Aaron Brown

Hedge Fund Risk Measurement

A depressing amount of material on hedge fund risk is as likely to teach good financial thinking as a red-car-on-curvy-road car ad is to teach automotive engineering

Investment risk measurement is a confused field. Risk, like nutrition, is an important but intangible component of heavily-marketed products. Sensible work is drowned out by shills. The confusion begins by not considering why you want to measure risk. Someone trained in Modern Portfolio Theory is apt to begin with Jensen's α.[1] Suppose you have a portfolio with expected return μ_p and standard deviation σ_p. You are considering adding a new asset with expected return μ_h and standard deviation σ_h and correlation ρ to your existing portfolio.

The standard deviation formula and a little algebra shows that if you add X dollars of new asset and sell $\rho\sigma_h/\sigma_p\, X$ dollars of your old portfolio, the new standard deviation will be

$$\sqrt{\sigma_p^2 - \sigma_h^2(1-\rho^2)}$$

which is always less than or equal to σ_h (unless you make X so big that you run out of old portfolio to sell). In order to restore the same total investment, without changing the risk, you have to put $(\rho\sigma_h/\sigma_p - 1)X$ dollars[2] into the risk-free asset (if the amount is negative, you borrow instead of invest). This will increase your expected return if

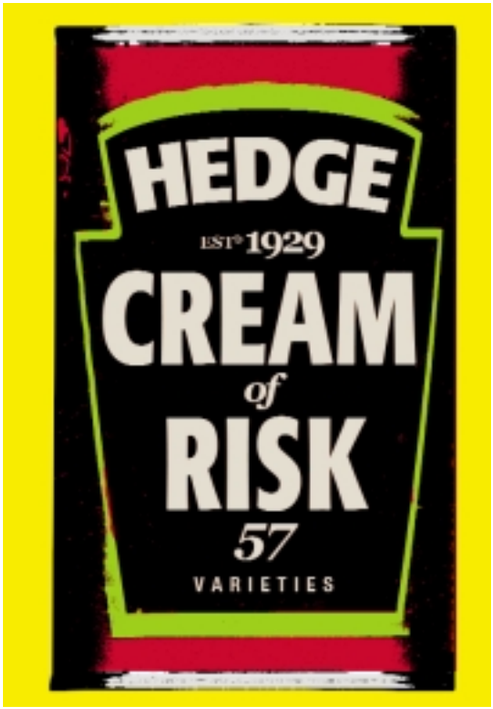

$$\alpha = r_h - \frac{\rho\sigma_h}{\sigma_p} r_p + \left(\frac{\rho\sigma_h}{\sigma_p} - 1\right) r_0 > 0$$

where r_h is the expected return on the new asset, r_p is the expected return on your old portfolio and r_0 is the risk-free borrowing and lending rate. this expression is known as Jensen's α.

With a lot of assumptions, this leads to the Capital Asset Pricing Model, with the market portfolio for p. But I don't need any of those assumptions to want investments that increase my expected return while decreasing my risk. The formula is a pure mathematical result. Some people don't like measuring risk by standard deviation, but in the context of a large diversified portfolio adding a small amount of a hedge fund, higher moments of the hedge fund's return distribution are insignificant.

The Trouble with alpha

Unfortunately, we cannot measure expected returns with any precision, except for the lowest risk investments. So α cannot be used to select investments. But it is still useful conceptually. For some investments, like liquid common stocks and many public mutual funds, we can measure $\rho\sigma_h$. Our long-term portfolio goals will tell us the ratio to use to convert that to expected return, 20 per cent is a reasonable figure for many investors. In that case, if I find a hedge fund with $\rho\sigma_h$ equal to 5 per cent, I will only invest if I expect the return on the fund to exceed the risk-free rate by 1 per cent. I still have to resort to non-statistical means to guess the expected return of the fund, but at least I have a risk-adjusted benchmark that will allow me to compare all investments on an equal basis.

This is the most common sales pitch hedge funds use, at least to institutional investors. In this view, the reason to measure risk is that it contributes to overall portfolio risk, which investors dislike.

The trouble with this approach, even conceptually, is there is no evidence hedge fund returns have statistical parameters. If each month's return is drawn from a different distribution, you cannot define the statistics. Of course, this is somewhat true for all financial assets. We know, for example, that high average return in the past does not mean high expected return in the future for stocks and mutual funds. We know that volatility moves up and down. Still, there is a strong predictable component to conventional financial assets. Some stocks and funds have demonstrably lower volatility than others and some asset correlations between large portfolios remain reasonably stable.

With hedge funds, it's not clear that predicting next month's return by looking at past returns of a fund is more accurate than looking at past returns of all hedge funds as a group.

