ARBA
FINANCE COMPANY LTD.

Hedge Fund Report

Jonathan BENCHIMOL

Summary

The dangers of Mean-Variance analysis for Hedge Funds: the Skewness Effect

Mean-Variance analysis provides a powerful framework for asset allocation and portfolio performance assessment. Most investors, consultants and even a few academics feel comfortable applying Mean-Variance analysis, and therefore Sharpe ratio, to assess the performance of any kind of asset. For instance, "Fung and Hsieh" claim that using a Mean-Variance criterion to rank Hedge Funds will produce rankings that are nearly correct. But a deep analysis of this process involve to believe that certain caveats must be respected in applying Mean-Variance analysis, as with any powerful analytical technique, otherwise erroneous conclusions may ensue.

The concept of variance (and therefore volatility and Sharpe ratios) is based on the idea of a returns distribution that is symmetrical around its mean return. Most of the time, investors even have in mind a normal distribution. The bell-shaped curve does indeed have the advantages of offering analytical tractability and well-known statistical properties. However, both theory and empirical research suggest that the normal distribution assumption is not applicable to Hedge Funds for the following reasons:

- Several Hedge Fund strategies are based on financial derivatives, whose returns themselves are not normally distributed but exhibit Skewness.
- Hedge Funds tend to use leverage to magnify returns. This results in frequent price jumps and returns distributions with a higher peak at the mean, a thinner midrange and fatter tails than a normal distribution (i.e. a large number of outliers); statisticians call this "leptokurtosis".

"For investors with non-quadratic preferences, mean-variance tools are inappropriate for portfolio construction and risk management, because some of the style factors involving dynamic trading strategies exhibit highly non normal distributions. Furthermore they may have nonlinear correlation with those of the nine buy-and-hold styles. Portfolio construction and risk management must take into account investor preferences and the joint distribution of the 12 investment styles." (Performance Attribution and Style Analysis: From Mutual Funds to Hedge Funds by William Fung and David A. Hsieh)

An essential question at this stage is therefore the applicability of the Mean-Variance framework to Hedge Funds, or more generally, to assets whose returns move away from the safe haven of a normal distribution. In particular, Mean-Variance analysis is known to produce difficulties with investments that exhibit highly asymmetric outcomes. The reason is straightforward. By using variance (or volatility) around the mean as a risk measure, one is assuming that investors are indifferent to the direction of risk and have no Skewness preference. Upward deviations from the mean are handled exactly the same way as negative deviations. Both are squared in the variance calculation process, so it is impossible to distinguish between positive and negative deviations having the same absolute value. But, investors most certainly do have a preference for upside risk and aversion to downside risk, a preference for positive Skewness. Mean-Variance analysis simply ignores this.

These could be regarded as simple statistical details. However, in practice, ignoring these details could easily lead to incorrect conclusions. To illustrate the potential dangers of using the Mean-Variance framework with option-based portfolios, consider the case of three very simple but popular strategies: investing in an index fund (long the index), selling a covered call (long the index, short a call on the index), and buying a protective put (long the index, long an index put). Can just adding derivatives systematically generate superior performance and add value?

If the market is efficient (stocks are always correctly priced since everything publicly known about the stock is reflected in its market price and analyzing stocks is a waste of time because all available information is already reflected in current prices: so complicated and abstract subject…) and the options are fairly priced, the answer is no. The returns and risk reduction derived from the sale or purchase of efficiently priced options when combined with the risk and returns of a diversified market portfolio used as collateral will be offsetting and should yield no excess risk-adjusted portfolio returns. However, in a Mean-Variance framework this is not the case, and systematically buying or selling fairly valued options does provide "alphas". Alpha measures the excess performance that is attributable to a manager's actions. Consequently, covered call selling dominates the long index, which itself dominates the protective put.

