MARKET INSIGHTS

September 2011

Using correlations in asset allocation

'History doesn't repeat itself, but it does rhyme' - Mark Twain

Tom Elliott
Global Strategist

Introduction

Measuring the correlation between assets has often been criticised as rear view mirror driving. The past is no guide to the future. But this is a naïve truism that misses the point. A historical analysis of financial markets shows that there are discernible relationships between individual assets, and between asset classes, and that these relationships, or correlations, do tend to persist over time. Knowledge of past correlations can therefore help identify diversification opportunities, which in turn can maximise a portfolio's return for a given level of risk, or minimise risk for a given level of return, within Modern Portfolio Theory (MPT).

But correlation tables within asset allocation have their limitations, and an understanding of these is important to avoid investment errors. This paper looks at two aspects of correlation tables that, in conversations with clients, have caused the most misunderstanding.

First, the choices of data frequency and of the time period used when calculating a correlation are both subjective decisions that carry with them potential hazards.

Second, stronger correlations among assets exist in bear markets than in bull markets, which is precisely opposite to the needs of investors. Combine this with the observation, from behavioural economics, that investors dislike losses more than they like gains of the same magnitude, and it is clear that some form of bias in the calculation of the correlation data may be appropriate to compensate for these two asymmetries.

The overriding message is to use correlations carefully. Understand the methodology used, know the volatility of a correlation as well as its value at any one time, and perhaps assume a stronger correlation than that produced by the maths to take into account certain asymmetries.

The underlying maths are not important to this paper, but are offered in summary form in the Appendix.

Exhibit 1 – **Correlations of returns (EUR)**

Upper-left value: weekly returns over previous ten years; lower-right value: weekly returns over previous one year.

	MSCI Europe	MSCI Japan	S&P 500	MSCI EM	MSCI World	Pan Europe bonds	EM debt	High yield bonds	Precious Metals
MSCI Europe	1.00	0.47 / 0.39	0.80 / 0.73	0.80 / 0.57	0.93 / 0.84	-0.20 / -0.20	0.37 / -0.09	0.53 / 0.49	-0.04 / -0.05
MSCI Japan		1.00	0.48 / 0.68	0.56 / 0.54	0.59 / 0.71	-0.05 / 0.05	0.33 / 0.47	0.35 / 0.16	0.00 / 0.26
S&P 500			1.00	0.68 / 0.62	0.95 / 0.95	-0.23 / -0.12	0.54 / 0.44	0.42 / 0.26	-0.05 / 0.16
MSCI EM				1.00	0.83 / 0.77	-0.22 / -0.05	0.49 / 0.33	0.51 / 0.38	0.10 / 0.27
MSCI World					1.00	-0.22 / -0.12	0.51 / 0.33	0.52 / 0.38	-0.02 / 0.18
Pan Europe bonds						1.00	0.07 / 0.33	-0.03 / 0.06	-0.01 / 0.17
EM debt							1.00	0.29 / -0.20	0.02 / 0.29
High yield bonds								1.00	-0.06 / -0.08
Precious Metals									1.00

Weekly returns over previous ten years

Weekly returns over previous one year

Note: Pan Europe bonds is the Citigroup Europe World Government Bond Index (WGBI). Emerging market debt is the JPMorgan EMBI Global, high yield bonds is the JPMorgan Domestic High Yield. Precious metals is the GSCI precious metals index. Ten year correlation values are calculated on weekly return data in euro for period 03/08/01 to 29/07/11, one year correlation values are calculated on a weekly basis in euro for period 06/08/10 to 29/07/11. Correlation measures the direction and degree of linear association between two variables. All indices are total return. Source: Bloomberg, J.P. Morgan Asset Management.

What are we measuring?

When we speak of the correlation of return and volatility between two assets, or asset classes, we are actually referring to the degree to which the two move in step with each other and the direction of that relationship. Importantly, what we are not saying is that the movement of one asset *causes* the movement of another. We are measuring how two assets respond similarly to the same external conditions.

The scale used is between +1, which would suggest a perfectly positive relationship with two assets moving in step and in the same direction, and -1, which would be a perfectly negative relationship with two assets moving in step in opposite directions. A score of 0 suggests no relationship at all, while anything over around 0.4 may be described as showing a relationship (assuming weekly or monthly data points are used).

