Hedge Funds: Truths and Myths

Prof. Luis A. Seco
Department of Mathematics
Rotman School of Business
University of Toronto
seco@math.utoronto.ca

Abstract:

Hedge funds have enjoyed increasing levels of popularity coupled with mysticism, opacity and obscurity, yet they are a trillion dollar industry worldwide. This paper aims at presenting some facts about the hedge fund sector, and to reflect on some of its myths. The paper reviews the industry, its products, its risks and proposes a way of thinking about them from the investor's viewpoint which takes into account their nature, very different to traditional investments.

Resumen

Los hedge funds, o fondos de cobertura, han despertado la atencion de los inversores con una aureola de misticismo y oscurantismo, y se han convertido en un sector de mas de $\$10^{12}$ a nivel mundial. En este articulo se intenta simplemente presentar algunos hechos sobre el sector, y reflexionar sobre algunos de los mitos. En este articulo revisaremos la industria, sus productos, riesgos y plantea una manera de pensar sobre los hedge funds, desde el punto de vista del invesor, que tiene en cuenta sus diferencias respecto a las inversiones tradicionales.

Executive summary

Hedge funds have become financial companions in our everyday life; their presence encompasses many of the activities that were traditional to the banking sector and are being abandoned due to increased regulation, many new activities which are a reflection of new capital markets created every year, and of course an ubiquitous search for opportunity, risk and profit in capital markets in general. They have experienced tremendous growth over the last decade as an asset base for institutional and individual investors.

The hedge fund's principal quality is their search for uncorrelated returns. Investors need uncorrelated returns for diversification. In an increasingly globalized world, economies are rapidly becoming highly dependent on each other's, and hedge funds owe their popularity largely to the fact that they present to the investors one of the few alternatives to economic-dependent, highly correlated investments worldwide.

Hedge funds obtain returns by taking risks; understanding their risks therefore becomes the key factor to understanding them as an industry. Their risk is the glass through which we must look at their returns. This presents a need to depart from traditional portfolio theories, which view risk and return as different, unrelated quantities, and the need for a new portfolio theory arises, which we address in the sections below.

We also address the main risk sensitivities; correlation, the principal benefit of hedge fund investing, will lead to the main sensitivity for us to study; if hedge funds are popular because of diversification, our main concern should be why and how that can go away; we will model it in the context of regime switching models, which provide an easy-to-understand picture of correlation breakdown phenomena.

Most of the article is devoted to understanding portfolios of hedge funds, investments that allocate assets to individual hedge funds, also called fund-of-funds. From this perspective, hedge funds become assets in which investors participate through shares in limited liability vehicles, very much like traditional equities. But we also consider hedge fund structures, which are becoming increasingly popular. Often designed to meet risk or return objectives, or to access hedge fund investments in compliance with regulations, they give rise to different risk and return profiles, which will be briefly discussed in the paper.

In an unusual burst of sanity, Don Quijote affirms that "...*para sacar una verdad en limpio, menester son muchas pruebas y repruebas",* stating that truth can only be found after profound and repeated verification. The world of hedge funds is full of giants and windmills, of risks and returns, of mathematics and experience, of qualitative and quantitative analysis, of headline news and private reports. The objective of the article is to present another point of view from which to observe an industry that has become too much of a legend.

1. Introduction to Hedge Funds

You start your breakfast with a glass of orange juice; perhaps the orange grower in Argentina who sold the oranges to the manufacturer obtained financing from the International Investment Group, a hedge fund in NY that specializes in commodity factoring. You then go to work in a car that you think you have leased from the Ford motor company; in reality, perhaps Ford Credit Inc. did nothing but resell the lease to a hedge fund who actually does all the work in terms of the management of the fixed income products that are the building blocks of a typical lease portfolio. When you arrive at work, you remember the days that Cerberus, another hedge fund, was trying to assist in a leveraged buyout for your company. After a hard days work, you get back home and relax watching a movie; its script was financed by a hedge fund, which specializes in investing in a diversified portfolio of Hollywood scripts, costly enterprises with a 1 in 20 probability of being accepted by a production studio, but that when rarely successful can offer returns in multiples of the original investments. And of course, your bank sold your home mortgage to a mortgage arbitrage manager a long time ago.