Aaron Brown is a Vice President of Risk Architecture at Citigroup. This article expresses his personal views and does not necessarily represent the views of Citigroup or any other entity.

Most funds use multiple strategies, and change them frequently. Funds-of-funds, as a group, seem to pick worse than average hedge funds, suggesting there is little predictability in returns. There is significant doubt about whether even the reported numbers are accurate for many funds, and we know they are not representative either of hedge funds as a group or of performance in between reporting dates. There is some evidence that you can predict the volatility of a hedge fund better by knowing where it stands relative to its high watermark, than by knowing which fund, or which type of fund, it is.

This situation seems strange only in finance. It does not occur to us to assign statistical properties to most things we buy. We know there is uncertainty when we order a meal in a restaurant, buy a car or lease an apartment. But we don't model that uncertainty. We make our judgment about what is the best decision.

Uncertainty still enters the picture. We tend to avoid risky purchase decisions, not because it contributes to our aggregate risk, but because it's easier to make a mistake.

Moving from restaurants to hedge funds, if we took this approach we would pick funds based on manager credentials, appeal of the strategy, soundness of the infrastructure, reputation of the associated parties and so forth. We could still use statistics, but we would look for indicators of successful hedge funds rather than assuming future performance would be drawn from the same statistical distribution as past performance. For example, we might notice that funds in which the manager has at least a third of her personal net worth invested do better than other funds.

The Risk of Error

Now the meaning of risk has changed. We're not projecting future fund performance, so we can't talk about the risk a fund adds to our portfolio. We're trying to pick good funds, ones that will beat our target return in the future. We don't care about the path the fund takes to get above the target. Instead, risk is the probability that we will pick a bad fund instead of a good one.

The most popular general measure of this kind of risk is maximum drawdown, the largest loss you could have suffered by buying and selling the fund on the worst possible past dates. This is also the payout on a lookback straddle option[3] which provides a much more reliable benchmark for most hedge fund returns than linear market factors.

We know there is uncertainty when we order a meal in a restaurant, buy a car or lease an apartment. But we don't model that uncertainty

There are two major complaints about maximum drawdown as a risk measure. First is that it penalizes a fund for doing well. If two funds have the same worst net asset value, the one with the highest earlier net asset value will look riskier. However, this is precisely the point. We want to avoid the hedge funds that look good at some times and bad at others. The ones that always look bad will never be mistakes, we won't pick them.

The other objection to maximum drawdown is it ignores the period over which the drawdown took place. If a fund loses 25 per cent in one month it gets the same value as a fund that loses 25 per cent over three years with ups and downs along the way. The statistical parameters that would generate these kinds of behaviors in random walks are quite different. But we're not concerned about the statistical parameters because we don't believe we're looking at random walks. We're worried about how erratic the fund's performance is, how likely we are to pick it even if it's bad, and how much we can lose in that case.

This approach to hedge fund risk management is the best-developed professionally. The people who extend credit to hedge funds, primarily trading organizations, worry a lot about good and bad funds, and are in a position to tell the difference. Their concern is more with blow-up, leading to counterparty default,[4] than poor long-term returns. But that is not as far from alignment with investors as it may appear. Blow-up risk is a major component to investor risk in a hedge fund, and often the hardest component to estimate. Moreover, the days in which trading desks cared only about P&L are gone. A hedge fund that disappoints its investors is not likely to be a good long-term customer, and poses reputational risks as well as credit risks. No good bank wants to be associated with a bad hedge fund.

This approach has always seemed reasonable to me. There are a number of reasons to start a hedge fund, not all of them in investor interest. Some managers are attracted by the free call option provided by the incentive fee. Others resort to hedge funds because they cannot persuade professionals to back their trading. Still others are deluded about their abilities, or use the hedge fund as a retirement hobby. Getting into one of these funds is the real risk. If you can instead find a talented, disciplined manager who takes other people's money seriously, you can put up with a lot of standard deviation.

This brings up an important difference between hedge funds and other investment. You cannot necessarily get into any fund you want. Most financial institutions will take anyone's money (subject to money laundering rules). The first thing everyone discovers when analyzing hedge funds is the ones with the most desirable statistical characteristics are either closed, or small and turn down more investors than they accept. It's like the old IPO problem: IPO's on average outperformed the market, but if you subscribed to all of them you underperformed the market. You ended up disproportionately with the bad ones; the good ones were oversubscribed. Moreover, the sort of questions prudent people like me ask make you unpopular with managers. One hedge fund manager told me, "I quit my job [at Goldman Sachs] and started a hedge fund to get away from people like you."[5] In my experience advising institutional investors, they are

much more afraid of being shut out of a hot fund than buying into a bad one.