Exhibit 1 shows both one- and ten-year correlations for a variety of asset classes. It should be mentioned at this point that while the one-year numbers are a useful guide to recent trends between asset classes, and help explain what has happened in the markets, they have little value in forecasting due to their greater volatility. Instead, longer time periods are used for forecasting potential risk and return within a portfolio.

We see several recognisable patterns, such as a generally high correlation between developed blue chip stock markets that suggests little diversification benefit. Meanwhile, European government bonds and precious metals have, respectively, a weak correlation and no correlation to the other asset classes shown, which suggests some diversification benefits in holding these assets within a portfolio.

Two sorts of problems affecting the use of correlations: time and asymmetry

1. Problems associated with time: the frequency of data points and the length of time periods used

It is generally considered that the more data points, and the longer the period being considered, the more reliable the correlation. But there are no hard rules, meaning that users of correlation tables need to look closely at the underlying methodology.

The frequency of data points Different correlation numbers will come from altering the frequency of data points (ie daily, weekly or monthly), also known as observations. As **Exhibit 2** shows, less frequent observations generally, though not always, deliver higher correlations.

Exhibit 2 – **Correlations using different frequencies**

Ten years to July 2011				
S&P 500 vs:	Yearly	Monthly	Weekly	Daily
GSCI Precious Metals	0.35	0.01	-0.05	-0.02
High Yield	0.75	0.47	0.42	0.15
Pan-Europe bonds	-0.11	-0.27	-0.23	-0.10

Source: As for Exhibit 1. Bloomberg, J.P. Morgan Asset Management.

This paper uses asset returns based on weekly observations, though monthly data points are often preferred if looking only at long time periods. What you cannot do is to compare a short-term correlation using weekly data points directly with a long-term correlation using monthly data points.

The length of time periods used How long should the observation period be? If too short a period is used, an investor risks relying on very volatile numbers. Yet if too long a period is used, the data may become meaningless due to structural changes that can take place within asset classes over the period.

Let's look at using short periods for calculating correlations. A short-term investor, looking for two assets that exhibit a low correlation, may feel that recent data are more meaningful than, say, ten years of data. They might focus on a one year correlation.

Exhibit 1 shows ten- and one-year correlations for a range of asset classes up to the end of July 2011. The investor sees that the correlation between the S&P 500 and the precious metals components of the Goldman Sachs Commodity Index (GSCI) was just 0.16. This would not be considered a meaningful relationship, suggesting diversification benefits.

But the pitfall of using a one-year correlation is the volatility. This is immediately apparent in **Exhibit 3 (next page)**, in which we see that the one-year correlation of the S&P 500 and the GSCI is highly volatile compared to the ten-year correlation. From being slightly negative at the start of 2011, the one-year has been over 0.5 for most of the subsequent period before falling sharply in July. Clearly, the forecasting value of a short period correlation to an investor is limited. Calculations based on longer periods have an advantage through eliminating near-term market 'noise'.

One of the explanations given for the credit crunch in 2008 was the lack of long-term trading data for new derivative financial products. This led to a dependency on relatively short periods of data, which did not include periods of acute market stress, to calculate correlations. Therefore the correlations were lower than they might otherwise have been, had market data stretching back to the Asian financial crisis and the collapse of the internet bubble been available. In turn, this led to an underestimation of the level of risk these assets had within calculations based on MPT.

But how long is the optimal period for a correlation calculation for forecasting purposes? A few economic cycles are usually thought necessary. However, a cut-off point has to be taken to avoid absurdities, since over time external economic changes alter relationships between financial assets, steadily reducing the value of data from a given point in time.

For example, as a developing economy matures, its stock market can be expected to become broader, more liquid and hence less volatile relative to its own history and to a developed country's stock market. This may result in a stronger correlation between the two stock markets over time.

