Polymetis LLC

The Statistical Mystique

"It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so."

Mark Twain

It is a scene replayed tens of thousands of times each day. Once aboard his plane, the experienced airline traveler works diligently to ignore the safety card in the seat pocket in front of him. Despite the best efforts of conscientious flight attendants to call attention to its utility and its significance, the intrepid traveler is too busy, too burdened, too exasperated, or simply too confident to review the requisite safety information. Besides, our sojourner has read the card one or twice previously, and has absorbed its contents via osmosis dozens, even hundreds, of times before.

Regrettably, hedge fund investors often act like these frequent fliers, ignoring valuable information because they are either too busy or too sophisticated to devote precious time and effort on something they already know well, or at least think they know well. This is especially true regarding the use of statistical data and information in evaluating managers. While this can lead to unfortunate outcomes, the situation is easily correctable.

Since the publication of Harry Markowitz's seminal "Portfolio Selection" in the Journal of Finance in 1952, the genesis of modern portfolio theory ("MPT"), the application of statistical and quantitative methods to asset management has increased dramatically. That growth has been enhanced in the past few decades by the exponential growth of computing power, the convenience of user-friendly software, the accessibility of vast amounts of data, and the proliferation of innovative and functional research.

This has been a positive net development, but it has not come without certain costs, particularly for the existing or potential hedge fund investor. The explosion in the availability of facts and figures, however, has not necessarily led to a commensurate increase in useful information. If anything, there may be too much data to process efficiently, and too many ways to slice and dice them. Statistics are neither informative nor useful just because they can be calculated. Furthermore, chasing after every possible measure consumes investors' most valuable resources, time and attention, and can lead to paralysis from analysis.

One useful tactic to combat this potential data and analytical overload is to, every now and again, refocus on the core quantitative elements embedded in the manager evaluation and portfolio construction process. It is useful, therefore, to focus on the "big three" of asset management statistics: mean, standard deviation, and correlation coefficient. Why these three? Not only are they the most widely-used and accepted investment statistics, but according to MPT, one can design efficient portfolios of assets using only these inputs.

Let us re-examine these three key measures with a fresh perspective, and with a minimum of mathematical theory, to see what they can tell us.

It's Just Not Normal

Before focusing individually on the three main statistics noted above, it is worthwhile to address the proverbial white elephant in the room. There exists a large body of academic and practical research that has demonstrated that hedge fund returns typically are not normally distributed, a key assumption underlying MPT. Instead of forming the classic bell curve, hedge fund returns, on the whole, form a curve that has negative skew. Skew is a measure of a distribution's symmetry. While the normal distribution is perfectly symmetrical (Figure 1 left), a negative skew distribution has a tail (or taper) that extends further to the left (Figure 1 right). In addition, hedge fund returns also exhibit excess kurtosis. Kurtosis is a measure of a distribution's peakedness. While a normal distribution is said to be mesokurtic (Figure 2 left), hedge fund returns typically exhibit leptokurtosis (Figure 2 center), which means they have a higher peak and fatter tails (a higher number of extreme events).

Figure 1

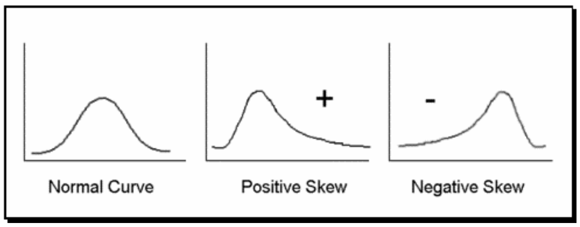

Figure 2

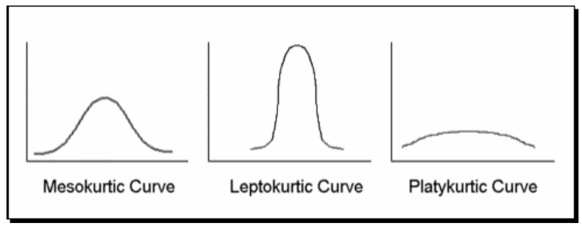

There are several reasons that account for this general condition. Many hedge fund strategies employ either leverage or derivatives which in extreme, or even near-extreme conditions, can cause disproportionate movements (non-linear) in return relative to regular underlying returns. This especially is the case with hedge fund strategies that implicitly or explicitly are short volatility (effectively betting that volatility will remain stable or fall). These are the so-called "short option" (aka "nickel in front of the steamroller") strategies which make consistent small gains, but which suffer occasional, though not rare, losses of considerable or even catastrophic proportion.

