White Paper

Identifying A Sponsor's Impact on Total Returns

Performance Attribution for the Total Portfolio

MORNINGSTAR®

Introduction

Performance attribution is a well-recognized quantitative approach to identifying the outcome of investment decisions. Sponsors of pensions, endowments, and foundations view performance attribution as an important tool in their investment manager due diligence process.

However, an investment manager's decisions are only part of a plan's success; the sponsor's own decisions also weigh heavily, but very few sponsors are able apply performance attribution to their own decisions. Also known as macro or balanced attribution, total portfolio attribution enables sponsors to:

- Identify the performance attributable to the strategic asset allocation policy
- Examine the outcome of deliberate deviations from policy weights
- Measure manager-picking skill in aggregate, or by asset class or investment style

This paper will discuss the methodology for conducting total portfolio attribution and outline seven critical challenges sponsors must consider when evaluating their options.

Background

For many years, sponsors and investment consultants have embraced performance attribution as a tool for evaluating investment managers. It is easy to see why: Both an analytical and communications tool, attribution analysis decomposes an investment manager's return versus the benchmark into pieces to explain the impact of decisions an investment manager makes, helping a sponsor distinguish between luck and skill.

Attribution is also one of an investment manager's most powerful tools for understanding and communicating performance drivers to investors and consultants. Attribution analysis has become an essential part of the quarterly performance-update meeting between investment managers and sponsors, identifying and conveying which countries, sectors, and securities were top contributors and which detracted value.

Intimately familiar with the benefits of performance attribution analysis, some sponsors have applied it to evaluate their own investment decisions to describe the value they add to total returns. The consulting field has also broadened the application of such self-evaluation to wealth advisors, managed accounts, and funds of funds (such as target-date funds).

In this paper, we will refer to the analysis of a sponsor's decision-making at the overall portfolio level as "total portfolio attribution."

Decisions in an Investment Portfolio

A sponsor's investment decisions fall into two general categories: strategic asset allocation and active management. The active management portion is further subdivided into tactical asset allocation and manager selection. Total portfolio attribution measures the impact of these investment decisions. The manager selection effect serves as a transition point between total portfolio attribution and micro, or traditional, attribution.

