More on the Mathematics of Returns-Based Style Analysis

Introduction

In my article "The Mathematics of Returns-Based Style Analysis" in our last newsletter, I explained how exactly we calculate the style attribution coefficients that StyleADVISOR displays in the style map view. Here's a brief summary of what I said:

> Let M be the manager return series and A_1, A_2, A_3, A_4 the return series of the chosen asset classes, i.e., the style indices. (The number four is chosen for simplicity; one could work with any number of indices.) Following William F. Sharpe's method of returns-based style analysis, we determine the style attribution coefficients c_1, c_2, c_3, c_4 in such a way that the variance of the series
>
> $$M - (c_1A_1 + c_2A_2 + c_3A_3 + c_4A_4)$$
>
> becomes minimal. If we refer to the expression $c_1A_1 + c_2A_2 + c_3A_3 + c_4A_4$ (that is, the weighted composite of the asset classes) as the "style benchmark," then Sharpe's method can be rephrased as follows:
>
> *The style attribution coefficients are determined in such a way that the variance of the excess return of the manager over the style benchmark becomes minimal.*

As I mentioned in my previous article, the mathematical details of actually calculating the style attribution coefficients are quite complex, but they are purely technical in nature and contribute little or nothing to the understanding of returns-based style analysis.

Now that we know *how* returns-based style analysis is done, the question arises as to *why* it is done in this way. In other words, in what sense do the style attribution coefficients that are calculated as described above represent the manager's effective asset mix? And why is it that the resulting style benchmark, that is, the weighted composite of the asset classes, can be considered a fair and appropriate benchmark for the manager? Those are the questions that I want to address in this article and my next one. As it turns out, Sharpe's method can be interpreted mathematically in two different contexts. Each of these interpretations provides support for the conjecture that Sharpe's method is the correct and appropriate one for determining a manager's effective asset mix. The first one views returns-based style analysis as a curve-fitting problem. The second one, which I'll talk about in our next newsletter, is more mathematically involved. It views returns-based style analysis as a calculation in a Euclidean space.

Let me also mention that a more detailed and more formal discussion of the mathematics behind William F. Sharpe's returns-based style analysis appears in a chapter that I contributed to the recent 3rd edition of T. Daniel Coggin's and Frank J. Fabozzi's Handbook of Equity Style Management, published by John Wiley & Sons.

Returns-Based Style Analysis as a Curve-Fitting Problem

The manger and index return series that make up the input of a returns-based style analysis can be easily graphed in a suitable coordinate system. This is done by using time as the x-axis and then placing the returns for each point in time on the y-axis, as shown in Figure 1. In reality, since a manager's returns are given daily, monthly, or quarterly, rather than continuously, the graph of the manager's return series would look more like a jagged sawtooth line than like a smooth curve. However, that is a subtlety that is irrelevant for the discussion at hand.

Now suppose we wanted to graph the manager's style benchmark—as calculated by Sharpe's method—in the same coordinate system as the manager itself. Recall that the style attribution coefficients are calculated in such a way that the variance of the difference between the manager and the style benchmark becomes minimal. In general, what does it mean for any series to have zero or near-zero variance? It means that the series is constant, or near constant. And what does it mean for the difference of two series to be constant, or near-constant? It means that the two series run parallel, or near-parallel. Therefore, we have the following sequence of equivalent statements:

The variance of the excess return of the manger over the style benchmark is minimal.

if and only if

The difference between the manager series and the style benchmark series is as close as possible to a constant series.

if and only if

The manager series and the style benchmark series are as close as possible to running parallel.

if and only if

The manager series and the style benchmark series have the same or very similar shape, without necessarily being close to each other.

Figure 2 demonstrates this situation: the manager and the style benchmark series are very similar in shape, but they differ by a considerable, near-constant vertical distance.

In summary, we can now give an equivalent, somewhat "visual" characterization of Sharpe's returns-based style analysis: *Sharpe's method determines the style attribution coefficients in such a way that the style benchmark series (that is, the weighted composite of the asset classes) most closely mimics the manager series in shape, without necessarily being close to the manager series.*

The underlying assumption here is that the *shape* of a manager's return series is determined by his style. The manager's skill, on the other hand, results in a *constant or near constant addition* to the returns, without significantly affecting the shape of the return series. Looking at things in this way also lays to rest one of the most frequently voiced objections to returns-based style analysis. Suppose you are looking at a period in time where growth outperformed value. Assume further that you are analyzing a pure value manager who, due to superior skills, did better than overall value indices. Would that manager not be labeled incorrectly as an average growth manager rather than an above-average value manager? The answer is no. Although the manager's return series may lie closer to a growth index than to a value index, its shape is still determined by the

manager's value style. That is what returns-based style analysis detects, thereby labeling the manager correctly as a value manager. The fact that the manager was skillful enough to do better than value in general will show up when looking at the excess return over the style benchmark. Thus, the manager's style will be recognized correctly and she will be given due credit for her skill.

