



Risk-Based Performance Attribution

Research Paper 004

September 18, 2015

Risk-Based Performance Attribution

Traditional performance attribution may work well for long-only strategies, but it can be inaccurate and even misleading when applied to hedge fund strategies. Risk-based performance attribution, while difficult to perform, is able to provide a more accurate picture of the drivers of hedge fund performance.

Simple Attribution

The easiest way to evaluate the performance of a fund is simply to compare it to a benchmark. For example, if, in a given year, a fund returns 18% and the S&P 500 returns 10%, then we could say that the fund outperformed the market by 8%.

If the fund in question is not highly correlated to the S&P 500, though, then this simple approach can be very misleading, especially in the short-run. For example, a fund with a low correlation to the market may underperform the market in years when the market is up significantly, but still outperform the market in the long-run. Most investors value low correlation, so how do we stop ourselves from unduly penalizing these low correlation funds?

This problem can be partially solved by choosing better benchmarks. In recent years, a number of hedge fund indexes have been created. While these hedge fund benchmarks will provide a better match in some instances, appropriate benchmarks will be impossible to find for many managers. More importantly, no matter how good the benchmark, this simple approach cannot tell us what factors are driving a manager's performance. The simple approach can only tell us that a manager is outperforming or underperforming a benchmark, it cannot tell us why.

A Better Approach

Another popular approach to performance analysis is to regress the returns of a fund manager against a benchmark. We end up with a formula like this:

$$R_t = \alpha + \beta R_{M,t} + \varepsilon_t$$

Here, R_t is the return of the fund at time t , $R_{M,t}$ is the return of the market, α and β are constants, and ε_t is an error term, which is zero on average¹. If a fund is highly correlated to the market and has similar volatility, then β will be close to 1, and this method will produce results that are very similar to the simple approach outlined in the previous section. If, on the other hand, β is positive but less than 1, then the fund will not fully participate in up markets, but will also not fully participate in down markets.

As an example, let's consider two managers. The first manager is running an index tracking fund, has a beta of 1.0 and generates no alpha. The second manager is a typical long/short manager, has a beta of 0.2 and generates alpha of 4% per year. Now assume that the market is up 20% one year and then down 10% the

¹ To be consistent with the capital asset pricing model (CAPM), we would first subtract the risk-free rate from both the fund and market returns. In this context α is often referred to as Jensen's alpha. To be concise and keep the formulas simple, we'll continue to use this formulation and to say "returns", with the understanding that we can easily substitute excess returns for simple returns.

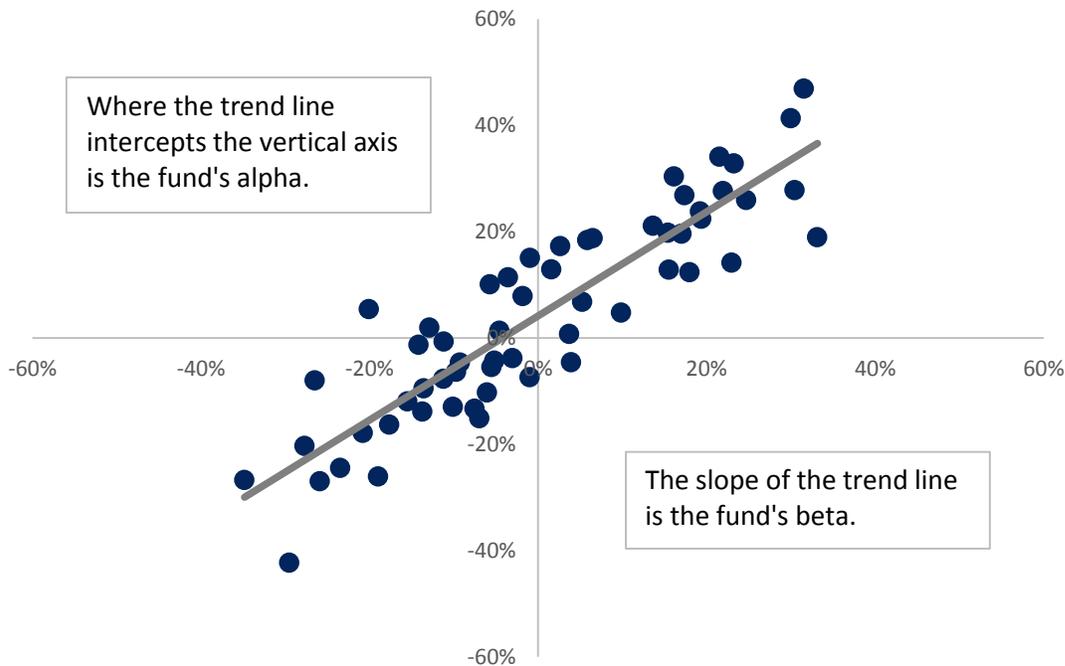
following year. What would we expect the returns for the two manger to be? Because ϵ_t is zero on average, we can ignore the final term in our equation. For the two funds, the expected returns would be:

Year	$\beta=1, \alpha = 0\%$	$\beta=0.2, \alpha = 4\%$
1	20.0%	8.0%
2	-10.0%	2.0%
Total	8.0%	10.2%

The $\beta = 1$ manager has the exact same performance as the market in both years, 20% followed by -10%. Because of compounding, the total return over 2 years is not +10%, but +8%. Now look at the low-beta manager. The low-beta manager underperforms in the first year, when the market is up significantly, just as we would expect. And, just as we would expect, the low beta manager outperforms in the second year when the market is down. Overall, the low beta manager returns 10.2%.

Interestingly, in this example the managers have the same average returns, but the low beta manager still outperforms over two years. In general, because of compounding, if two managers have the same average returns, but one has lower volatility, the lower volatility manger will outperform the higher volatility manager. This example is not meant to be an argument for low beta strategies over high beta strategies, however. What the example *is* meant to show is that if we know a manager's beta to an index we are better off for two reasons. First, we are better off because we have a better understanding of the managers' performance. In the first year, we understand that, even though the low beta manager did worse in absolute terms, he added more alpha, which should be very valuable in the long run. Second, we have a better understanding of the managers' risk. The total return of the managers was similar over the past two years, but if the market is down 20% next year we would expect the high-beta manager to lose significantly more than the low-beta manger (-20% versus 0%).

In the preceding analysis we assumed that the beta of each manager was known. In practice we need to estimate beta. The traditional approach to estimating beta is to regress the manager's historical returns against those of an index. While the math is straightforward (this type of analysis can easily be performed in Excel), in practice we are often faced with a very serious problem, namely a lack of data. The median hedge fund age is roughly three years. If we only have access to monthly returns, this will give us just 36 data points to work with. As we will see in the next section, even when we do have enough data points, if a manager's exposures are changing over time —as they almost always are— this standard regression approach can still be very misleading.



Risk-Based Attribution

Now imagine two different managers: the first is a macro trader, who is either long or short S&P 500 futures; the second is a fundamental long/short manager, but, unlike in our previous example, this manager has no net exposure to the S&P 500 ($\beta = 0$). The macro trader changes his exposure frequently. One day the macro trader is long, the next day short, the next long again. Importantly, the macro trader is just as likely to be correct—that is long when the market is up, and short when the market is going down—when the market is going up as when the market is going down. The long/short manager is always neutral.

At the end of the year, the market is flat, but both managers have returned +10%. The macro manager managed to generate positive returns in a flat market by being correct more often than not about the direction of the market. The long/short manager managed to be up in a flat market because the longs in his portfolio outperformed his shorts.

If all we did was look at the returns of the managers using regression analysis, we would conclude that both managers had a beta of zero and alpha of +10%. This should seem strange. The two managers are employing completely different strategies, and have completely different ways of creating value for their investors, yet both managers look exactly the same to the traditional regression-based approach to performance analysis. The reason both managers look the same is that even though the macro manager is highly correlated with the market on any given day, both managers are, on average, uncorrelated with the market.

The traditional regression-based approach to performance analysis is often referred to as *ex post* performance analysis. This is because we are evaluating the manager's performance based only on data that we observe after the fact, the realized returns of the manager and the benchmark. Risk-based

performance analysis, rather than looking only at realized returns, also takes into account the risk of the manager each day, that is, how the manager is positioned relative to the benchmark. Because risk-based performance analysis makes use of data that is available before the returns are realized it is often referred to as *ex ante* performance analysis.