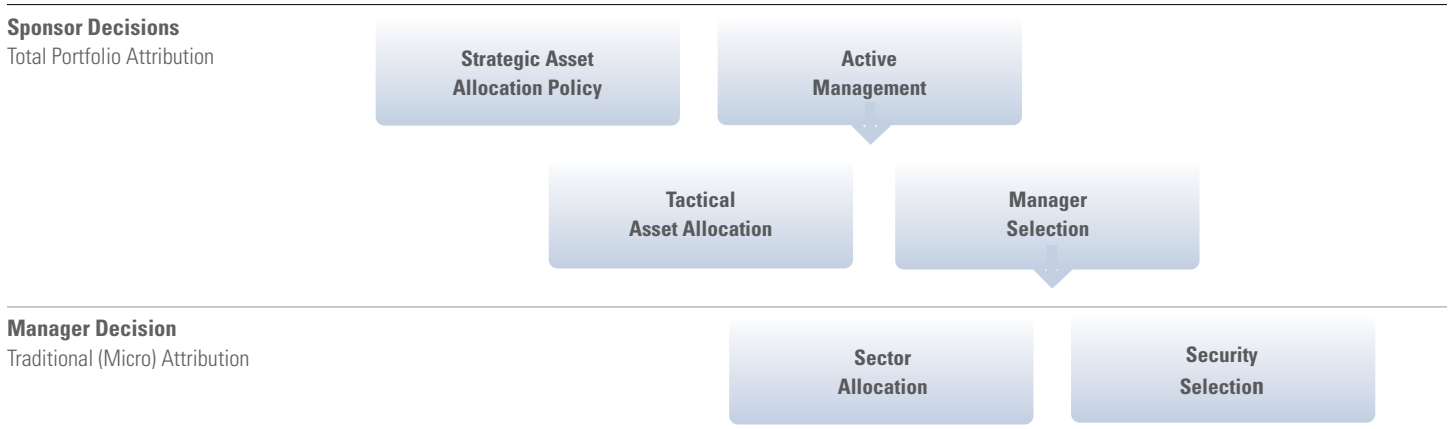

Strategic Asset Allocation
Strategic asset allocation refers to a long-term asset allocation policy based on the investor's objective, risk tolerance, and other constraints. This is usually expressed as target weights in a list of asset classes. In terms of performance evaluation, this is often referred to as the policy benchmark, and it consists of combining policy weights with passive returns from conventional indexes or investable benchmarks. This is the return an investor receives from passive investing, and it serves as an anchor for evaluating a sponsor's skills in active portfolio management.

Active Management:
Tactical Asset Allocation
This refers to a sponsor's intentional deviation from the long-term policy allocation to capitalize on his or her bullish and bearish views on different asset classes. The attribution result that measures the impact of this decision is called weighting or allocation effect. Depending on the attribution model, this term may also be a reflection of the timing of a sponsor's rebalancing decisions. The weighting effect is an informative and intuitive number. Overweighting an asset class that has outperformed the total is a good decision, just as it is a good decision to have been underweight in an asset class that has underperformed. The converse is also true. An attribution analysis pinpoints whether decisions about over- or underweighting particular asset classes added ordetracted value. Aggregating weighting effects from all asset classes provides an overall assessment of a sponsor's skills in tactical asset allocation.

Active Management:
Manager Selection
Another portfolio decision is manager selection—deciding which managers to hire and fire and the timing of such events. In attribution terms, the manager selection effect represents the portion of excess return attributable to the sponsor's skill in selecting active managers who outperform their benchmarks. Intuitively, the manager selection effect is positive when a manager achieves a positive excess return over the designated benchmark. However, an investment manager's contribution to the portfolio may be small even if he or she significantly outperformed his benchmark if the manager was allocated little money. In this sense, total portfolio attribution is always focusing on the actions of the plan sponsor. Taken in aggregate, the manager selection effect provides an assessment of the sponsor's skill in picking active managers. Aggregating manager selection effects at the asset class level can identify whether the plan sponsor is more effective at picking managers in particular asset classes. It could even suggest asset classes that are better off being passively managed.

Challenges to Conducting Total Portfolio Attribution

Total portfolio attribution is a largely underserved topic. Due to the scarcity (and expense) of commercial solutions, sponsors who recognize the power of attribution often resort to proprietary spreadsheets. The limitations of spreadsheets are many, and analysis that can be easily accomplished in a software system can be difficult and error-prone using a spreadsheet.

What's more, sponsors using spreadsheets for total portfolio attribution have found guidance scarce, as there is little coverage in the form of industry or academic literature. This is especially true when compared with micro attribution on equity or fixed-income managers.

Planning for Total Portfolio Attribution

The following seven critical areas should be considered when evaluating options for total portfolio attribution:

1. Data Collection

Both proprietary spreadsheets and software systems require current data. Here is one place that a software system containing public investment data offers a considerable advantage over a spreadsheet. Data gathering is no small task, especially when a sponsor invests in multiple types of investment vehicles such as separate accounts, mutual funds, exchange-traded funds (ETFs), insurance subaccounts, individual equity securities, etc. The benefits of attribution might be minimized or lost entirely if laborious manual data collection and entry prevent regular analysis. In addition, a software system with an easy process to import nonpublic information (such as proprietary funds, private equity, private real estate, etc.) adds value as a central data aggregator even when all the plan's investments are proprietary.

2. Portfolio Changes

A portfolio undergoes changes during a month. Assets might be reallocated among investment managers, the roster of investment managers might change, and cash flow might occur. While midmonth valuation, time-weighted rate of return on subperiods, compounding of attribution results, and manager changes can be easily accomplished in a software system, you risk errors using a spreadsheet. For example, let us assume that a portfolio has recently added a new investment manager, funded with proceeds from a terminated manager and a client contribution. A software application would merely require that the plan sponsor enter the portfolio update date, select the performance stream that represents the new investment manager, enter the funding amount of the new manager and either zero out or remove the terminated manager. The steps are just as easy when the plan sponsor prefers to update data via importing.

