The biggest "CIO[1]" Risk:
"Alpha" Risk misManagement

by
Leslie Rahl, Richard Horwitz, Erin Simpson

Summary

Despite "beta" risk being the dominant source of systematic behavior in the hedge fund industry, the greatest source of market-driven hedge fund blow-ups is "alpha" risk. A significant percent of alpha risk results from hedge funds' tendency to be net-long Complex, Illiquid, and Opaque ("CIO") securities. The hedge fund industry does an inadequate job of managing alpha risk. Finally, this article defines a risk-adjusted measure of alpha, naturally called the "Alpha Ratio".

What are Alpha and Beta Risk?

Let's first define alpha and beta risk. To do that we must first define alpha and beta. Beta is the return enjoyed by hedge funds driven by their exposure to the underlying markets. Over a long period of time, markets reward investors for accepting the inherent risk. For example, the equity market has, over a long period of time, returned an average 11% and long-term bonds have returned 8%, both in excess of the risk free rate that averaged 5%. However, these beta exposures can be gained through ETFs or futures, and do not justify the significant fees hedge funds charge.

Alpha is the additional return generated in excess of that which comes from market betas. It is broadly generated by:

- Security selection
- Beta timing
- Complex, Illiquid, Opaque (CIO) premium

The first two drivers are based on skill, and ultimately represent a zero sum game – if you do not have superior skill no excess returns should be generated. The last is rewarded by the market to compensate investors for taking the risk of holding more complex, less liquid, and more opaque securities. As with betas, CIO returns are not a zero sum game. CIO securities are not favored by the typical, risk-averse investor. They therefore generally trade at a discount and generate a premium return. While some hedge funds can definitely generate alpha by superior security selection and beta timing, a significant amount of hedge fund alpha is driven by exposure to CIO risk, the risks related to such securities. Therefore, hedge funds, in general, go long more complex, less liquid, and more opaque securities and hedge by shorting simpler and liquid common securities.

Although hedge funds present themselves as alpha generators, the alpha they generate comes with a significant amount of beta. The correlation of hedge fund indices and the equity market

[1] Complex, Illiquid, Opaque.

is .86. This implies that the vast majority of systematic hedge fund behavior is driven by market betas. However, the HFR Composite Index since after the fall of 1998 demonstrates a worst month of 3.5% and the largest drawdown of 6.4%. This level of risk in the index is inconsistent with the hedge fund blow-ups that repeatedly occur. Even if a fund is highly geared to this systematic behavior, the largest loss driven by this exposure would still be relatively small.

Hedge Fund Blow-ups

Let's start exploring hedge fund blow-ups by reviewing some of the biggest market-driven[2] hedge fund blow-ups:

- Mortgage back funds in February 1994 (Askins) and October 1998 (MKP and Ellington) going long MBS derivatives hedged by collateral and Treasuries

- Relative value funds in September 1998 (Long Term Capital Management, DE Shaw) going long the convergence and Russian yield curve trades

- MBS arbitrage funds in October 1998 (MKP and Ellington) going long MBS derivatives hedged by collateral and Treasuries

- Long/short sector funds when the tech bubble burst (Andor, Paramount) going net-long micro/small cap tech stocks

- Credit funds in May 2005 going long equity and short mezzanine traunches of CDOs

- Amaranth in March 2007 going long the March and short the April natural gas contracts

- Credit funds in spring 2007 (Bear Stearns, Basis, Sowood, BNP) going long "high-quality" sub-prime CDOs hedged with more liquid credit instruments

- Systematic equity funds in August 2007 (Global Alpha, Tyke) going long value and short growth

The primary cause of these hedge fund blowups was alpha risk. The common factor in all of these blowups was aberrant relative value behavior, exacerbated by leverage, rather than just directional moves in the underlying markets. Unfortunately, relative value strategies have residual basis risk and are not fully immunized from the vagaries of the financial markets. In all of the above blowups, a primary source of the alpha risk was CIO risk.

In taking exposure to CIO risks, many hedge funds routinely earn a small premium in exchange for accepting the risk of a low probability adverse event. In other words, these hedge funds are effectively selling insurance (alternatively, shorting volatility). This is the cause of their "left-tail" (leptokurtotic) behavior (their distribution of returns is a large number of relatively small gains intermixed with a very small number of relatively large losses). Since the gains happen

---

[2] Operational risk is part a source of many hedge fund problems. Valuation risk is addressed later in the article.

significantly more frequently than the losses, it can take some time (often years) for an investor to see evidence of this extreme left-tail tendency.