To see how risk-based performance analysis works, let's look at the performance of the macro manager again. On any given day, the manager is either long or short the S&P 500. The macro manager only has market risk, and no stock-specific, or idiosyncratic, risk. When the macro manager is long, we would expect him to make money if the market goes up and lose money if the market goes down. At the end of the day, then, we would attribute all of the macro manager's P&L to the market, and none to alpha. We would do this every day, and every day there would be some P&L attributed to the market and none to alpha. At the end of the year, we would add up all of the daily P&L attributions. In this example the macro manager is up 10% at the end of the year, and all of this 10% is attributed to the market, and none to alpha. The risk-based performance analysis is telling us something that the traditional approach could not, that the macro manager's performance is entirely due to market exposure.

For the long/short manager, the opposite would happen. Each day, we would look at the manager's exposure to the market and see that the manager had no market exposure (remember, this long/short manager always has a beta of zero). If the manager made or lost money on any given day, all of the performance would be attributed to alpha. At the end of the year, the long/short manager is also up 10%, but all of the long/short manager's performance is attributed to alpha. The long/short manager's performance is entirely due to stock selection.

The risk-based approach to performance analysis clearly tells us something about the performance of a hedge fund manager which we would not be able to discern using traditional *ex post* analysis. Maybe the traditional approach is providing useful information that we could not get from the risk-based approach, though. Maybe we should use both approaches.

One argument for the traditional regression-based approach is that the traditional *ex post* approach tells us how the manager was positioned on average. In our example, if both the macro manager and the long/short manager are uncorrelated with the market on average, that is all that we should care about. This argument is misleading. First, even though the managers in our example had similar performance last year, they will react very differently to a market crash. In a crash the macro manager will make a lot if he happens to be short that day, but lose a lot if he happens to be long. The macro manager is only neutral on average. With the macro manager it is like drowning in a river that is only three feet deep on average. The long/short manager, by contrast, is not only neutral on average, but neutral every day. The long/short manager should be immune to a market crash. The risk-based approach, not surprisingly then, is a more accurate indicator of risk than the traditional approach. Second, the risk-based approach produces all of the data we need to assess a manager's average exposure to the market, and more. In the current example, the risk-based approach requires us to evaluate each manager's exposure to the market factor each day. With the risk-based approach, we know what a manager's market exposure is on average, how variable that exposure is, and what the exposure was in the most recent period (i.e. how the manager is positioned now). The risk based approach tells us everything that the traditional approach tells us and more.

A potential criticism of risk-based analysis is that the difference between the macro manager and the long/short manager are obvious. The difference between the two managers could easily be determined

in the course of standard due diligence, simply by talking with the managers. Why go to all the trouble of performing risk-based analysis, when it is obvious what factors must be driving the managers' performance. While the managers in our current example have clearly delineated strategies, hedge funds in practice are rarely so simple. Managers in practice are likely to be somewhere in between, with both market and idiosyncratic risk. What's more, their risk profiles are likely to be changing over time. A long/short manager will be low beta one day, neutral the next, and high beta after that. Only a risk-based approach to performance analysis can accurately determine the profit attribution of a manager with a changing risk profile.

One of the greatest advantages of risk-based performance analysis is that it can be run on a daily basis. Traditional ex post analysis can only provide attribution on average, over a long period of time (months or years). Risk-based performance analysis allows you to understand what factors were driving your performance every day, week and month, and to analyze your portfolio at any level, from the top all the way down to the security level.

Risk-based performance analysis can be extended to any number of factors. We can include sector and country factors for our long/short manager, and fixed income and commodity factors for our macro manager. Traditional analysis can also be extended to multiple factors, but it is extremely susceptible to spurious correlations. With risk-based performance analysis, because we have access to the actual positions in the portfolio, it is much easier to prevent spurious results. This is yet one more advantage of risk-based performance analysis.

Conclusion

For hedge funds, risk-based performance analysis offers a number of benefits compared to simple benchmark analysis and traditional ex post analysis. Risk-based performance analysis allows us to understand the actual drivers of hedge fund performance, which allows both managers and investors to make better, more informed decisions.