In contrast, a spreadsheet requires significantly more steps:

- Insert a new row in the spreadsheet for the new manager
- Update formulas to read the row
- Break the month into two sub-periods
- Determine the percentage allocation to all investment managers
- Collect sub-period returns of all investment managers
- Make two attribution calculations
- Accumulate attribution results

As you can see, in a spreadsheet analysis, a simple exercise that involves two investment managers demands that data for all managers in the portfolio be updated.

3. Multiperiod Analysis

Multiperiod accumulation of attribution results requires careful evaluation and can be highly mathematical for both traditional (micro) and total portfolio attribution analyses. Attribution analysis is often calculated on a daily or monthly basis, while the evaluation period usually spans several days, months, or even years. This requires that daily or monthly attribution results be accumulated over the time period. However, the arithmetic method of attribution favored as the industry standard presents a problem when being compounded.

Below is a simple example using excess return to illustrate that the arithmetic method of adding and subtracting numbers does not work in a multiperiod setting without what is known as mathematical smoothing:

Time Period	Portfolio Return %	Benchmark Return %	Excess Return %
January	12.00	10.00	2.00
February	3.00	-1.00	4.00
January and February	15.36	8.90	≠

- Compounding the portfolio return from January to February is simply (1+12/100) x (1+3/100) – 1 = 15.36%.
- Compounding the benchmark return using the same formula yields 8.90%.
- Therefore, the excess return is 15.36% – 8.90% = 6.46%.
- However, the excess return of 6.46% cannot be computed by adding 2.00% for January to 4.00% for February, as this yields 6%. Neither can it be obtained by compounding these two numbers, resulting in (1+2/100) x (1+4/100) – 1= 6.08%.
- In attribution analysis, the excess return is the sum of attribution results, so it suffers from the same issue identified above.

Since 1999, much has been published on this subject in the *Journal of Performance Measurement*, and several multiperiod mathematical smoothing methodologies have been proposed. Most are currently in use in the industry. Some methods are naïve and, therefore, are simple to implement, such as pro-rating among various attribution results the difference between the excess return calculated from compounded returns and the one computed by sum of monthly attribution results. The popular methodologies are more thoughtful and have greater mathematical rigor.

We performed a careful evaluation of existing arithmetic methodologies in both time-series and cross-sectional analyses to identify the one that best preserves the relative results. That study is its own topic. Here, we cite one example from the study to illustrate that the choice of method matters. Because a large data set is required for this type of analysis, we used equity mutual funds and performed traditional attribution rather than total portfolio attribution, but the same issue applies to both situations.

Below are the cumulative attribution results with various multiperiod arithmetic methodologies, using monthly data. The analysis was conducted using the Morningstar US Market index as the benchmark, and Morningstar sectors as the security classification. You will immediately notice the discrepancy in results among methodologies:

Cumulative Attribution Results Jan 1998–Sep 2009

Methodology	Weighting Effect %	Selection Effect %
Frongello/Wilshire (a.k.a Portfolio Cumulative)	-2.74	8.05
Pro-Rating	-1.95	7.27
Menchero	0.37	4.95
Modified Frongello	1.39	3.93
Cariño	1.45	3.87
Reverse Frongello (a.k.a. Benchmark Cumulative)	4.53	0.79

The weighting effects have different signs, so some of the methods imply that the investment manager has made poor weighting decisions over time while others demonstrate good weighting decisions. The largest gap in this case is 7.26% between the Frongello/Wilshire method and the Reverse Frongello method. Although one might argue that the discrepancy is not large considering that the analysis spans 11 years and three quarters, the highest discrepancy during the time period can be almost 3% per year.

This discrepancy should give pause, and can certainly make you wonder which method to select. And, for those who are currently using one of these methods—especially the Frongello/Wilshire method and the Reverse Frongello—it implies potentially arriving at the wrong conclusion.