As the alpha of these CIO exposures are generally very small, leverage is often used to amplify them, to generate satisfactory hedge fund returns. For strategies based on CIO exposures, this leads to a potentially dangerous combination of illiquidity and leverage. This potentially explosive mixture is the cause of most of the blowups above.

"Valuation risk" is highly interrelated with alpha risk. The fundamental problem is that price discovery (valuation) of CIO securities is difficult and imprecise. Securities that are primarily driven by beta can be easily marked to model. CIO securities, primarily driven by alpha, are difficult to value as they, by definition, are complex, illiquid, and opaque. Therefore, many blowups (e.g., Lipper, Beacon Hill, Wood River, Lancer, Manhattan Fund), in which the root cause was alpha risk, have resulted in subsequent valuation issues, covering up the initial losses. Furthermore, as this difficulty in valuing CIO securities permits managers investing in them to "smooth" returns ("Return Smoothing Practices", *the hedgefund journal*, September 2007), top-down (based on the monthly return series) measures of volatility of funds with material exposure to CIO securities are often understated.

Who is responsible for guarding against alpha risk? The conclusion of Presidents Working Group was that the ultimate responsibility was "most appropriately borne by investors." *The Agreement between PWG and U.S. Agency Principals on Principles and Guidelines Regarding Private Pools of Capital* published in February 2007 stated: "Private pools of capital can be an appropriate investment vehicle for more sophisticated investors. Because these pools can involve complex, illiquid or opaque investments and investment strategies that are not fully disclosed, the risks associated with direct investment in these pools are most appropriately borne by investors with the sophistication to identify, analyze and bear these risks."

The misManagement of Alpha Risk

Having identified alpha risk as the primary cause of most hedge fund blowups, let's now address how well alpha risk is managed by hedge funds. The unfortunate answer is: not very well. The fundamental problem is that risk management techniques have been developed by the sell-side to focus on structural risk exposures or betas. The fundamental approach is statistical risk management. Statistical risk management applies historical return behavior to the current portfolio to determine what the potential returns of today's construction would have been over recent history. The Basel Accord mandates that banks and broker/dealers use Value at Risk or VaR to calculate their capital requirements to cover the market risks they are taking. VaR is excellent at permitting an institution to aggregate risk across asset classes. The vast majority of risk of the securities issued by these institutions is beta (including interest rate duration). Financial institutions, in general, do not have significant alpha risk exposures (excepting recent losses by Merrill Lynch, Citibank, and Bear Stearns resulting from sub-prime mortgage holdings). The primary objective of applying VaR is to aggregate the risk of the portfolio of outstanding positions and to permit the institution to hedge residual risk, if they so choose. Therefore this methodology is primarily focused on beta risk management.

The fundamental problem with applying the VaR methodology to hedge funds is that statistical risk management by definition is backwards looking. Were history to repeat, and therefore statistical relations to hold true, statistical risk management by itself would be adequate. However, almost every financial crisis has resulted from some market dislocation, or a discontinuity of historical relationships. Over the last couple of decades, there has been a once in a lifetime crisis every few years:

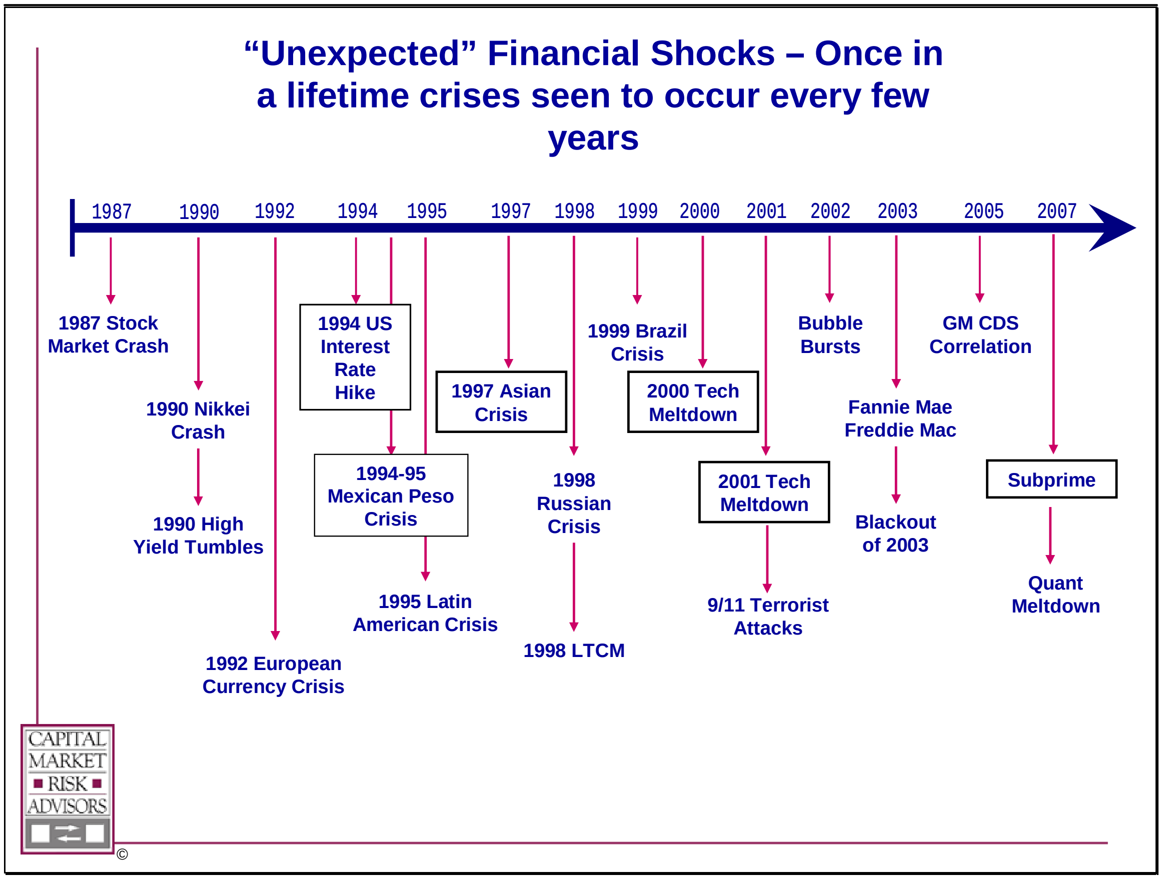

While some crisis occurs with regular frequency, each specific event is idiosyncratic and very different in nature. Shocking a portfolio based on prior extreme events is unlikely to be informative about tomorrow's discontinuity (although stress testing specific scenarios about which you are concerned is useful). Therefore, statistical risk management does not handle the alpha, or left-tail event risk to which hedge funds are exposed.

VaR has not proven to be effective for the majority of hedge funds strategies. It works very well in global macro (directional traders that are very similar to the prop desks of banks). Equity strategies (long/short, market neutral, activist) have not embraced VaR as they generally develop their own spreadsheets to identify the style/industry alpha exposures. Credit (distressed, high yield, structured, convertible arb) and event (capital structure arbitrage, merger arbitrage) strategies do not use VaR because many of their holdings are relatively illiquid and the idiosyncratic risks of the lower quality credits in which they participate are equity-like. Finally, while VaR can work for CTA's, they utilize the SPAN margin calculated by the exchange.

Given that statistical risk management does not work well for many of these strategies, how does one manage alpha risk? The simple answer is "diversification". As alpha risk is by definition idiosyncratic, it cannot be predicted. Investors are seductively drawn to CIO exposures, as large losses are generally preceded by a long series of consistent gains. Therefore, you cannot inoculate yourself from it, as it is inextricably linked to the alpha generated by hedge funds. Instead, you can diversify your exposure and keep your concentration to each specific exposure to an acceptable level. Furthermore, while one can easily hedge beta risks, alpha risks cannot be hedged.

Let's talk a little about diversification. Diversification is the primary lever in risk management in hedge funds. Risk is neither good nor bad, as long as it has a prospect of being well rewarded. However, a concentrated exposure to a risk, especially to a level to which you cannot sustain a very unfavorable outcome, is bad. This is consistent with the concept that there is no risk that you should be unwilling to accept to some level of exposure as long as the risk is attractively enough compensated.