From an abstract point of view, a hedge fund is a business that turns investment risk into investment return. After the securitization revolution of the nineties, anywhere you look there is a hedge fund. Most of them, however, invest only on stock markets, taking long and short positions on individual stocks and stock indices. Hedge funds are not regulated entities, and therefore can take both long and short positions, use arbitrage, leverage, buy and sell undervalued (overvalued) securities, trade options, bonds, OTC products, and in general invest in almost any opportunity in any market where it foresees *inefficiencies* and/or substantial gains at reduced risk.

The hedge fund sector was created by the private banking and family office business in the seventies, as investors no longer accepted relative return results, and demanded absolute return strategies; hedging a stock portfolio with an index was the simplest way of doing it.

The Hedge Fund Sector is estimated to be a $1 trillion industry and growing at about 20% per year with approximately 10000 active hedge funds. From 1999 to 2004 hedge fund assets grew 46% (mutual fund assets grew 5%). The reason for their growth is varied: bank regulations have created a fertile territory for unregulated firms to operate and take the business that traditionally belonged to the banks; the creation of new capital markets (such as energy, insurance bonds, leasing, factoring, etc.) created opportunities that smaller firms can take advantage of faster and better than the larger banks. And of course, there is perhaps a share of unskilled managers who are attracted to a profitable, high-fee business: 1% of the assets as management fee is normal, and 10% of the absolute gains over a benchmark is standard.

Many hedge fund strategies have the ability to generate positive returns in both rising and falling equity and bond markets. Hence, they provide a long-term investment solution, eliminating the need to correctly time entry and exit from markets. Common wisdom is that inclusion of hedge funds in a balanced portfolio reduces overall portfolio risk and volatility and increases risk-adjusted returns.