4. Multiple-Decision Hierarchy

Attribution analysis is more challenging when the sponsor makes tiered decisions, such as first having a broad, 60% equity, 40% fixed-income allocation, then diversifying the equity portion into various domestic and foreign asset classes, followed by allocations to investment styles as shown in the diagram below:

Hierarchy

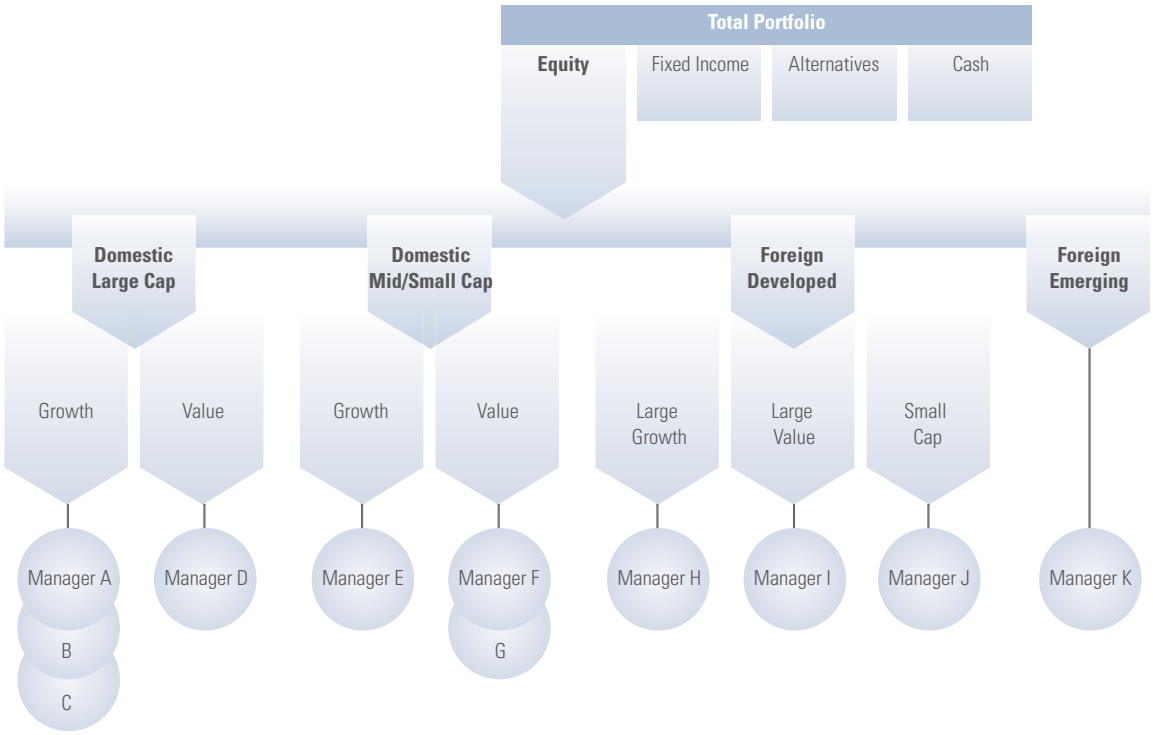

In a multilevel hierarchy, it is important to anchor a decision at the correct hierarchical level; otherwise an overweight in the prior decision could make a subsequent decision appear overweight and vice versa, thus arriving at the wrong conclusion. To see how this can happen, let us examine a simplified example below.

Example of Multiple-Decision Hierarchy

Broad Asset Class / Asset Class	Actual Weigth %	Policy Weight %	Benchmark Return %	Asset Class Weighting Effect %
Alternatives	5.45	10	-0.02	0.00
Real Estate	2.00	5	5.34	0.04
Commodities	3.45	5	-5.38	-0.04

In this example, real estate and commodities are the only two asset classes within the alternatives broad asset class. With one quick glance at the table, one might conclude that commodities is underweight compared with its policy weight, as the actual weight in the asset class is 3.45% while the policy weight calls for 5%. Given that commodities underperformed alternatives with returns of -5.38% versus to -0.02%, it would appear to have been a good decision to underweight an underperforming asset class. In other words, using the conventional Brinson & Fachler methodology, the weighting effect for commodities would be (3.45% - 5.00%) x [(-5.38%) – (-0.02%)] = 0.08% which is a positive number, demonstrating a good decision. But is the wrong conclusion.