Comparison with Related Methods

Using somewhat more abstract mathematical terminology, the main result of the previous section can be rephrased as follows: *Returns-based style analysis determines the style benchmark by solving the following curve-fitting problem: find the linear combination of the asset classes that is the best fit to the manager's return series, where the criterion of best fit is the shape of the curve.* There are of course other criteria of best fit that one might consider for this curve-fitting problem. For example, it has been suggested in the past to minimize the sum of the squares of the difference between the manager series and the style benchmark series. Geometrically, this amounts to using closeness as the criterion of best fit, that is, one would try to place the style benchmark as closely as possible to the manager series. Figure 3 shows a style benchmark that is obtained in this way. If style analysis were done in this way, then indeed the objection discussed in the previous section would be valid: in times where growth outperforms value, an above-average value manager would be erroneously labeled as an average growth manager.

A number of other criteria of best fit exist, e.g., one might try to maximize the correlation between manager and style benchmark. None of these methods, though, results in the clear distinction between a manager's style and his skill that one obtains using William F. Sharpe's method of minimizing the variance of the excess return.

Summary

William F. Sharpe's method of returns-based style analysis rests on the assumption that a manager's style determines the shape of her return series, whereas her skill results in a constant or near-constant addition to the returns. Using shape as the criterion of best fit, one can thus infer the manager's style and skill by means of a curve fitting process.
In our next newsletter, I will explore an entirely different mathematical context in which returns-based style analysis can be interpreted, namely, the theory of Euclidean spaces. This interpretation confirms in yet another way the plausibility and usefulness of Sharpe's method.

