%	0.016%
Technology	–0.0051%	–0.007%
Telecommunications	–0.0119%	–0.006%
Consumer, Non-Cyclical	0.0354%	0.026%
Portfolio	**0.100%**	**0.100%**

Figure 12. Contrasting the Allocation Effects

Holdings- versus Transaction-Based Attribution

Before wrapping up this introduction to attribution, we should briefly touch on the issue of holdings- versus transaction-based attribution.

Holdings-based attribution uses the beginning period holdings to calculate the weights that are used in the model, while the transaction-based approach takes into consideration any transactions which may occur during the period (e.g., buys, sells, income) to enhance the accuracy of the results. In addition, some holdings-based models ignore the activity when calculating returns, while the transaction-based will take this activity into consideration. As a result, if there is turnover during the period, the holdings-based approach may (a) fail to reconcile to the excess return, resulting in a residual, which is the difference between the excess return and the sum of the attribution effects, and (b) may provide inaccurate results; while the transaction-based method (a) should reconcile and (b) should produce highly accurate results. The Spaulding Group's research has found that most firms prefer a transaction-based approach.

Whether you're using a holdings- or transaction-based model, it's important that the data is clean. While there are lots of contributors to problems with data, one of the main contributors is the failure to reconcile to the custodian's official books and records.

Conclusion: Putting Attribution to Work

In this brief paper, we've introduced the concept of attribution. We've discussed how the concept is not a new one and that it actually employs a similar strategy-execution-analysis approach that we find in many other areas. We've also noted that model selection is important, because different models can yield different results; thus, understanding how the model behaves is an important part of the decision process. And we've analyzed some examples of how the two most commonly used attribution models differ in their approach as well as in the results they yield. Armed with this information, you should be able to assess the value of performance attribution and how it can help you communicate the sources of performance more effectively to your clients and prospects.

For More Information

If you'd like to learn more about this subject, you may want to consult The Journal of Performance Measurement, the leading resource for new thinking on this subject. The Spaulding Group also publishes a free monthly newsletter that you can sign up for or simply access from our website: www.SpauldingGrp.com. Training classes on performance measurement and attribution are also available.

About the Author

David Spaulding, CIPM, is an internationally recognized authority on investment performance measurement. He's the president of The Spaulding Group, and founder and publisher of *The Journal of Performance Measurement*. He has written three books on performance measurement, was a contributing author to a fourth, and co-edited and contributed to a fifth. He consults to clients throughout the world on investment performance issues and teaches classes on performance measurement and attribution. He earned a BA in mathematics from Temple University, a MS in systems management from the University of Southern California, and a MBA in finance from the University of Baltimore. He is pursuing a doctorate in finance and international economics at Pace University.

About The Spaulding Group

The Spaulding Group is the leading provider of investment performance measurement products and services—performance measurement is our passion. With offices in the New York and Los Angeles metropol-

itan areas, we offer a wide range of services including consulting, GIPS® verification, system design, software searches, training, publishing, conferences, and interactive forums. Through our newest venture, PerformanceJobs.com, we offer recruiting services specifically for the performance measurement industry.

About Advent Software

Advent Software, Inc., a global firm, has provided trusted solutions to the world's leading financial professionals since 1983. Firms in more than 50 countries rely on Advent technology to run their mission-critical operations. Advent's quality software, data, services, and tools enable financial professionals to improve service and communication to their clients, allowing them to grow their business while controlling costs. Advent is the only financial services software company to be awarded the Service Capability and Performance certification for its service and support organizations.

Bibliography

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.

Campisi, Stephen, "Debunking the Interaction Myth," *The Journal of Performance Measurement*, Summer 2004, pp. 63–70.

Laker, Damien, "What is This Thing Called Interaction?," *The Journal of Performance Measurement*, Fall 2000, pp. 43–58.

Spaulding, David, "Demystifying the Interaction Effect," *The Journal of Performance Measurement*, Winter 2003/2004, pp. 49–54.

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

TSG, 2007, "Performance Measurement Attribution Survey." (The Spaulding Group: Somerset, NJ).

ADVENT SOFTWARE, INC.

[HQ] 600 Townsend St., San Francisco, CA 94103 / PH +1 800 727 0605

[NY] 1114 Avenue of the Americas, New York, NY 10036 / PH +1 212 398 1188

[HK] Level 8, Two Exchange Square, 8 Connaught Place, Central, Hong Kong / PH +852 2297 2280

[UK] One Bedford Avenue, London WC1B 3AU, UK / PH +44 20 7631 9240

WPPA1010