The double life of correlations

Given the importance many investors place on correlation when it comes to portfolio construction and assessing the diversification potential of assets, Kevin Kneafsey of the Schroders' Multi-Asset team analyses the uses – and misuses – of correlation. In this work he writes about the Dr. Jekyll and Mr. Hyde characteristics of correlation, and demonstrates why it makes sense to be more thoughtful about the use and impact of correlations for asset allocation and position sizing.

Introduction

One of the most frequently used—and misused—statistics in investing is correlation. As with many other helpful things, correlation is commonly misused with the best of intentions. Both managers and investors tend to calculate correlations where data is most plentiful, rather than to assess the relationships that impact our portfolios the most. A surplus of data gives us more confidence in the results, but it may tell us very little about the relationship in which we are most interested. This is especially important because we rely heavily on correlation to assess diversification potential of assets and in portfolio construction. Are we giving correlation more control over our portfolios than is warranted?

To answer this question in a meaningful way, we must think deeply about what correlation is and think carefully about how we use it. Here we consider some of the challenges correlation presents and suggest ways to overcome these challenges. We demonstrate ways in which correlation is used and misused, how it is applied, and where we believe it falls short of our expectations. It appears that correlations lead a kind of double life. Like *The Strange Case of Dr. Jekyll and Mr. Hyde*, this may be a case of good and evil. With correlations, the impact on our portfolios depends on the data used and how we apply the results. Correlation can be a valuable portfolio tool, if applied sensibly. Where we calculate it mindlessly and apply it indiscriminately, it will continue to present a source of instability and vulnerability in our portfolios.

What is correlation?

Correlation is the covariance between two series, scaled by the product of the volatility of each series. Values range from +1.0 (highly correlated), to -1.0 (inversely correlated). Since the standard deviation of any series must be non-negative, the sign of the correlation is determined by the sign of the covariance.

To better understand correlation, we need to better understand covariance. The covariance of two series measures how they vary around their respective average values. If x is greater than the average x value when y is greater than the average y value, then the covariance will be positive. If, on the other hand, the x's are greater than their average when the y's are less than their average, and vice versa, then the covariance will be negative.

Starting with the basics

Correlation measures the ***average relationship*** between two variables ***over some time period***, where the relationship is based on the ***frequency of the data*** (daily, weekly, monthly, etc.) and the mean value of each series ***over that period***. Correlations often are unstable, in part due to the fact that the mean of each series against which each observation is compared may change as the period changes and as the frequency of observation changes.

Figure 1 plots a very simple example of two series, with a correlation of +1.0. The deviations above and below their respective mean values are aligned just so. The chart on the right extends that on the left, by adding just one more data point to each series. This additional data point for each series is a whopper, but it helps make a couple of interesting points:

- The correlation flips from +1.0 to almost -1.0. This highlights the instability of correlations.
- The average of each series changes with the period considered. Adding one more observation to the data window causes the mean to shift up for series 2 and down for series 1. Consequently, even over the period where the correlation had previously been +1.0, the correlation is now negative, as all of the original values of series 2 are now below the new mean of series 2, and all of the original values of series 1 are now above the new mean of series 1.
- Correlations are susceptible to outlier data. The one additional data point was indeed a whopper, shifting correlation from +1.00 to -0.94.

Figure 1: Structure of an insurance-linked instrument

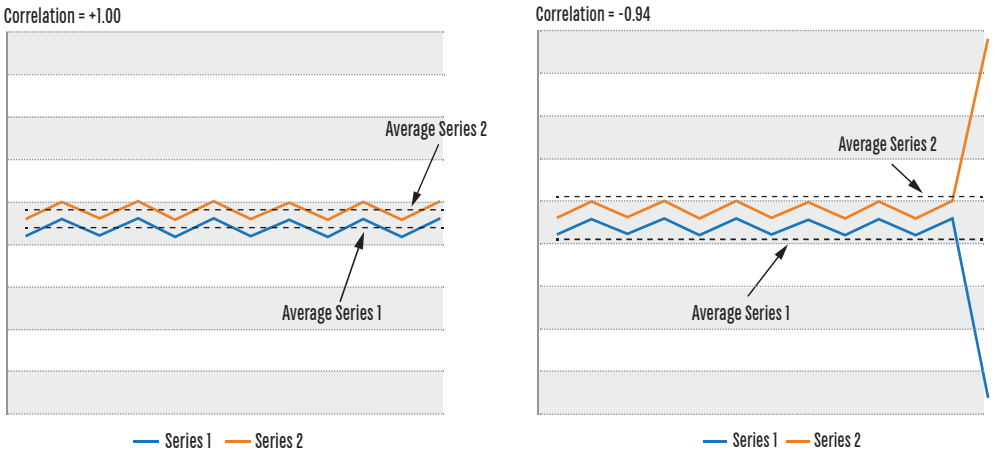

Source: Schroders. Values reflect simulated results and are shown for illustrative purposes only. Does not reflect any actual portfolio.