Exhibit 3 – **Rolling correlations between S&P 500 and GSCI precious metals**

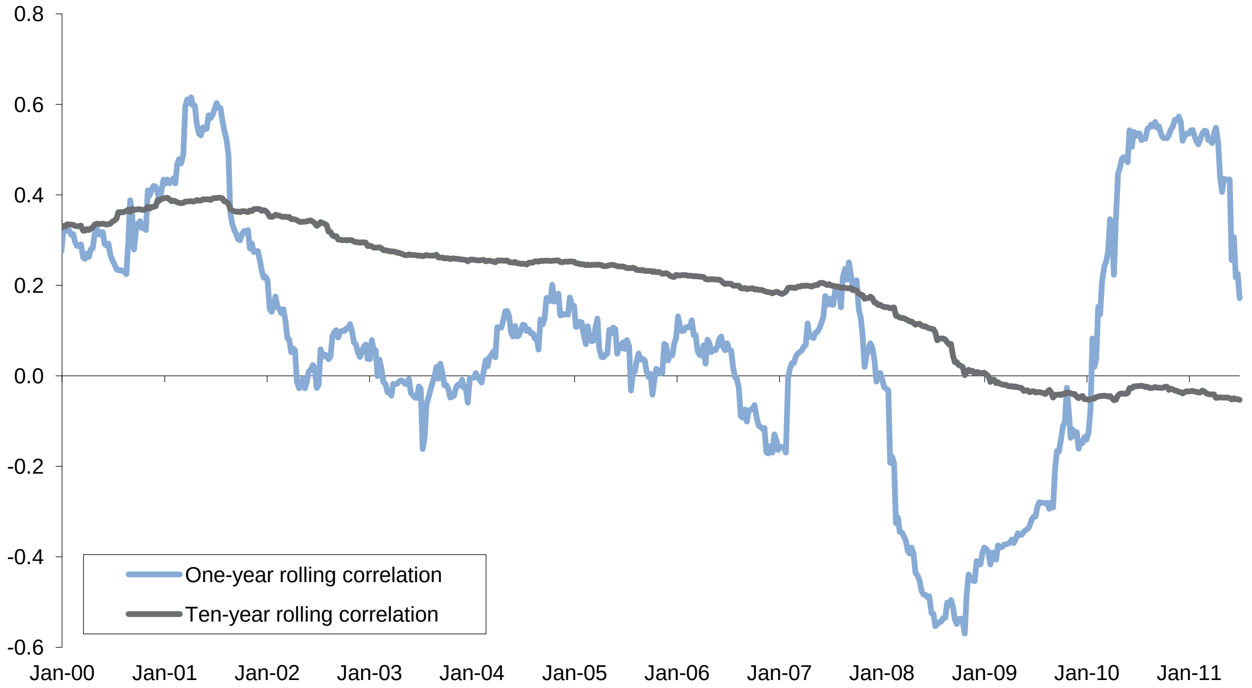

*total return indices in EUR
Source: Goldman Sachs , S&P, J.P. Morgan Asset Management.

Therefore, if we are using a correlation number as a guide to the future relationship between two assets, how relevant is previous stock market price history if 30 years ago the stock market consisted of just a cement works, a bank and a distillery? Should we just look at the previous five years?

Clearly the investor has to make a compromise, and periods of around five and ten years are considered normal.

Within developed stock markets the increased globalisation of financial and capital markets have led to stronger correlations in recent decades, another example of a long-term structural change taking place.

2. Problems associated with asymmetry: changing correlations during market cycles and investor sentiment towards losses

Two different asymmetries need to be borne in mind when considering correlation tables, both of which suggest that an investor may wish to apply a greater weighting to short-term periods of stress within a long-term correlation.

Changing correlations during market cycles We see stronger correlations between assets during times of market stress and high volatility, and weaker correlations during times of market strength and low volatility. This asymmetry happens because during periods of market stress investors are susceptible to bouts of panic (at this point one must put aside the assumption within MPT, that investors are at all times rational). Consequently, investors may sell *all* risk assets – not just those immediately affected by the crisis - as they scramble for safe havens, such as cash and G3 government bonds. During bull markets we tend not to see general flights into risk assets, but some favoured over others and hence a greater diversion of returns and lower correlations.

This relationship is demonstrated in **Exhibit 4**, which looks at the correlations of various asset classes against the DAX index of German equities before, during and after the global financial crisis that followed the collapse of Lehman Brothers. All but one of the correlations against the DAX strengthened during the crisis, with the greatest increase occurring in assets that traditionally had no, or only a low, correlation with equities, such as precious metals (eg, gold), commodities, high yield and emerging market debt.

