A Brinson Model Alternative: an Equity Attribution Model with Orthogonal Risk Contributions

This paper presents an algorithm for equity attribution that circumvents some of the less intuitive features of the standard Brinson model. The proposed alternative framework provides an unambiguous division of returns by source of risk, in a manner consistent with first-principles interpretation.

Andrew Colin, Ph.D.
is Director of Fixed Income Research for the StatPro Group plc. He holds a Ph.D. in Mathematics from the University of St. Andrews and has worked or consulted for a wide range of banks and other companies. He also is adjunct professor at the Faculty of Business, Queensland University of Technology, where his team researches risk in financial and engineering asset management.

INTRODUCTION

The Brinson (or sector) attribution model and its successors (Brinson, Hood, Beebower, 1986; Brinson, Fachler, 1985) are widely used in portfolio management as tools to assess the skills of the manager in running an equity investment portfolio. These models decompose a portfolio's excess return above the benchmark into returns arising from stock selection and asset allocation decisions.

In naïve terms, stock selection is deciding *which* stocks to hold, while asset allocation is deciding *how much* of each stock to hold. Stock selection is a binary, yes/no type of decision, while asset allocation forms a continuous, constrained set of decisions.

Let us put this another way. Consider an investment "tipping" newsletter, where a group of stocks are singled out by an expert as being a good buy. These recommendations form pure stock selection advice. How much of each stock to buy is never suggested, since this decision cannot be made in isolation but must take into account the other holdings of the investor.

If one based an investment strategy on the use of such recommendations, the ideal attribution analysis would be one that measures the return generated by the use of the newsletter (stock selection) and the return generated by decisions on amounts bought (asset allocation). Sectoral attribution cannot provide such an analysis, and this is one of the motivations for the algorithm presented in the paper.

In addition, the sector attribution model does not clearly discriminate between the returns generated by these two types of decisions. Specifically, there are two ways in which the model falls short:

1. The returns attributed to stock selection in the Brinson model also include the effects of asset allocation decisions at the subsector level. This extra asset allocation return is not split out from the stock selection return, even though it arises from a quite different source of risk.

This is illustrated in the following example. Suppose that a manager decides that a \$100 million portfolio should contain 15% by market weight of IT stocks. The benchmark weight is 10%, so the portfolio is overweight IT stocks by 5 percent. This is an asset allocation decision, and the return generated by this decision is classed under asset allocation return. Asset allocation does not form a decision on which specific IT stocks to hold; it just indicates how much exposure there should be to a particular class of securities.

The manager now decides to invest the 15% of his funds allocated to IT into Microsoft and Borland stock. He purchases \$10 million of Microsoft and \$5 million of Borland stock. The returns from this decision are attributed to stock selection, even though the return has been generated by a mixture of stock selection and asset allocation decisions, *i.e.,* to buy Microsoft and Borland, and *not* to buy IBM or Dell and to purchase stock in the above amounts rather than \$2 million of Microsoft and \$13 million of Borland. To put it another way, the decision of how much exposure there should be to individual stocks is a continuation of the sector-level asset allocation decision and should be included in the asset allocation returns. This is not the case in the conventional sector attribution model.

2. Despite the qualitatively different nature of the stock selection and the asset allocation decisions, the Brinson framework does not clearly discriminate between the two due to the appearance of an *interaction* term. To quote Bacon (2004), “A flaw of both Brinson models is the inclusion of the interaction or other term. Interaction is not part of the investment process; you are unlikely to identify in any asset management firm individuals responsible for adding value through interaction.” To put it another way, the returns are not orthogonal, since the returns from both decision types can overlap.

The Brinson Attribution Model

If one is managing a portfolio from the stock level upwards, one cannot decide upon the *quantity* of each stock to buy without first deciding *which* stocks to buy. In this case, the stock selection decision precedes the asset allocation decision. If the portfolio is managed in this manner, it uses a bottom-up investment process.

A commonly used alternative approach is top-down investment. The investor first *partitions* the portfolio and benchmark into sectors or buckets based on industry sector, large or small capitalization, or some similar classification. The use of such a partition has, in fact, nothing to do with the portfolio’s risk or overall return; it is an externally imposed means of allowing investment decisions to be taken in a hierarchical fashion.