Indeed, the call writing has reduced the upside potential (undesirable) while protective put buying has reduced downside risk (desirable). Mean-Variance analysis is myopic to aspects of the returns distribution beyond mean and variance. Similarly to the Skewness bias, a kurtosis bias (or extreme event bias) can be identified in the Mean-Variance framework. Since Mean-Variance analysis ignores asymmetry, it will identify some dominance, but this simply a statistical illusion. The conclusion is that managers who have relatively high (but negative) Skewness are likely to perform better than their counter parts.

Out of this come three points that are relevant to Hedge Funds. First, Hedge Funds often use options in their portfolios and this affects their returns distribution. Hedge Funds managers might be tempted to accept negatively skewed returns in exchange for improving the mean or the variance of their portfolio. Out of the money call writing is a typical "portfolio enhancer" in that respect. If they are lucky, Hedge Fund managers implementing this strategy may survive for years, boast superior Sharpe ratios (as the premiums flow directly in the bottom line with no apparent increase in the volatility), attract several investors and cash in hefty performance fees before being seriously hit once. In fact, out of the money option writing has already been the source of several disasters in Hedge Fund industry.

As an illustration of the dangers of Skewness, we should mention Victor Niederhoffer, and his three funds (Global Systems, Friends, and Diversified). Among other things, Niederhoffer's strategy consisted in selling short out of the money put options on the S&P 500. This led to steady gains from the end of 1993 to the end of 1996. The three funds had a compound annual return of approximately 40% and Sharpe ratio close to 3. This ranked as one of the best ever performances for directional commodity trading advisers over a three-year period. However, things deteriorated in 1997. The three funds had to file for bankruptcy following the crash that occurred on October 27, because they were unable to face the exercise of the previously sold put options. This is not surprising considering that option writers are essentially taking an almost unlimited future risk in exchange for what is usually a rather small immediate gain (the premium).

The second point was only recently established by academic research, when several authors pointed out the option-like nature of hedge fund returns:

- Fung and Hsieh (1997) were the first to suggest that trend following funds had a return profile similar to a straddle (i.e. long a put and a call) on US equities and that the global macro style behave like a straddle on the US dollar.
- Mitchell and Pulvino (2000) pointed out that merger arbitrage strategies exhibit a payoff similar to writing an uncovered put option on the market index. Consequently, the two authors have created a risk arbitrage index that investors can trade. The index is now managed by AQR Capital Management.
- Agarwal and Naik (2001) observed that the returns of relative value arbitrage funds are significantly related to the returns of call options on the Lehman High Yield index, deep out-of-the-money call options on the Russell 3000 index, and at-the-money put options on the MSCI Emerging Markets index.
- Fung and Hsieh (2000) benchmarked trend-following hedge funds by using lookback straddles, which deliver the performance of a perfect foresight trend follower.
- Amin and Kat (2001) developed a new performance measurement approach that is based on combining S&P 500 options to approximate hedge fund returns distributions.

This does not necessarily mean that hedge funds are actually using options, but more generally that payoffs they offer can be replicated by passive strategies that involve options on standard asset classes.

Finally, the third relationship is linked to hedge fund mangers' remuneration contracts, which often involve a performance-related fee. This is equivalent to the investor writing and granting some free call options on the fund's asset to the manager. For instance, if the incentive fee is 20% of profits, then the investor is short one-fifth of call option. The exercise price of the option depends on the hurdle rate and high water mark provisions, if any. The expiration of the option corresponds to the end of the period used to calculate the incentive fee. As a consequence of this option-like nature of performance fees, even if the prefee returns do not exhibit option-like characteristics, the postfee returns will.

Hedge funds and options are clearly not dissociable. For that reason, I believe that caution should be the watchword in using mean variance analysis with hedge funds. The free lunch it seems to offer may just be a statistical mirage rather than a first-class meal. In particular, when considering the various levels of Skewness and kurtosis, there is a high probability that some of the strategies depart sufficiently from normality to negate the overall conclusions. Then overreliance on Sharpe ratios might skew allocations toward suboptimal long-term strategies.