Exhibit 4 – **Correlations before, during and after the credit crunch crisis**

DAX correlation with	Long term Jan 00 – Jul 11	Pre-crisis Jan 00 – Jun 08	Crisis July 08 – Jul 09	Post-crisis Aug 09 – Jul 11
FTSE 100	0.82	0.78	0.91	0.78
CAC 40	0.91	0.89	0.96	0.92
S&P 500	0.74	0.71	0.82	0.70
MSCI World	0.85	0.83	0.93	0.81
MSCI Europe	0.92	0.91	0.96	0.92
MSCI Japan	0.39	0.39	0.54	0.15
MSCI Asia ex Japan	0.63	0.58	0.81	0.48
MSCI EM	0.74	0.67	0.89	0.66
German Bunds	-0.41	-0.38	-0.44	-0.58
Pan Europe bonds	-0.25	-0.32	-0.12	-0.22
EM debt	0.34	0.30	0.54	-0.01
High yield bonds	0.45	0.31	0.65	0.50
Commodities	0.21	0.02	0.54	0.44
Precious metals	-0.01	0.05	-0.23	0.13

Source: Bloomberg, J.P. Morgan Asset Management.

The exception in the table is for pan-European government bonds, where we see the negative correlation weaken during the crisis. This reflects the fact that the Italian government bond market exhibited a positive correlation with German equities and other risk assets over the six weeks following the Lehman Brothers collapse (they fell in value, and the yield increased). This was in contrast to German Bunds which rose in value over the same period, benefiting from their safe haven status.

The sharp rise of correlations between previously low, or non-correlated, assets surprised many investors and led to criticism at the time of MPT as being a form of rear-view mirror investing. However, as the opening remarks in this paper suggest, to ignore patterns of past behaviour among asset classes because they can at times go sharply askew seems to be wilfully ignoring an important source of data just because it is not perfect.

Since the increase in correlation between assets occurs precisely when it is least wanted by an investor, and may wipe out a leveraged investor, perhaps some adjustment within the calculation is warranted, which will give greater weight to past periods of market stress.

Furthermore, some adjustment may be warranted by the fact that gains made during a bull market often take longer to accrue than losses of a similar magnitude made in downturns. This leads to a bias in favour of the data covering a bull market, and to lower correlations for the period as a whole, than if markets moved up and down over similar durations.

Investor sentiment towards losses An additional reason why correlation numbers may understate potential problems arises from investor 'loss aversion'. The field of behavioural economics has attempted to quantify the degree to which investors dislike a loss - even a minor one - much more than they like a gain of a similar amount. As **Exhibit 5*** shows (next page) a kink has been introduced into the standard utility curve (in the bottom left square), to demonstrate this asymmetry of response. The chart shows that the level of utility, or happiness, arising from a gain of a certain sum of money (say, EUR 5), is less than the unhappiness we feel when we lose the same amount of money, with the link being the particularly sharp loss aversion seen at the initial stages of loss.

*Eg, 'Loss Aversion, Asymmetric Market Co-movements and the Home Bias'. Kevin Amonlirdviman and Carlos Carvalho, Federal Bank of New York Staff Reports no.430 February 2010.