Big trading banks have the credit and market sophistication needed to evaluate a hedge fund's risks, and can demand the necessary access. Even the proudest hedge fund managers know they have to answer questions from trading counterparties. The risk and financial controls of the fund can be compared to the bank's own controls over its proprietary trading.

Reverend Thomas Bayes to the Rescue

Investors do not have this kind of access, which puts them in a dilemma. Past performance information of hedge funds is not reliable for estimating statistical parameters of future returns, if indeed such parameters even exist. Investors cannot get the inside information that might help tell good funds from bad. Fortunately there is a middle way, Bayesian selection.

Bayesians distinguish between prior and posterior risk. Funds have statistical parameters, but they may not be constant and we don't know them. We study past returns not to predict future returns directly, but to gather information about the statistical parameters. We pick funds that improve the expected return and reduce the risk of our portfolio considering both the risk that we have misestimated the parameters (picked a bad fund) and the risk that the random walk will go against us (even a good fund might produce a bad outcome).

To begin, a Bayesian has to define carefully the population of funds to pick from. He could set strict criteria, such as 10 years history, at least $1 billion under management and principals with at least 20 years of trading experience. This will give him a small group of funds. If he loosens the criteria, he will get more funds to pick from, but also more uncertainty about the parameters.

Once the population is defined, the Bayesian will have some subjective belief about the probability distributions of the funds. The performance information is then used to update the subjective prior distribution to form a posterior distribution. This distribution, which combines a random walk with uncertainty about the parameters of the random walk, is used for decision-making.

Without going through a full Bayesian analysis, it's possible to consider some popular hedge fund risk measures. Suppose that we are looking for hedge funds with 1 per cent expected return per month, 3 per cent monthly standard deviation and zero correlation with the stock market. We have 36 monthly[6] net asset values after fees for each fund.

There are six major worries to a hedge fund investor.

The fund has a low expected return.

The returns have a high volatility.

The returns have higher than promised correlation with the conventional asset classes.

Negative skew, smoothed returns and martingale strategies all make blow-ups more likely

The returns have a large negative skew. The manager might be writing a lot of out of the money options, or following a dynamic strategy with a similar risk profile.

The returns are smoothed. This can happen from dishonest bad marks or honestly in illiquid markets.

The manager is following a martingale strategy of increasing risk after losses and decreasing it after gains.

The first three worries are straightforward parameter misestimations. We want one set of parameters and we get another, worse, set. The last three are all connected with blow-up risk. Bad returns in a fund can set off a death spiral. The high leverage means positions must be trimmed if equity declines. That can lead to illiquid positions being sold at a disadvantageous time in the market, when counterparties know the fund has to sell. Investors tend to pull out[7] which forces more selling. The fund manager often does not want new money, even if it is available, when she is below her high watermark. In fact, the manager may be better off letting the fund die and starting a new one with a high watermark reset. All of these things lead to further worse returns, which keeps the spiral going. Negative skew, smoothed returns and martingale strategies all make blow-ups more likely.

Scorecard

Suppose the universe of hedge funds consists of 100 good funds (1 per cent expected return, 3 per cent standard deviation and zero correlation) and 100 funds that are identical except their expected return is only 0.5 per cent. If I throw out every fund with less than 1.64 per cent average monthly return over 36 months, I will be left with 10 of the good funds and only 1 of the low return funds.

Unfortunately, if I also have 100 of each of the other five types of bad funds[8], I will be left with 25 smoothed funds, 13 high volatility funds, 10 funds correlated with the stock market and 8 each of funds with negative skew and run as martingales. I started with 1 fund in 7 good and ended with 1 in 7.5 good. I have more chance picking a good fund at random from the ones I threw out than the ones I am left with.

It is possible to test directly for the other types of bad hedge funds. Table I shows the results. The first column shows the results described above if I eliminate funds based on average return. The second column shows the result of throwing out funds with high standard deviation of return. This test works well against volatile funds, but terribly against funds with negative skew, and also badly against smoothed and martingale funds.

Throwing out funds that are correlated with the stock market works a little better. It gets rid of most of the funds correlated with the market, without retaining disproportionate numbers of other bad funds. The Shipiro-Wilk[9] test for Normality does an excellent job against martingale funds, and also gets rid of some of the nega-

Table I
Number of funds, by type, left in the sample if I throw out funds based on the risk measure listed at the top of the column.