If the way to manage alpha risk is diversification, how should diversification be managed? The key to managing the diversification of alpha risk is to understand all of the basis exposures one is taking. There are many sources of basis risk:

- Instrument (e.g., cash versus derivative)
- Asset class (e.g., equity, interest rate, commodities)
- Position type
- Liquidity group
- Geography
- Industry
- Market cap
- P/E group
- Maturity
- Credit quality
- Debt seniority and capital structure
- Structural versus synthetic versus primary

As we have already said, the market disruptions that caused the above blow-ups were historically unprecedented. One must fundamentally understand a fund's basis exposures and be comfortable that the fund could tolerate a significant discontinuity in each. Furthermore, an investor in multiple funds must be able to aggregate these exposures across his entire portfolio to understand his aggregate exposure to these alpha risks.

Finally, as many of these exposures are driven by CIO premia, and there is always a risk of a broad flight to quality as occurred in the fall on 1998, one must similarly feel comfortable that they could tolerate a general market flight from CIO securities. A good proxy for CIO risk is bid-offer spreads. In general, the more complex, illiquid, and opaque a security is the larger the bid-offer spread. A systematic, if not precise, way to test for potential flight to quality risk is to stress the bid-offer spreads. While doing so does not provide an exact measure of an individual portfolio, consistently doing so across portfolios provides a very relevant relative measure of exposure to CIO risk.

The Alpha Ratio

Having defined both alpha and alpha risk, we propose a new measure that combines them. The Alpha Ratio is a risk-adjusted measure of alpha returns. The Alpha Ratio is fundamentally the Sharpe Ratio calculated on the alpha return stream. While the Sharpe Ratio adjusts returns for volatility, it does not recognize correlations. This is incomplete as an investor is not indifferent between two funds with the same Sharpe Ratio but with very different betas. The Alpha Ratio implicitly incorporates correlations by isolating alpha. The Alpha Ratio is calculated as the ratio of the annualized alpha of a fund divided by the standard deviation of the alpha returns. The alpha is not reduced by the risk free rate (as are the returns of the Sharpe Ratio) because alpha is already calculated as the non-beta returns in excess of the risk free rate. The variance of the alpha returns is simply the product of the variance of the returns of the fund multiplied by one minus the R-square of the regression that removes the betas. Therefore, the standard deviation of the alpha returns is the standard deviation of returns multiplied by square root of one minus the R-square.

As Bill Sharpe has openly recognized issues with his own ratio because of the non-normality of hedge fund returns, how can we propose to fundamentally apply his methodology to alpha? The answer is relatively simple. While several years of monthly returns provide adequate data to fit the dual parameters of the normal distribution, they are woefully inadequate to define a return distribution given kurtotic and skewed behavior. Measures that work for non-normal distributions (e.g., Omega, CVAR) can be very misleading because of the limited amount of data. However, we believe that while less than perfect, the standard deviation of returns is still a meaningful, though imperfect, measure of risk.

This belief is based on research on the drivers of drawdown, the ultimate measure of hedge fund risk. Unfortunately, hedge fund track records are so short that even for the longest records only one or two major drawdowns have occurred. Therefore, there is inadequate data to analyze the largest drawdown with any statistical significance. However, our research has sought to identify the drivers of drawdowns to establish the best proxy for the largest drawdown, the ultimate behavior that data limitations preclude us from directly measuring. We performed a cross-sectional analysis of a universe of 700+ hedge funds with a full decade of returns (subject to all of the statistical biases to which such datasets are unfortunately exposed). For each fund we calculated the largest drawdown as the dependent variable and the following measures of return as potential independent variables:

- Standard deviation
- Downside deviation
- Serial correlation
- Beta
- Skew
- Kurtosis

Analyzing the relationship of largest drawdown of each fund to its standard deviation alone results in an R-square of .76, a very high explanatory power (better than downside deviation). Adding serial correlation increases the R-square to .78. Adding the additional three factors (including the skew and kurtosis which should explicitly recognize the non-normality of returns) does not increase the R-square. In other words, across a large universe of hedge funds with long track records, the standard deviation of returns is a superior proxy for drawdowns. Adjusting for serial correlation slightly improves the explanatory power. Incorporating higher moments of the distribution does not add value. We therefore believe that, despite its deficiency, the Sharpe Ratio, and its new-found sibling the Alpha Ratio, are meaningful measures of risk-adjusted returns, with the caveat emptor that significant anomalies can exist because of extreme non-normal behavior.