Exhibit 5 – **Investor marginal utility**

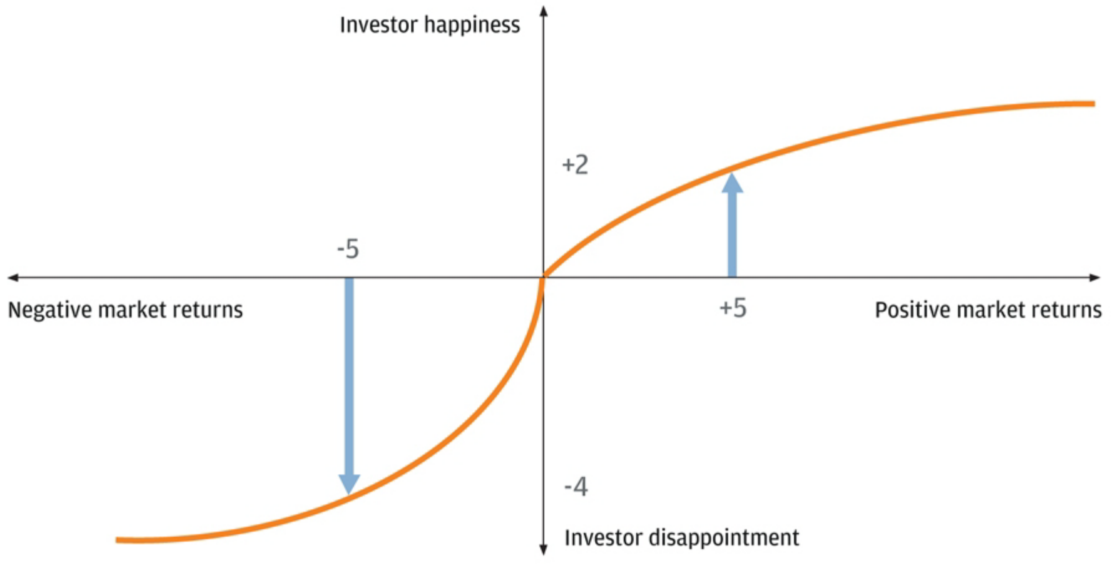

Source: J.P. Morgan Asset Management.

Loss aversion confirms the need for a conservative bias in correlations, but this time the problem stems from slightly irrational human behaviour.

Alternatives to correlations

Correlations are not the only tool available to investors to identify relationships between asset classes. But they are perhaps the most straightforward to understand and to use, once their limitations are understood. For investors interested in other ways of expressing co-movement between assets, cointegration (a means of quantifying the long-run structural relationship between variables) and copulas (based in probability theory) are two of a number of methods that may be considered. However, they too rely on assumptions that must be properly understood before being used.

Summary

Correlations between assets have a key role in portfolio construction, helping investors to maximise returns and minimise risk within a diversified portfolio. However, investors need to be aware of the methodology used to calculate correlations, such as the frequency of data, the time periods observed, and the consistency of comparisons from different sources. Correlations based on short-term data have little forecasting value due to their volatility, and longer periods that include periods of market stress are preferable.

An asymmetric tendency for assets to have a higher correlation during a period of market stress, compared to a period of market gains, suggests that investors should assume correlations to be higher than the numbers suggest. This is particularly true if the investor is leveraged. This is echoed by studies into 'loss aversion', which suggest that a short, sharp downward turn in the market may also deter an investor who has only lost a small proportion of his wealth from future investing.

Appendix: The maths

Correlation defines the strength of the relationship between two (random) variables. Correlation tables commonly use the return for each variable over a set period, eg ten years, with a defined frequency, eg daily, weekly or monthly returns.

The correlation of returns of two assets is calculated by dividing the covariance of the two returns (of assets X and Y), by the product of their standard deviations and the number of observations. This produces a scaled measure of movement bound between -1 and +1 which is a more useful indicator than covariance itself.

Calculating the correlation *(φx,y)*

$$\varphi_{x,y} = \frac{Cov(x,y)}{s_x s_y n}$$

The covariance is the degree by which the two variables move together and is defined as the sum of the product of each variable less its average/mean value for the specified time period, ie sum[(X – mean X)(Y- mean Y)].

Calculating the covariance *(Cov (x,Y)*

$$Cov(x,y) = \sum_{i=1}^{n} [(x_i - \overline{x})(y_i - \overline{y})]$$

Where:

φ(x,y) = correlation of x and y (rho)

Cov(x,y) = covariance of x and y

s_x and s_y = sample standard deviation of x and y

$\overline{x}$ and $\overline{y}$ = average x and average y over sample

n = number of observations in the sample

Tom Elliott, *vice president,* is a global strategist within the Investment Marketing Team at JPMorgan Asset Management, responsible for investment communications through the Guide to the Markets suite of products. An employee since 1995, he worked in the Global Multi-Asset Group (GMAG) until 2006 and before that he was head of the Investment Writing Team. Previously, he worked at Euromoney Publications as a feature writer for a year and prior to that he spent four years at Greig Middleton & Co. as a graduate trainee and securities analyst. Tom obtained a BA in History from Sussex University and an MSc in Economic History from the London School of Economics.