The returns of a number of hedge fund strategies also exhibit considerable serial correlation. Serial correlation, or autocorrelation, is a phenomenon whereby a data series exhibits a high degree of correlation with itself over time. Practically speaking, in the hedge fund space, this means that any month's return will be strongly dependent upon the previous month's return. As several studies have shown[1], the most likely cause for serial correlation arises from the pricing of illiquid securities. In any event, serial correlation is a further negation of the normality assumption and inconsistent with the efficient market theory.

Even though hedge fund returns in general are not normally distributed, consideration of the big three remains relevant. First, MPT, although imperfect, remains an extremely practical construct that is both intuitive in concept and wide-spread in use. Second, there is some research[2] suggesting that the negative impact of using the normality assumption is not always excessive. Third, the deviation from normality in the hedge fund return space is not uniform across strategy. Some strategies are more normal than others. In any event, since MPT and the big three are in use, and will continue to be used for the foreseeable future, it makes sense to take things as they are. Paraphrasing Dorothy Parker, they may not be normal, but they are common.

[1] Getmansky, Lo, and Makarov (2003) , De Souza and Gokcan (2004), Brooks and Kat (2001), amongst others.
[2] Hood and Nofsinger, The Normal Equivalent: Evaluating Non-Normal Portfolio Characteristics, January 2006

I Mean It This Time

Let us begin by reviewing the equation used to calculate the mean of a data sample (Figure 3). The mean, commonly known as the average, is the sum of all the elements in the sample divided by the number of elements in the sample.

Figure 3

$$mean = \frac{\sum_{k=1}^{n} x_k}{n}$$

The equation for the mean is the least complicated, most intuitive, and the best understood of the three. The mean serves as the best guess estimate for future results.

Valid questions have been raised about the aggregate performance of hedge fund indices due to performance reporting and index construction issues, but these do not apply at the manager level. The major concern affecting the utility of the mean is the sample size. The absolute minimum traditionally considered to be statistically meaningful is 30 elements, or two and one-half years of monthly data. Substantially longer track records (larger sample sizes) are desirable.

Of course, this raises a not inconsequential question. Does a long track record constitute a single sample? If a manager has a track record in excess of fifteen years, for example, are the data points from eight to fifteen years ago relevant if the manager's strategy has evolved over time or underlying market conditions have changed drastically? It calls into question whether recent data is drawn from the same distribution, or whether the entire track record may comprise several separate samples. So, even for the most straightforward of the three statistics, and in a best case (large sample size) scenario, there can exist a potential problem in application.

Deviation Without Standards

Now, let us review the equation used to calculate the standard deviation of a sample (Figure 4). Often described as the dispersion of the data centered around the mean of a sample, the standard deviation is calculated using the sum of squares method.

Figure 4

$$S = \sqrt{\frac{\sum_{k=1}^{n} (x_k - \bar{x})^2}{n-1}}$$

The numerator calculates the distance or deviation from the mean for every single data point in the sample. These distances are squared so that dispersion one way (above the mean) does not cancel out dispersion in the opposite direction (below the mean). The sum of the squares is scaled or standardized by dividing by the number of elements in the sample[3]. This insures that the amount of deviation is not dependent upon the size of the sample, and so samples of different sizes can be compared on an equivalent basis. Finally, the square root of the dividend (numerator divided by denominator) is calculated in order to offset the earlier squaring of the individual dispersions.

The standard deviation equation is not only a bit more complicated than that of the mean, it actually contains the mean (the $\bar{x}$ or x-bar symbol in Figure 4) as an element in its calculation. So, any potential issues regarding the validity of the mean are thus incorporated into the standard deviation. Practically, the standard deviation provides a way to bound expected outcomes. For example, in a normal distribution, about 68.3% of the values will fall within 1 standard deviation of the mean (i.e. the mean ± the standard deviation). Similarly, about 95.5% of the values will fall within 2 standard deviations of the mean. The remaining 4.5% of the sample is split equally between the tails on the right and the left. One implication is that an extreme negative outcome (no better than, but potentially much worse than, the mean minus 2 times the standard deviation) can be expected to appear about 1 out of every 44 times. Applied to the hedge fund space, this means a "blow-up month" roughly every 3 ½ years is highly probable.