Time for a reality check

These are dramatic and interesting examples, but what about the real world? To ground our correlation insights in reality, we examine the relationship between the US stock and bond markets. Figure 2 plots the correlation between the S&P 500 and 10-year Treasuries using daily data and a rolling 250-day window (the approximate number of trading days in a year). This provides some evidence of the instability of correlations as over the period 1963–2015; the equity-bond correlation ranges from a high near +0.70 to a low of -0.70.

Figure 2: Stock and bond correlation fluctuates significantly

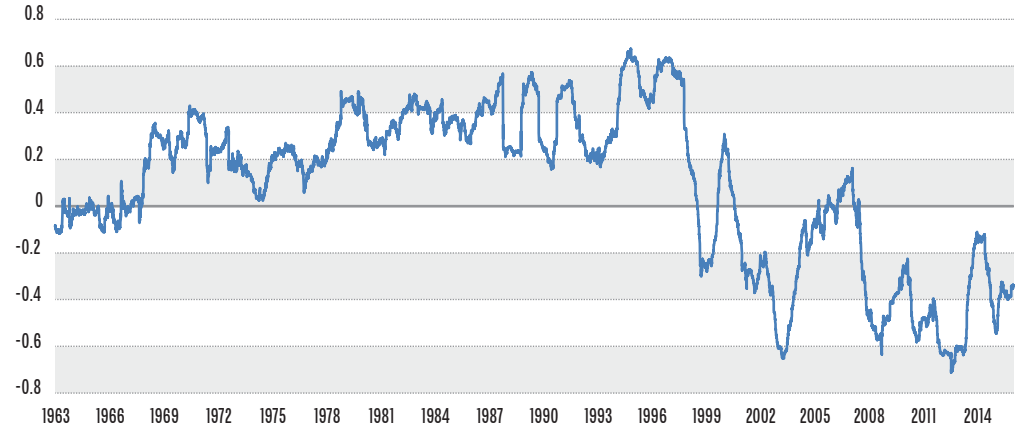

Source: S&P 500 and 10-year US Treasuries; Schroders, Bloomberg and Federal Reserve Bank of St. Louis, 1963-2015.

Figure 3: Stock and bond averages fluctuate broadly

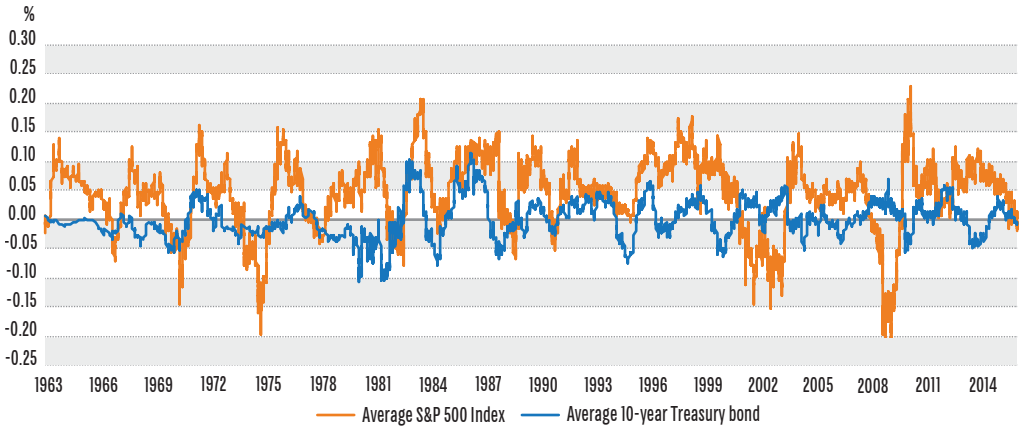

Source: Rolling 250-day average daily returns for S&P 500 and 10-year US Treasures; Schroders, Bloomberg, and Federal Reserve Bank of St. Louis, 1963-2015.

Outliers do influence the results. On Black Monday, October 19, 1987, the S&P 500 fell by more than 20%, while 10-year Treasuries rallied 0.50%. The correlation measured using 250 trading days through the Friday before Black Monday was +0.55. Rolling that 250-day window forward one day to include Black Monday lowers the correlation to +0.29. Given that correlation is bounded between +1 and -1, a drop of 0.26 represents 1/8th of the entire span of the correlation—all driven by data from a single day.[1]

1 This is not precisely true. When a 250-day window is rolled forward, it picks up a new day (in this case, Black Monday), and the first day of the series rolls out of the window. To see the impact of this day that rolled out of the window, we looked at the correlation over the 250 days up through the Friday before Black Monday, and we compared it with the correlation over the 249 days up through the Friday before Black Monday; that is, letting one day fall out the back of the window without moving one day ahead to pick up Black Monday. The correlations are 0.5469 and 0.5467, respectively. They are virtually identical, so the day that rolled out had almost no impact. We can safely argue that the lion's share of the drop from 0.5469 to 0.2880 is driven by the inclusion of Black Monday.