Correlation solely captures linear covariations

Correlation analysis is the most ubiquitous concept in portfolio optimization and asset allocation, but it is probably also one of the most misunderstood. What does correlation really measure? Statistically, correlation is a measure of the "linear" dependence between two random variables. In case of perfect linear dependence (which does not necessarily imply perfect equality), correlation should be equal to one. In the case of totally independent variables, correlation should be equal to zero. In the case of perfect linear dependence opposite in sign, correlation should be equal to minus one.

As soon as the relationship between two variables is not linear, correlation will provide inaccurate results.

Hedge fund indices are not hedge funds

The last criticism we will address in this report relates to the use of hedge fund indices rather than true hedge funds at the asset allocation level. By construction, hedge fund indices are summary statistics for a large universe. They tend to aggregate the performance of several hedge fund managers, and therefore internally diversify risks. They are also subject to well-know return biases (e.g. survivorship bias, selection bias, stale pricing). Consequently, hedge fund indices exhibit considerably lower risks and higher returns than individual hedge funds. This is particularly problematic when considering equally weighted indices based on a large number of funds. These are useful indicators of the industry's average behavior, but are not investable in practice. First, rebalancing is impossible due to lockup periods and fees; second, the minimum investment requirements preclude any individual from holding a portfolio of several hundred hedge funds.

A simple definition of Skewness and Kurtosis

Skewness measures the distribution asymmetry. A risk-averse investor does not like negative Skewness.

Kurtosis measures the fat-tail degree of a distribution. A risk-averse investor prefers a distribution with low kurtosis (i.e. returns not far away from the mean). A distribution to be normal should have an excess kurtosis equal to 0. In case of positive Skewness, it is possible to have a high excess kurtosis (higher than 5) and to have no risk of extreme negative returns in the future. The extreme returns will only be positive. This is possible if the Skewness is positive. As soon as the Skewness begins to be negative, the impact of a high excess kurtosis is significant for a risk-averse investor.

The most interesting thing is when the return distribution has a Skewness lower than -1 and an excess kurtosis higher than 1. In this case, the probability to have sudden high negative returns increases. For a distribution with a Skewness of -1 and an excess kurtosis of 5 (technology stocks, media stocks, telecom stocks or hedge funds have this kind of excess kurtosis level), a mean-variance approach will conclude that the investor will not lose more than -3.5% in the next 1 day with 99% probability. An approach, which accounts for Skewness and kurtosis, gives -7.4% losses in the next 1 day with 99% probability, which is exactly what one observes on the equity market. The difference is huge: more than 100% of risk underestimation with the use of the mean-variance.

In conclusion, for optimization and simulation, the approach must account for volatility, Skewness, and kurtosis as soon as they are significant and as soon as the investor is risk averse.

A simple definition of Covariance

A statistical measure of correlation of the fluctuations of two different quantities. In finance, covariance is applied to the annual rates of return of different investments, to measure the correlation of their year-to-year fluctuations in performance.

The definition is

$$Cov(r_1, r_2) = 1/n * \sum (r_{1\,i} - r_{1\,ave}) * (r_{2\,i} - r_{2\,ave})$$

where the terms $r_{1\,i}$ and $r_{2\,i}$ are actual values of the annual rates of return of two investments, taken over several years, n is the total number of values of $r_{1\,i}$ and $r_{2\,i}$ used, and $r_{1\,ave}$ and $r_{2\,ave}$ are the average values of $r_{1\,i}$ and $r_{2\,i}$.

An important property of the efficient frontier is that it's *curved*, not straight. This is actually significant (in fact, it's the key to how diversification lets you improve your reward-to-risk ratio. To see why, imagine a 50/50 allocation between just two securities. Assuming that the year-to-year performance of these two securities is not perfectly in sync) that is, assuming that the great years and the lousy years for Security 1 don't correspond perfectly to the great years and lousy years for Security 2, but that their cycles are at least a little off -- then the standard deviation of the 50/50 allocation will be *less* than the average of the standard deviations of the two securities separately. Graphically, this stretches the possible allocations *to the left* of the straight line joining the two securities.