As noted above, hedge fund returns typically show a degree of serial correlation, due in some combination to illiquidity and to subjective pricing methods such as mark to model (a theoretical price) or securities priced to quote (a "phony" price offered by a dealer who would never actually transact at the price). The impact varies considerably by strategy[4]. The net result is an understatement of the true volatility of the underlying returns. As a measure of risk, standard deviation, as well all other risk measures (such as the Sharpe ratio) calculated using the standard deviation, consistently will underestimate risk.

It's Not Too Late To Correlate

Finally, let us review the equation used to calculate the correlation coefficient ("correlation") of one investment to another (Figure 5). This obviously represents an even more complicated equation than that of either the mean or the standard deviation. Bearing in mind that the value of correlation ranges from -1.00 to +1.00, let us concentrate on one very important aspect of the calculation.

Figure 5

$$r = \frac{\sum_{i=1}^{n} (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\left(\sum_{i=1}^{n} (X_i - \bar{X})^2\right)\left(\sum_{i=1}^{n} (Y_i - \bar{Y})^2\right)}}$$

By looking at the equation, we can see that the denominator is a square root. Harkening back to our knowledge of algebra, we recall two relevant facts. First, the value of a square

[3] For extremely technical reasons, the sum of squares actually is divided by the sample size minus one (the n-1) rather than n, the total number of elements.
[4] Pedro Gurrola Perez, An Approach to the Non-Linear Behavior of Hedge Fund Indices Using Johnson Distributions, (2004).

root must be a positive number. Second, the sign of a fraction will be positive when both the numerator and the denominator share the same sign. Since, the denominator in the correlation equation must be positive, it is the value of the numerator that will determine whether the correlation will be positive or negative. If the numerator is negative, the correlation will be negative. If the numerator is positive, the correlation will be positive.

Let us focus further on the numerator while recalling still one more relevant fact from our knowledge of algebra. Note that the numerator is the sum of a series of products (the x dispersion or differential times that of the y). Just as with the division inherent in the fraction, the product of two numbers will be positive when both are positive or both are negative. They will be negative when one is positive and one is negative. The x and the y values must be opposite for the correlation to be negative. This makes intuitive sense. Now, all this review of algebra brings us to the salient point. The values are calculated relative to their own means, the x-bar and the y-bar, respectively, not to a common reference point. This can lead to a faulty understanding of the relationship.

The simple but extreme illustration shown in Table 1 and Table 2 below will reinforce the point. Table 1 shows 18 months of performance data for three hypothetical managers. Glance at the return information without looking at the correlation matrix that appears at the top of the following page in Table 2. In your mind, estimate what you would expect the correlation to be for the three possible pairs: Mgr 1 – Mgr 2, Mgr 1 – Mgr 3, and Mgr 2 – Mgr 3. Remember, the correlation ranges in value from -1.00 to +1.00. If the correlation is close to 0, it (essentially) means there is no relationship between the variables. If the correlation is positive, it means that as one variable gets larger, the other gets larger. If the correlation is negative, it means that as one gets larger, the other gets smaller. The closer the correlation is either to +1.00 or -1.00, the more closely the two variables are related. At the extreme values, +1.00 or -1.00, the two variables are said to be either perfectly correlated or perfectly inversely correlated. Now, make your estimates and consult the correlation matrix in Table 2.

Table 1

Month	Mgr 1	Mgr 2	Mgr 3
1	1%	-1%	2%
2	1%	-1%	2%
3	1%	-1%	2%
4	1%	-1%	2%
5	1%	-1%	2%
6	1%	-1%	2%
7	1%	-1%	2%
8	1%	-1%	2%
9	2%	-2%	1%
10	1%	-1%	2%
11	1%	-1%	2%
12	1%	-1%	2%
13	1%	-1%	2%
14	1%	-1%	2%
15	1%	-1%	2%
16	1%	-1%	2%
17	1%	-1%	2%
18	1%	-1%	2%