By imposing this partition on the universe of investible stocks, it is possible to reverse the order in which the previous decisions were taken. A broad-brush asset allocation decision is made by assigning desired holding weights to each sector, and using this as a constraint on further investment. Typically, such weights are measured relative to the corresponding weights in the benchmark, so we talk of being over-weight or under-weight with respect to a given sector.

In the Brinson framework, the excess generated by this over- or under-weighting at the sector level can be measured and is referred to as asset allocation return; it is measured by

$$r_{AA} = \sum_{i \in S} \left(w_i - W_i \right) \times b_i \quad , \tag{1}$$

where w_i and W_i are the weight of the portfolio and benchmark for sector i, respectively, and b_i is the return of the benchmark for sector i.

Once the sector weights have been chosen, the composition of each sector is decided. This process involves both stock selection choices and further asset allocation decisions within the sector. The returns by both these decisions at this level are aggregated together as stock selection returns. In the Brinson framework, these returns are measured by

$$r_{SS} = \sum_{i \in S} W_i \times \left(r_i - b_i \right) \quad , \tag{2}$$

where r_i is the return of the portfolio for sector i, and the other terms are as defined for Equation 1.

The sum of these two terms does not sum to the active return, defined as the return of the portfolio above that of the benchmark. To fix this, one can either introduce an additional interaction term

$$r_I = \sum_{i \in S} \left(w_i - W_i \right) \times \left(r_i - b_i \right) \quad , \tag{3}$$

where the terms are as defined above, or use the portfolio sector weight w_i in Equation 2 so that

$$r_{SS} = \sum_{i \in S} w_i \times \left(r_i - b_i \right) \quad . \tag{4}$$

The Brinson framework allows a hierarchical decomposition of sectors into subsectors, so that transport stocks can be decomposed into railways, aircraft, etc. In principle one could continue to break the portfolio down so that eventually each sector contains exactly one security. However, this is not particularly useful, since then the return of the security and the benchmark sector will be the same, and the stock selection return will be zero.

An Alternative to the Brinson Attribution Model

Suppose that the set of benchmark stocks B represents the investible universe (we cover the case of stocks that lie outside the benchmark below). The benchmark has non-zero holdings in every stock within B, with return

r_i and weight a_i. The weights $\{a_i\}$ satisfy $\sum_i a_i = 1$, and the return of the benchmark R_B is given by $R_B = \sum_i a_i r_i$.

We now select the stocks for our portfolio *P*, where $P \subset B$. The first step is to build an intermediate portfolio P' composed of only these stocks, *but with the same relative holdings that they have in the benchmark*. This is similar to the construction of a semi-notional portfolio in the Brinson framework, but at the security level instead of the sector level. Here, we are simply considering the effect of modifying the benchmark holdings by including some stocks and excluding others; we are expressly not considering asset allocation effects. The weights in this portfolio are given by

$$w_i' = \frac{W_i}{\sum_{i \in P'} W_i} \quad , \tag{5}$$

for $i \in P'$, and

$$w_i' = 0 \quad , \tag{6}$$

for $i \notin P'$.

The return due to the stock selection decision is the difference in return between this new portfolio P' and the benchmark, given by

$$R_{SS} = \sum_{i \in B} \left(w_i' - W_i \right) \times R_i \quad , \tag{7}$$

The return due to the asset allocation decision is the difference in return between P' and the portfolio *P*, given by

$$R_{AA} = \sum_{i \in B} \left(w_i - w_i' \right) \times R_i \quad . \tag{8}$$

In both (7) and (8), the sum is over all securities in the benchmark.

For instance, suppose the portfolio and benchmark are composed of five stocks with the following weights and returns. Note that returns are the same for portfolio and benchmark, since we are working at the security level. If a security is traded at a non-end of day level, then the return of that security in the portfolio may differ to the return of the same security in the benchmark. In this case, we may attribute the difference in return to trading effects, and proceed as before using the benchmark return.

We decide to build a portfolio from stocks 1 and 3. Based on their relative holdings within the benchmark, their weights within P' will be 0.3/(0.3+0.1) = 0.75 and 0.1/(0.3+0.1) = 0.25, respectively.

The difference in return between the benchmark *B* and this intermediate portfolio P' is the return due to the stock selection decision. The benchmark return is the sum of the performance contributions, which is (0.3 * 2%) + (0.3 * -2%) + (0.1 * 1%) + (0.1 * -2%) + (0.2 * 0%) = -0.1%. The return of the intermediate portfolio is (0.75 * 2%) + (0.25 * 1%) = 1.75%. Therefore, the return due to stock selection is 1.75% - (-0.1%) = 1.85 percent.