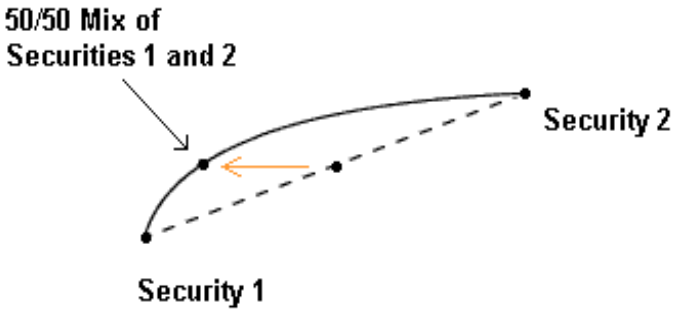

In statistical terms, this effect is due to lack of covariance. The smaller the covariance between the two securities (the more out of sync they are) the smaller the standard deviation of a portfolio that combines them. The ultimate would be to find two securities with *negative* covariance (very out of sync: the best years of one happen during the worst years of the other, and vice versa).

Conclusions

There have been several ad hoc attempts to solve the above biases of mean-variance optimization and obtain a more robust framework for asset allocation and performance analysis. A first attitude consists in measuring explicitly the bias of traditional measures. If a meter reading is systematically 20% high, the meter can still be used by applying the necessary correction. This was the approach pioneered by Leland (1999), who suggested a Skewness adjustment, and Lhabitant (1998), who explicitly computed the Sharpe ratio bias. However, relying on biases tools is not very satisfactory.

Several practitioners prefer replacing variance by another risk measure in the optimization process. In particular, professional investment consultants recommend using the Sortino ratio, which is defined as the incremental return over a minimum acceptable return (MAR), divided by the downside deviation below the MAR; deviations are only based on return occurrences with value below the MAR. In my opinion, a better alternative is to use a drawdown i.e. a statistic that captures the deepest loss a portfolio has experienced over some historical period. Drawdown measures considerably change the perception of risk, as well as the usefulness of correlation. In the return/volatility approach of Markowitz, diversification results from the linear correlation between traditional portfolios and alternatives. In the return/drawdown approach, diversification comes from the low correlation during the downside movements.

More recently, several academics have suggested using value at risk as a risk measure. Since value at risk is the maximum loss over a given interval in “normal” market conditions, optimizing according to expected return and value at risk is equivalent to increasing the return while limiting the risk of a major loss, which is in line with investors’ preoccupations.

Nevertheless, most of these studies tend to confirm that there is a compelling argument for the inclusion of hedge funds in traditional investment portfolios. Though highly stylized, the analysis of efficient frontiers in a risk/return environment (not necessarily a volatility/return environment) provides a good foundation for assessing the corresponding value added. However, the benefits may not be as great as expected. For instance, Favre and Galleano (2000) observe that the diversification effect claimed by hedge fund managers is not true for the extreme negative market returns; and investing in a diversified portfolio of 53 large hedge funds increases the value at risk (which accounts for Skewness and kurtosis) by 13% to 27%. It is therefore rather illusory to quest after a systematic target optimal proportion of hedge funds in asset allocation. The benefits will vary greatly, depending on the nature of the traditional portfolio already held; the type of hedge fund strategy considered; and the risk measure.

Consequently, the pros and cons of adding hedge funds to traditional assets should be assessed on a case-by-case basis and using the necessary information and tools. The talisman of most investors and advisers is still too often a high Sharpe ratio (high returns per unit of risk taken). This is a shortsighted attitude; in addition to the traditional risks inherent in stocks and bonds which should not be overlooked. Clearly, the difficulties associated with estimating their risks and returns make edge funds poor candidates for strategic asset allocation optimization, but not necessarily poor candidates for asset allocation.

Last but not least, investors should remember that with hedge funds, they are investing in a strategy and not an asset class, which means that their investments depend on the ability of the manager and the constancy of the investment strategy.