Table 2

	Mgr 1	Mgr 2	Mgr 3
Mgr 1	+1.00	-	-
Mgr 2	-1.00	+1.00	-
Mgr 3	-1.00	+1.00	+1.00

Do the results surprise you? The values in Table 2 are quite correct. Allowing some liberty, my guess is that, quants and cheaters aside, you correctly estimated the Mgr 1 – Mgr 2 correlation, but guessed exactly opposite to the results for the Mgr 1 – Mgr 3 and Mgr 2 – Mgr 3 correlations. How is this possible when Mgr 1 and Mgr 3 were positive in every month while Mgr 2 showed negative results in each month? Doesn't a perfect correlation of +1.00 means that investments will move in perfect synchronicity, up and down together, while a correlation of -1.00 will mean that the two will move in opposite ways, one always zigging while the other is zagging?

Remember, correlation measures the relationship between the managers relative to their own reference point (i.e. their own mean). The mean or average return for the three managers were +1.06%, -1.06%, and +1.94%, respectively. We now can see in every instance except for Month 9 that Mgr 1 was below his mean. Conversely, Mgr 2 and Mgr 3 outperformed their mean in every case except for Month 9. So, even though the returns for Mgr 1 and Mgr 3 were quite alike, one always was zigging while the other was zagging relative to each one's own mean. Conversely, Mgr 2 and Mgr 3, whose returns were almost polar opposites, did actually move in perfect tandem relative to their respective means.

The correlation statistic also faces other problematic issues in its application to the hedge fund arena. For example, correlations may not be stable over time, changing in response to underlying financial market conditions or to changes in manager style or execution. Also, as the underlying distribution strays from normality, the efficiency or desirability (explanatory power) of the correlation statistic decreases rather quickly.

Be cautious when using and calculating correlations. To paraphrase Inigo Montoya, correlation may not mean what you think it means.

Lies, Damn Lies, and Statistics

Thus, we have explored our understanding of the "big three" statistics, especially with regard to some of their quirks, limitations, and shortcomings. Statistics are like power tools, very helpful and effective when used properly, and potentially dangerous and destructive when used improperly. It makes sense to handle them with care since we have learned the following.

- Statistics can be calculated from data that may be inappropriate, misleading, or inaccurate.
- Statistics may not be useful predictors as future results may not come from the same sample as the historical track record.
- Statistics are easy to misinterpret or misconstrue.

- ➢ Statistics or analyses derived from other statistics will suffer both from the flaws of the source as well as from their own limitations.

- ➢ Statistics are sensitive to a number of underlying assumptions regarding their use, especially with regard to the shape of the distribution.

So, other than being aware of these limitations and taking care not to rely too heavily on these statistics, what should the intrepid hedge fund investor do? There are a variety of other measures to consider, though there are no easy or complete solutions.

- ☑ One can employ more complex statistics that incorporate higher statistical moments like skew and kurtosis. Some examples include Kappa, Modified VaR, and Omega. These can be powerful analytical tools, but they have their own limitations which can offset their utility.

- ☑ One can employ a mosaic of statistics to overcome the limitations or biases embedded in any individual statistic. The danger is that different statistics may not be independent of one another, but effectively may tell the same story in a different way.

- ☑ One can employ common sense or qualitative overlays to adjust for known or perceived limitations in the data or statistical method. This, however, can add a layer of subjectivity that offsets the rigor of statistical analysis. The overlay also may make matters worse, not better.

- ☑ One can apply other quantitative or statistical methods to adjust for non-normality issue limitations. For example, one can transform data or inputs, or even attempt to normalize distributions. These efforts can range from the relatively simple to the relatively complex. Proceed with caution.

There is no one size fits all solution. In the end, the use of statistical analyses in hedge fund evaluation and portfolio construction is a process that must be managed carefully. Ultimately, the best defense is a well-constructed, comprehensive program implemented by experienced and knowledgeable practitioners.

Additional Information

Polymetis LLC is an investment boutique founded in 2009 by Thomas Kuntz, CFA to provide premier due diligence, research, and investment advisory services in the alternative marketplace.

For more information please contact:

Thomas Kuntz
Polymetis LLC
5 Lyons Mall, #153
Basking Ridge, NJ 07920
908-350-7087
tk@polymetisllc.com