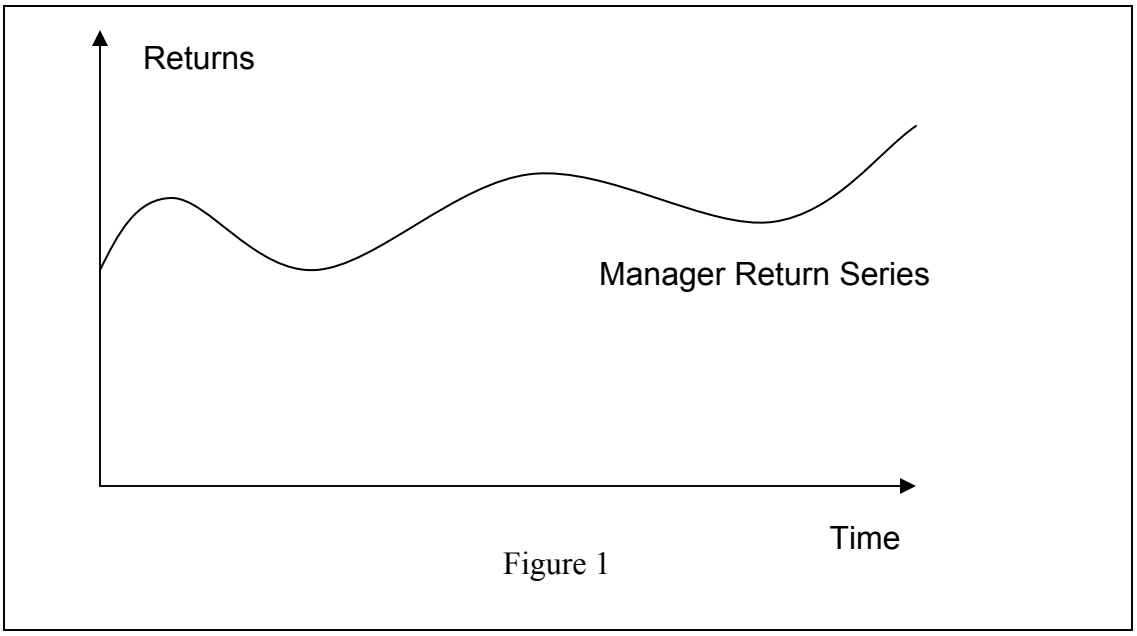

Figure 1

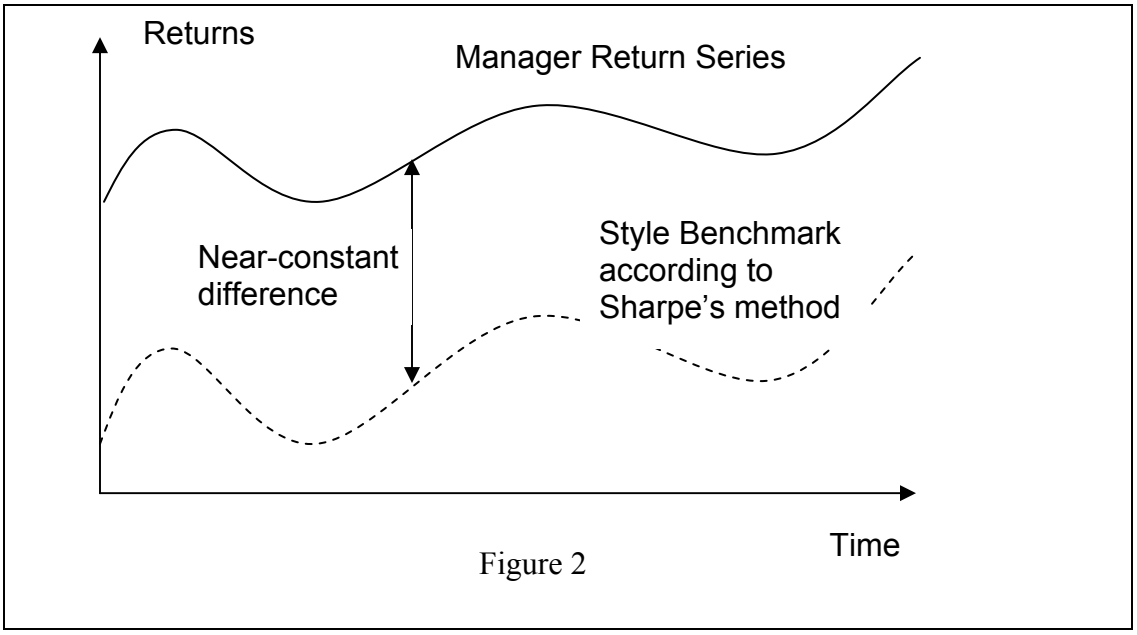

Figure 2

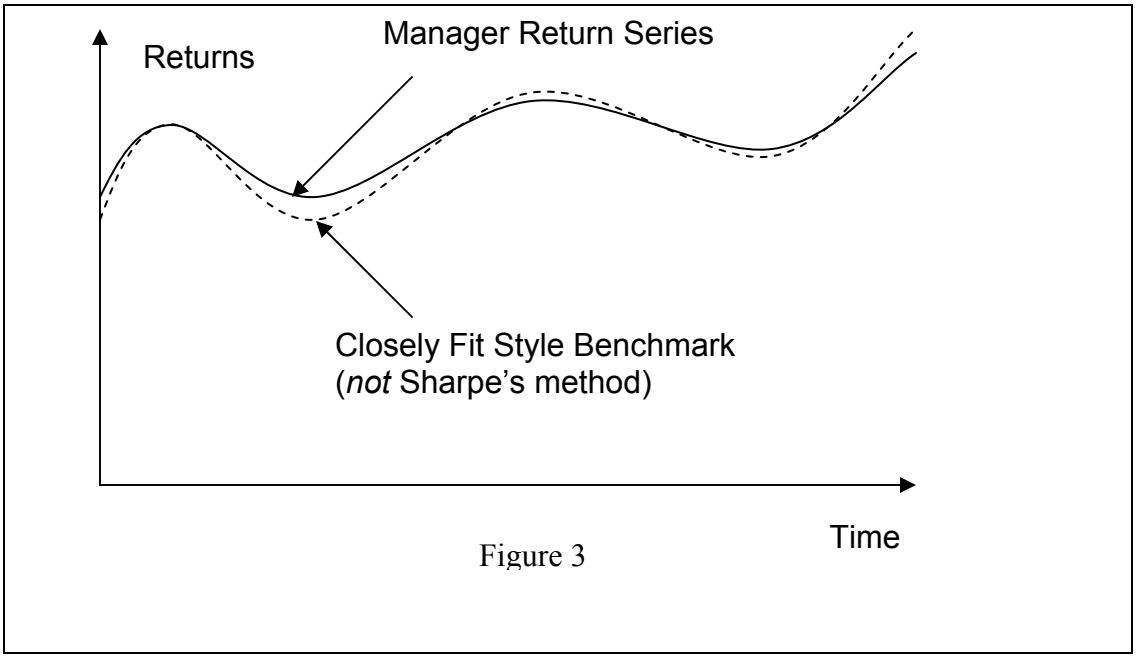

Figure 3