Multi-period decision attribution

Andre Mirabelli, Feb 2018

Brinson Attribution was originally conceived as a single-period evaluation of the impact of sector effects vs. issue effects. Its development has moved along two dimensions. The first is its expansion to include multi-period evaluations. The second is toward a clearer understanding of its purpose as an evaluation of controllable investment decisions.

There have been a large number of attempts to address the first aspect alone. These include models by David Carino, Jose Menchero, Andrew Frangello, Stephen Campisi and a significant number of various other models I presented in the 'Structure and Visualization of Performance Attribution' in the 2000/2001 issue of the Journal of Performance Measurement. Numerous additional articles on the topic have been published in that journal since that time. Generally, these methods calculate attributes for each day in isolation. Then, noting that these terms do not roll up to the desired multi-period active return, since they ignore temporal compounding, they in one way or another smooth the error over the obtained terms. They thereby warp the original meaning of these terms without supplying an alternate meaning for the terms thus obtained other than that they are formally adjusted single-day terms that now roll up to the desired results.

However, these models have the following problem. They are each caught on one of the horns of a dilemma that can be seen from the following two–period attribution example.

In the first period the fund makes 10% return and the benchmark makes 0% return.
So at the end of the first period the active return is 10%.
By the end of the first period the fund is made into an exact replica of the benchmark.
In the second period the benchmark makes 10% return and, since the fund is an exact replica of the benchmark during the second period, the fund also makes 10% return in the second period.
Thus, for the whole of the two periods, the fund made 21% and the benchmark made 10% (as seen by compounding the single period results for the fund and benchmark). Therefore, at the end of the second period, the active return is now 11%.
Where did this additional 1% active return come from? To what decision should it be attributed?

The versions of these author's methods fall into one of two categories:

1. They assign the 1% to the first period, even though at the end of the first period they said that the first period achieved only 10% active return. Now they have to go back and revise the evaluation of what happened in the first period and what was reported for it in the past, to say that 11% active return has to be attributed to the first period.
2. They assign the additional 1% active return to the second period even though the fund was an exact replica of the benchmark throughout the second period.

The first source of this problem is the lack of a clear formulation of the problem that multi-period decision attribution intends to address. Such a proper formulation is the evaluation of the impact of each controllable investment decision made each day upon the active return up to that point in the decision process. Here 'the impact of a decision on the active return" is defined as the change in the active return as one goes from the end of the implementation of the previous decision on a day to the end on that day of the implementation of the decision being evaluated. This is in opposition to the above flawed methods which instead begin with terms that evaluate the impact upon each day in isolation. Since these do not sum

to the active return for the period, these flawed methods then take the error (1% in the above case) and smooth it across the other terms. Some put it all into the active return of the first day and some all into that of the second, and some into combination thereof.

In a correct analysis of this simple case, the 1% is assigned to the Invest Effect. But in more robust examples, some of the difference between the single day effects and the effects on the evolving active returns is distributed among the other Attributes. However, this is now not done to simply distribute the error of an incorrectly formulated method. Instead it is done in order to correctly model the answer to the properly formulated attribution question about the impact of each decision. In this well-formulated economically meaningful approach, no “Interaction” terms or choices between Brinson/Fachler and Brinson/Hood/Beebower arise, since they are not part of any decision evaluation process.