Correlation and diversification

Correlation is often used to assess the diversification potential between various assets. Correlation is helpful for this task, but it requires some thoughtfulness. Consider two series with zero correlation. Intuition would suggest that that an equal-weighted combination of these two series would be very diversifying.

What we find, however, is that the correlation between series 1 and the equally weighted portfolio of the two uncorrelated series is +0.95. What happened to diversification? When we equally weighted these two series with zero correlation, we created a series that looks shockingly similar to one of the original series. In case you are curious, the equally weighted portfolio of series 1 and 2 has a correlation with series 2 of only +0.08 – it is almost zero. This is a valuable clue as to what happened.

Because correlation is calculated by scaling the covariance by the product of the volatilities of each series, some good things happen as well as some bad things. One good thing is it makes correlations much easier to interpret than covariances since correlations are bounded between -1.0 and +1.0. However, this aspect of correlation calculations makes them blind to differences in volatilities of two series.

In the example above, series 1 is approximately three times more volatile than series 2. This happens to be roughly the relationship between the volatility of equities and bonds. Diversification is about balancing risk—not capital—so equal capital allocations resulted in a portfolio that was very skewed toward the more volatile asset, or series 1. While correlation is important for diversification, it is helpful only if it tells us something about the future relationship between assets, at a frequency that matters to us. Even then, it tells us nothing about how to weight the assets to maximize the benefits of diversification.

What is the solution?

The first step in overcoming the shortfalls of correlation is to recognize that the correlation calculation is a mindless crunching of the numbers; consequently, it is subject to garbage in = garbage out. Our solution must start with a better understanding of the complexities of correlation and an active engagement of our brains in this process. It is important to answer a few questions before we crunch any numbers. The first few questions relate to a thoughtful decision about what data to use in the calculation of correlations and why:

1. At what frequency do we care about the relationship between two assets? Our data should be consistent with that frequency.
2. What periods in the past are most representative of what we think the period in the future we care most about will look like? Mindlessly pushing data through the correlation meat grinder is unlikely to provide enlightening results.
3. Given the frequency of interest and the number of historical periods similar to our expected future state of the world, how much data do we have to crunch? The amount of representative data should be an important factor in how much faith we put in the output of the correlation calculation.

We must also determine the sort of relationship we expect between the two assets over the future period of interest. This expectation is our “prior” and we will compare it to what the data suggest. To make an educated statement about the relationship expected between two assets, it is important to think about what drives asset returns. The following simple example may help to clarify this.

For consistency and simplicity, consider again the relationship between equities and Treasury bonds. Let the period of interest be the next one to three years, so we will focus

on fundamental drivers of equity and bond returns rather than more short-term drivers. For this example, we consider three drivers: real economic growth surprises, real interest rate surprises, and inflation surprises.

Figure 4: Correlation depends on how the asset class responds to surprises in fundamentals

Positive Surprise to	Equities	Bonds	Sign of Correlation
Real economic growth	+	–	–
Real interest rates	–	–	+
Inflation	–	–	+

Source: Schroders.

Figure 4 shows the expected response of each asset class to positive surprises in each fundamental factor. It is considered a positive surprise when the fundamental factor is larger than anticipated. The table shows that equities respond positively to positive real economic growth surprises and negatively to higher-than-expected real interest rates and inflation. Government bonds are expected to respond negatively to positive surprises in any of these fundamental factors. The consequence is that, if we expect economic growth surprises to dominate over the next one to three years, then we would expect the equity/bond correlation to be negative. If, however, we expect surprises to real interest rates and/or inflation to dominate, then we would expect positive correlation over the next one to three years. If we are unsure which will dominate, we would expect correlations closer to zero.

If the idea of relying on macroeconomic expectations to derive more precise correlation statistics makes you uncomfortable, consider the alternative: As many investors do, take a bunch of data from the past, without any thought about its relevance to the future. Select a frequency that offers ample data, regardless of the frequency of interest of the relationship between the assets. Put this through the correlation meat grinder and trust the output completely. Doesn't it make more sense to start with a prior about what you expect to see, then think carefully about what data are most relevant for your decision, and then proceed with a combination of strong intuition and careful analytics?

Correlation forecasts: High expectations for a single statistic

One of the most common uses of correlations is in portfolio construction, forecasting how assets will move relative to each other to help determine asset allocation. Typically, the recent past is used as a forecast of the near-term future. Given what we now know about the tendency for correlation to exhibit instability, variability, and vulnerability to outlier data, is this a sensible approach to one of the most important decisions in investing?