Careful examination reveals that commodities is a second-level decision, as the sponsor first decided how much to allocate to the alternatives broad asset class, then determined how to split the allocation between real estate and commodities. It is the alternatives broad asset class that is underweight, with 5.45% actual weight compared with the policy of 10%. This decision to underweight makes real estate and commodities appear to be underweight as well.

For a correct interpretation of results, the actual portfolio weights of a lower-level decision must be anchored to the weights of the prior hierarchical level. In this example, if the 5.45% actual weight in the alternatives broad asset class were to be split 50/50 between real estate and commodities, as the policy weights specify, there would have been 2.725% in each asset class. So, intuitively, we can see that commodities is actually overweight, allocated at 3.45%, while real estate received just 2% weight.

Mathematically, to anchor the policy weight of 5% in commodities on alternatives, it would be 5% x 5.45% / 10% = 2.725%. In other words, a 5% policy weight translates to a 2.725% effective weight. Now we can see that the actual weight of 3.45% in commodities is larger than the anchored policy weight of 2.725% for the benchmark, demonstrating an overweight. Modifying the conventional Brison & Fachler weighting effect formula to account for anchoring, the result becomes [3.45 – (5% x 5.45% / 10%)] x [(-5.38%) – (-0.02%)] = -0.04%, as shown. Taking another look, we observe that overweighting an underperforming asset class was a poor decision, and it makes sense that the asset class weighting effect of -0.04% is a negative number.

This example demonstrates that the methodology must be hierarchical in structure, or you can easily arrive at the wrong conclusion. The investment policy statement drives the complexity of the asset allocation decision tree in terms of layers of decisions and number of classifications. A meaningful attribution analysis should closely match the decision-making of the sponsor, and as a result it is important to make sure your attribution system can be fully customized. Check also that there is sufficiently robust methodology to handle the additional mathematical complexity of a hierarchical structure, especially in a multiperiod setting.

5. Premium/Discount in Closed-End and Exchange-Traded Funds

Closed-end and exchange-traded funds have two types of prices, one based on the net asset value (NAV) and the other based on market price. The difference between these two prices is referred to as "premium" when the market price is above the NAV and "discount" when it is below. The premium or discount widens or narrows over time. The return differential caused by change in premium/discount is not insignificant. For example, based on the Morningstar Global Closed-End database, the average return differential due to change in premium/discount (expressed as absolute numbers) is 3.3% for the month of September 2009, and double-digit numbers are common.

In the Morningstar Global Exchange-Traded Funds database, this number is 1.3%. Although there are few ETFs with a double-digit premium or discount, over 100 funds exhibit more than a 5% difference between returns based on NAV and those based on market price. Many of these have NAV-based returns in positive territory while market-price-based returns are negative and vice versa, and choice of which return to use can change the conclusion about a fund's performance.

When attribution is performed on returns calculated based on NAVs, this correctly reflects the contribution from fund managers, but it does not represent the economic value of the underlying investment. When attribution is performed on returns computed based on market prices, this correctly reflects the economic value of the underlying investment, but the premium/discount that is outside the fund manager's control is incorrectly attributed to a contribution by the fund manager. What is the solution to this dilemma? Isolate the return attributable to change in premium/discount from the manager selection effect. The attribution effect associated with change in premium/discount is informative. It is positive when the premium has widened or discount has narrowed, and it is negative when the premium has shrunk or discount has grown. This type of analysis requires a system that tracks both types of returns. If your plan includes these investments, make sure any attribution system you select has this capability, or plan for extra data collection and methodology in your spreadsheets.