Type of Hedge Fund	Average Return	Standard Deviation	Correlation Coefficient	Shapiro-Wilk	Autocorrelation Coefficient
Good	10	10	10	10	10
Low Return	1	10	10	10	10
High Volatility	13	2	10	10	8
Correlated	10	10	2	10	10
Negative Skew	8	58	10	5	14
Smoothed	25	32	11	6	2
Martingale	8	18	10	2	22

Table II

Type of Hedge Fund	Alpha	Sharpe Ratio	Maximum Drawdown	Average Drawdown	Worst Return (VaR)[10]	Maximum Regret
Good	10	10	10	10	10	10
Less Return	1	1	3	2	4	6
Extra variance	13	8	5	5	4	9
Correlated	8	10	10	10	10	10
Negative Skew	8	27	39	36	39	12
Smoothed	24	27	9	12	29	8
Martingale	8	16	19	19	19	10

tive skew and smoothed funds. Throwing out funds with high autocorrelations gets rid of most of the smoothed funds, but favors martingales and negative skew funds.

One solution is to combine two or more of these measures to try to eliminate more bad funds. In my experience, this does not work well. Instead, we can look at some general-purpose measures that are effective against more than one kind of bad fund. Table II shows some popular choices.

Almost all the measures are effective against low return funds, and only alpha fails to spot high volatility funds. None do much either way against correlated funds, while negative skew funds, smoothed and martingale funds tend to do better than good funds on most measures.

If I had to pick a measure to put my faith in, I would use maximum drawdown, but also set a minimum level for the maximum wealth the drawdown is computed from. That is, I want a fund that attained a high maximum net asset value per share relative to the initial value, and didn't fall too far from that maximum. That is effective against every kind of bad fund except correlated funds, and correlation can be tested separately. In principle, it could collect funds that once took a lot of risk and are now resting on their laurels. However, in my experience, that is not common in the hedge fund industry. Managers who want to rest liquidate or sell their funds. Managers with funds near their high watermark are anxious to collect some incentive fees and you maximize the value of the call option by maximizing volatility.

Unfortunately, everyone wants the funds that have great past returns and are now near their peaks. These are the hot funds. You have trouble getting in, and you will pay the highest fees.[11] It can be more productive to consider hedge fund risk measures that find overlooked funds.

There's no good answer to measuring the risk of hedge funds from information available to most investors, including institutions. The data are not available and no amount of clever analysis can make up for that. However, if you start with a clear theory and apply rigorous standards, you can do better than random; and random is better than most people will do following their instincts; and following instincts is better than blindly applying risk measures developed for the stock and public mutual fund markets. You may not do well, but a lot of people will do worse.

REFERENCES

1 Jensen, M. C., 1968. "The Performance of Mutual Funds in the Period 1946-1964." *Journal of Finance*, 34, 389-416.

2 $\rho\sigma_h/\sigma_p$ is commonly denoted β.

3 Goldman, M., H. Sosin and M. Gatto, 1979, "Path Dependent Options: 'Buy at the Low, Sell at the High,'" *Journal of Finance*, 34, 1111-1127. Also see Fung, W. and Hsieh, D., 2001, "The Risk in Hedge Fund Strategies," *The Review of Financial Studies*, 14, 313-341.

4 Hedge funds are of more concern than other customers because they typically use more leverage, pursue riskier strategies, have no access to additional capital and cannot be hedged in the credit derivative market. Smaller funds may rely on smaller and less professional middle and back office staffs than other trading customers.

5 I assume he did not mean I reminded him of the masters-of-the-universe white-shoe investment bankers at Goldman, but the annoying nerds in risk management.

6 Few hedge funds give more than monthly data, and some give only quarterly or annual. Many do not report at all, except to investors.

7 Funds typically have lock-up provisions to reduce this risk, but they do not eliminate it.

8 I made each bad fund 1 standard deviation worse than the good one. For example, the good fund has a mean return of 1% per month and standard deviation of 3% per month. Over 36 months, the standard deviation of the average return is 0.5%, so I gave the low return fund 1% - 0.5% = 0.5% expected return. By setting the cut-off 1.28 standard deviations above the mean (1% + 1.28*0.5% = 1.64%) I eliminate 90% of the good funds and 99% of the bad.

9 Shapiro, S. S. and Wilk, M. B., 1965. "An Analysis of Variance Test for Normality (Complete Samples)," *Biometrika*, 52, 3 and 4, pages 591-611. Unfortunately, this test works so well because the good fund data all came from a simulated Normal distribution. In practice, good funds generally have harmless types of non-Normality.

10 It's not possible to estimate VaR reliably from 36 monthly returns, but if we had to we would make some parametric adjustment from the worst monthly return among the 36.

11 An important point that is often neglected is to adjust past fund returns for the current fee structure, or in some cases, the fee structure that applies to you. Many managers raise their fees after success, and if you restate past performance with the higher fees, it was not successful. Also some managers let different investors in on different terms, the reported returns may not be the ones you would have received.