To answer this question, let's return to the commonly used stock-bond correlation, and use a 250-day correlation as the forecast of the correlation over the following 250 days, during 1962–2015. The table in Figure 5 lists the percent of observed deviations over the entire period and over the first and second half of the period that exceed certain thresholds. For example, the column labeled 0.20 shows that, over the full period, nearly 1/3rd of the realized correlations were outside the range of the forecasted correlation +/- 0.20. Over the second half of the period, one-half of the realized correlations were outside this range. For example, if the forecasted correlation was +0.30, half time the realized correlation did not fall in the range +0.10 to +0.50. Clearly, data from the recent past has not been particularly useful in forecasting even the near-term future for a volatile relationship like that of equities and bonds. This is especially true recently.

Figure 5: Recent stock/bond correlation has not been very predictive

Percent of forecasts deviating from actual by more than +/-					
	0.10	0.20	0.30	0.40	0.50
Full period	60%	31%	17%	8%	3%
First half	45%	14%	3%	0%	0%
Second half	76%	49%	32%	16%	7%

Source: Schroders, Bloomberg, and Federal Reserve Bank of St. Louis, based on the S&P 500 and 10-year US Treasury yields between 1962-2015.

Is the same true for more stable asset relationships? Consider the correlation between the exchange rate of dollar/pound and dollar/euro. Again calculating 250-day correlations, the mean correlation (since the euro was introduced in 1999) is 0.65, ranging from just below 0.40 to just above 0.80. In this case, data from the recent past was a much better predictor of the near future – the historical correlation of spot exchange rate percent changes was a relatively good predictor of near-term future correlations, based on daily exchange rate data from November 1998 through mid April 2016. While 1/3rd of realized correlations fell outside the forecast correlation +/-0.10, only 3% were outside +/-0.20.

Figure 6: Recent correlation of USD/GBP and USD/EUR spot rates is more predictive

Percent of forecasts deviating from actual by more than +/-			
	0.10	0.20	0.30
Full period	33%	3%	0%

Source: Schroders and IMF, between November 1998-April 2016.

This is encouraging from a forecasting point of view, but it may still prove problematic for portfolio construction. For example, take a long/short portfolio in which the exposure to each currency pair is scaled to a target volatility. The position sizes with a correlation forecast of 0.40 are half the size of those with a correlation forecast of 0.85. This relationship, between holding size and correlation, is highly nonlinear. When the correlation forecast climbs to 0.95, the positon size required to achieve the target risk is three times that required with a 0.40 correlation. Given the difficulty in forecasting correlations and the instability of correlations, it is surprising how much influence they have on portfolio construction. We believe position sizes should reflect conviction and the opportunity for gains, and less so noisy correlation forecasts.

Emerging from the double life

There is a Dr. Jekyll and Mr. Hyde relationship with correlations that we have in investing. On the Hyde side, there is the mindlessness of the correlation calculation and the data chosen, the instability of correlations, and the difficulty in forecasting correlations accurately. Here, we give in to the basic human impulse to simplify complex choices in investing and to avoid the discomfort that goes with uncertainty.

On the Jekyll side, we give correlations vast authority over our portfolios, impacting such things as asset choice, position sizing, and our expectations of diversification potential.

Conclusion:

How should we deal with the double life of correlation, which at times is very useful but can also be very problematic? In our view, this Dr. Jekyll and Mr. Hyde nature of correlations must first be recognized and then managed with two steps:

Step 1: Take more ownership of the correlation forecast. This means being more thoughtful about the data employed, including the period covered and the frequency of the data. It also means starting with an intuitive prior about the correlation against which the data-driven correlation can be compared.

Step 2: Control the impact that correlation has on portfolio construction. This is especially important where the impact of correlation on holdings is large and nonlinear.

Such an approach of combining strong intuition and careful analytics gives us a clearer perspective to recognize the ways in which correlation can improve our portfolios and how it could harm them.

Kevin Kneafsey, PhD
Multi-Assets Investments, Senior Adviser

Kevin Kneafsey joined Schroders in October 2012 as a Senior Adviser, based in the US, providing support to the Multi-Asset Investments & Portfolio Solutions team.

Previously, he was Head of Research for Blackrock's Multi-Asset Client Solutions (BMACS) team with primary responsibility for Market Advantage (BlackRock's optimal beta strategy), MuSt (BlackRock's internal hedge fund of funds, and a systematic trading strategy iShare. Before that, Kevin was the lead strategist for Global Ascent, BGI's flagship global macro hedge fund. Kevin joined Barclays Global Investors in 1994 and held several research, strategy and portfolio management positions including Chief Strategist, Global Market Strategies Group and Head of Research, Client Solutions.