6. Distinguishing Between Skill-Based and Non-Skill-Based Numbers

This is particularly important because the primary purpose of attribution is to analyze the impact of investment decision-making. Numbers unrelated to skill should be isolated; for example, costs such as investment management, investment consulting, advisory services, and other fees affect the final take-home return but do not

represent investment decisions. Traditional attribution does not address these costs, but a good total portfolio attribution system should provide this as an option. This requires that the system track returns both gross-of-fees and net-of-fees, or support import of the missing component.

Another non-skill-based statistic that must be isolated is benchmark misfit, also known as style bias. The concept of benchmark misfit was introduced by Bailey in 2001, and we have adapted and broadened its application in total portfolio attribution. In the example below, we have assumed that the U.S. large-cap asset class is implemented with both a value and a growth manager. Each of these managers should be compared to a style-specific benchmark, for example, the Russell 1000 Value for the large-cap value manager. Let us further assume that this manager outperformed the Russell 1000 Value, but underperformed relative to the S&P 500.

Manager Selection Effect and Manager Benchmark Misfit

Name	Benchmark	Actual Weight %	Portfolio Return %	Benchmark Return %	Manager Selection Effect %	Manager Bechmark Misfit %
U.S. Large-cap	S&P 500	29.20	-6.07	-8.43	0.70	-0.01
Large-cap value manager	Rus1000Val	16.00	-10.16	-11.50	0.21	-0.49
Large-cap growth manager	Rus1000Gr	13.20	-1.10	-4.81	0.49	0.48

A positive number of 0.21% for the manager selection effect indicates that the large-cap value manager outperformed his designated index, the Russell 1000 Value. However, this manager underperforms the benchmark assigned to the overall U.S. large-cap asset class, and thus is a negative contributor at the asset-class level. We decompose the underperformance compared to the S&P 500 into two components: a positive manager selection effect for outperforming the style-specific benchmark and a negative benchmark misfit measure to represent that the value style of investing is currently out of favor (indicated by the Russell 1000 Value underperformance relative to the S&P 500). The large-cap value manager's ability to outperform his designated benchmark is a skill metric, while value investing being out of favor is the result of the market—out of the large-cap value manager's control.

In the example, the manager selection effect is the weighted difference between the manager's return and his or her designated benchmark, or 16.00% x [(-10.16%) – (-11.50%)] = 0.21%. The benchmark misfit is the weighted difference between the return of the manager's designated benchmark and that of the asset class benchmark, or 16.00% x [(-11.50%) – (-8.43%)] = -0.49%. In the example, instead of doing a diagonal analysis comparing the manager's performance of -10.16% to the S&P 500's -8.43%, we do a horizontal comparison of the manager versus Russell 1000 Value to obtain the manager selection effect, followed by a vertical comparison of Russell 1000 Value versus the S&P 500 for the benchmark misfit.

Benchmark misfit exists when a manager's benchmark is not the same as that of the asset class. Therefore, a benchmark misfit would not appear if the investment policy separates growth and value into two asset classes, each with separate target policy weights and benchmarks. In addition, there is no benchmark misfit if both value and growth managers are compared to the asset class benchmark instead of style-specific benchmarks. It may be simpler to ignore benchmark misfit, as some attribution models and literature do, but one must consider whether doing so is consistent with the asset class definition and manager performance evaluation guidelines stated in the investment policy statement.

7. Attribution as Communications Tool

Last but not least, more than being an in-house analytical tool, an attribution report should be an effective communications tool to a sponsor's clients, investment committee, or board. As such, attribution results should be conveyed in combinations of charts and tables designed and labeled to meet the needs of audiences of varying levels of sophistication. A friendly report can command more attention and immediately improve the odds of productive dialogue. A good reporting tool should be easy to customize and offer options for explanatory copy and inclusion of logos, etc. Not all software providers can accomplish this, and limitations are particularly pronounced for proprietary spreadsheets.

Conclusion

There are many factors to consider when making a build-or-buy decision, or contemplating transition from proprietary spreadsheets to a commercial software application. Some of these are related to operational efficiency, such as data collection and spreadsheet maintenance. Clear reports to convey results to clients, investment committees, and boards are extremely important. Most challenging of all are the nuances in methodology, as an incorrect application of methodology can lead to erroneous conclusions about the quality of investment decisions.