The next step is to include the effect of the asset allocation decision. Suppose that the weights of stocks 1 and 3 are 0.8 and 0.2, respectively. Then the portfolio's return is (0.8 * 2%) + (0.2 * 1%) = 1.8%, and the return due to asset allocation is 0.05%.

The sum of the asset allocation and stock selection term is 1.90%, which is the difference between the portfolio and benchmark return, as required. There is no interaction term.

These results are fully consistent with intuition. We made an excellent stock selection decision by buying the only two stocks from the benchmark with the positive returns, so we would certainly expect to see a positive stock selection return. After taking stock selection into account, we also made a fairly conservative asset

Table 1: Sample Portfolio Structure

Stock	1	2	3	4	5
Portfolio weight	0.8	0.0	0.2	0.0	0.0
Benchmark weight	0.3	0.3	0.1	0.1	0.2
Return	2%	-2%	1%	-2%	0%

allocation decision by over-weighting the highest performing stock (1) by 0.05 and under-weighted the next highest (3) by 0.05, leading to the slightly positive asset allocation return of (0.05% * 2%) + (-0.05 * 1%) = 0.05 percent. Again, this added value, since we over-weighted a higher performing stock.

At no point have we needed to define or use a partition to run this attribution analysis. The analysis is made entirely at the stock level. From these stock-level returns, we can slice and dice the relative attribution returns to provide any level of insight required (and subsequently to aggregate attribution returns at sector levels, if required).

Investing in Non-benchmark Stocks

An off-benchmark stock can best be handled by viewing it as having a zero holding in the benchmark. In this case, Equation 7 implies that the stock selection return of all such stocks is always zero and that the excess return generated by holdings in such stocks is due to asset allocation. This is consistent with the approach proposed by Laker (2006).

DISCUSSION

How much does this scheme match user expectations? We give two examples to illustrate common situations.

First, suppose that the user makes no stock selection decisions and includes all the benchmark stocks in the portfolio. In this case the intermediate portfolio P' will be identical to the benchmark, so the difference in returns between the two (which we have identified as the stock selection return) will be zero. This is precisely as expected, since any value added will be generated purely by asset allocation decisions.

Compare this to the Brinson scheme. In this case, the stock selection return will only be zero if the sum of returns in Equation 2 is zero. The returns of each sector in the portfolio will be affected by the asset allocation decision, so this is unlikely to occur.

Second, suppose that the user only takes stock selection decisions without asset allocation. Strictly speaking, this is not possible, since the inclusion or omission of any stock will affect its asset allocation. Both the asset allocation and stock selection decisions are, in fact, asset allocation decisions. The stock selection is just a question of whether a stock's holding should be set to zero or otherwise.

However, in the Brinson framework we can match the exposures within each sector as described above; in which case the first term in Equation 1 will be zero and no return due to asset allocation will be recorded.

Is this approach consistent with conventional equity attribution schemes? The answer is: it can be made so. If we impose a partition schema onto our portfolio and benchmark, we can certainly use a top-down investment process by assigning particular weights to each sector within the managed portfolio. Each sector then has to be managed as a sub-portfolio, applying stock selection and weight selection; with the constraint that the total amount of stock held be consistent with this externally imposed constraint.

In the new scheme, this will not be the case, and there will be return due to asset allocation even if the sector weights are matched. The reason is that the management of each sector also involves asset allocation decisions, even though the effect of these is rolled into the stock selection return in the Brinson framework. We contend that this new scheme conveys a clearer picture of the effects of decisions made at all levels.

ACKNOWLEDGEMENTS

The author thanks Carl Bacon, Frédéric Bardoux, Mathieu Cubilié and Damien Laker for their comments on an earlier draft of this paper.

REFERENCES

Bacon, Carl, Practical Portfolio Performance Measurement and Attribution, Wiley, 2004.

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

Brinson, Gary, Randolph Hood, and Gilbert Beebower, "Determinants of Portfolio Performance," *Financial*

Analyst Journal, July/August 1986, pp. 39-44.

Laker, Damien, "Performance Attribution with Zero-weighted Sectors," *The Journal of Performance Measurement*, Summer 2006, Vol.10 No. 4, pp. 33-46.