As you evaluate potential solutions, be sure to include the following in your assessment:

- The extent and flexibility of data management capabilities
- How portfolio changes and multi-period accumulation of attribution results are accommodated
- The ability to support attribution on a portfolio with multiple asset classes, investment mandates, and managers
- The separation of skill-based measure from non-skill-based measures such as fees
- Clear and understandable reports, graphics, and other communication materials

The Morningstar Direct℠ Advantage

Total Portfolio Attribution from Morningstar Direct has the mathematical rigor sufficient to handle complex portfolios. It enables users to identify the sponsor's impact on total returns, decomposing the performance impact attributable to:

- the strategic asset allocation policy itself
- the outcome of deliberate deviations from policy weights done to take advantage of market conditions
- manager-picking skill in aggregate, or by asset class or investment style

All Morningstar databases are included: stocks, hedge funds, separate accounts, insurance products, mutual funds, closed-end funds, and exchange-traded funds. Morningstar Direct has an intuitive interface for importing proprietary investments, fees, and other nonpublic data. What's more, Morningstar's attribution report templates are dynamic, effective tools for communicating the sources of portfolio performance to clients, boards, investment committees, and internal stakeholders.

An All-in-One Solution

In addition to Total Portfolio Attribution, Morningstar Direct also features traditional attribution to assess the performance of investment managers for due diligence and ongoing monitoring. It includes powerful searching and filtering capabilities, with tools to automate portions of your work. You will find sophisticated peer analysis options, as well as performance reporting. In short, Morningstar Direct has everything you need to simplify your processes, enhance productivity, and save money. Join the thousands of thoughtful people at institutions worldwide who have consolidated multiple tools with Morningstar Direct.

Learn More

Visit our website at http://global.morningstar.com/WhitePaperTPA

To arrange a no-risk demonstration of Total Portfolio Attribution in the convenience of your office:
Call: +1 866 910 0840
E-mail: information@morningstar.com

About the Author

Cindy Sin-Yi Tsai is senior research analyst for Morningstar, Inc. Before joining Morningstar in 2008, she was director of research, performance, and product development for Altair Advisers. Prior to forming Altair, Tsai specialized in portfolio asset allocation modeling and investor risk analysis as vice president of product development for West Hills Institutional. Tsai holds a bachelor's degree in business economics, with honors, from the College of Wooster in Ohio, and a master's degree in finance from the University of Illinois at Urbana-Champaign. She holds the Chartered Financial Analyst (CFA) and the Chartered Alternative Investment Analyst (CAIA) designations.

References

Bailey, Jeffery V., "Benchmarks and Attribution Analysis for the Total Fund," Benchmarks and Attribution Analysis, *AIMR Conference Proceedings*, 2001, no. 3, pp. 28-37.

Bonafede, Julia, Steven Foresti, and Peter Matheos, Ph.D., "A Multi-Period Linking Algorithm That Has Stood the Test of Time," *The Journal of Performance Measurement*, Fall 2002, pp. 15-26.

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-US Equity Portfolio Performance," *Journal of Portfolio Management*, Spring 1985, pp. 73-76.

Cariño, David R., "Combining Attribution Effects Over Time," *The Journal of Performance Measurement*, Summer 1999, pp. 5-14.

Campisi, Stephen, "Balanced Portfolio Attribution," *The Journal of Performance Measurement*, Winter 2008/2009, pp. 24-35.

Frongello, Andrew Scott Bay, "Linking Single Period Attribution Results," *The Journal of Performance Measurement*, Spring 2002, pp. 10-22.

Frongello, Andrew Scott Bay, "Readers' Reflections," *The Journal of Performance Measurement*, Winter 2002/2003, pp. 7-11.

Levecq, Christian, "An Exposure-based Attribution Model for Balanced Portfolios," *The Journal of Performance Measurement,* Summer 2004, pp. 14-22.

Menchero, Jose G., "An Optimized Approach to Linking Attribution Effects Over Time," *The Journal of Performance Measurement,* Fall 2000, pp. 36-42.