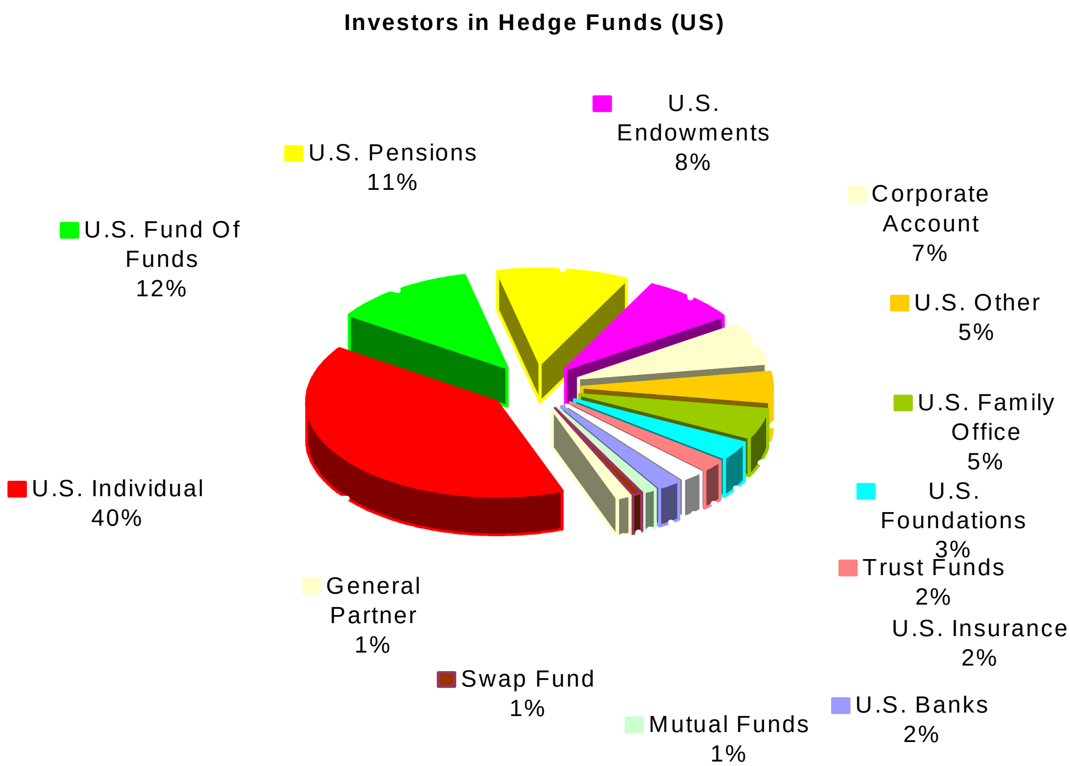

Chart 1.1

The hedge fund sector not only has improved by extending the investor group to institutions –as the previous graph shows- but also a wide range of hedging strategies are available to hedge funds: Investing in anticipation of a specific event - merger transaction, hostile takeover, spin-off, exiting of bankruptcy proceedings, etc.; Investing in deeply discounted securities - of companies about to enter or exit financial distress or bankruptcy, often below liquidation value, commodity trading and financing, fixed income arbitrage, mortgage arbitrage, risk arbitrage, convertible securities, etc.

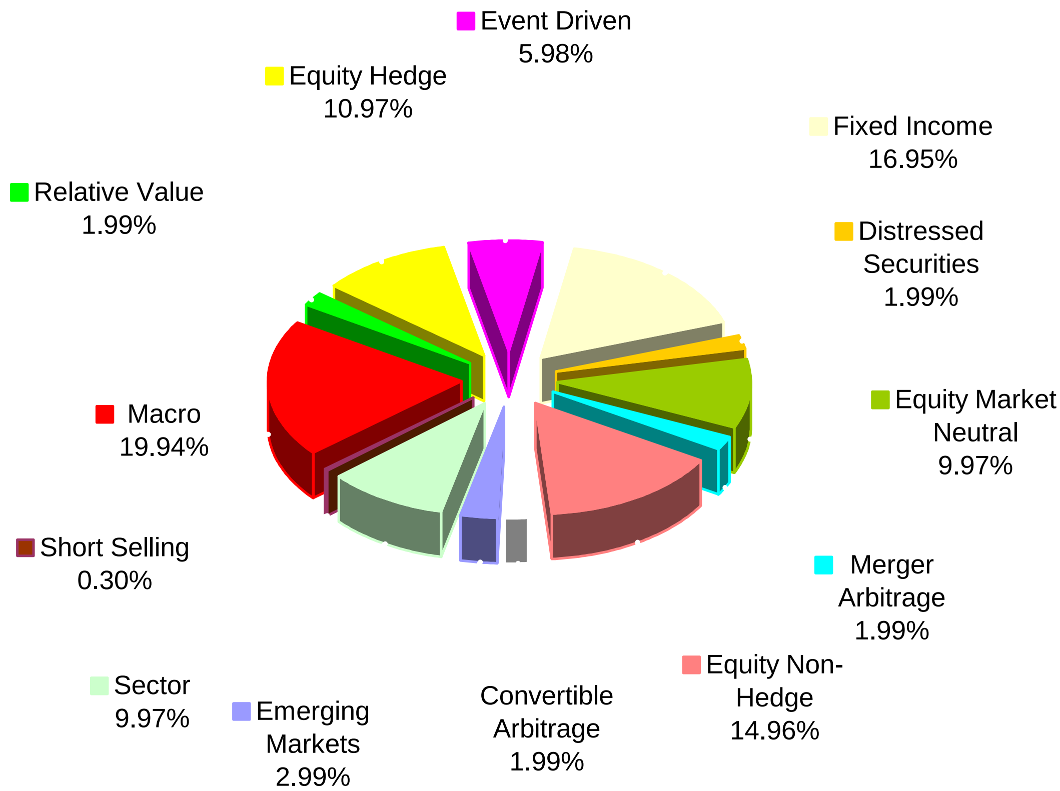

Chart 1.2

Hedge funds utilize a variety of financial instruments to reduce risk, enhance returns and minimize the correlation with equity and bond markets. An increasing number of pension funds, endowments, insurance companies, private banks and high net worth individuals and family offices allocate assets to hedge funds to minimize overall portfolio volatility and enhance returns.

The key to the diversification benefits of hedge funds is their correlation, to other markets and amongst themselves. While not an asset class, they are a return stream class in themselves as they present features not present in other investment universes.

The following two charts present two contrasting histograms; in the first, we see the distribution of correlations amongst the Dow Jones 30 stocks (435 pairs in total, and 435 correlation numbers); in the second, we see the same correlation histogram, this time for a collection of hedge funds. While we clearly see that Dow stocks have a predominant correlation of about 90%, for hedge funds it is more like 20%, with quite a few pairs with negative correlations.

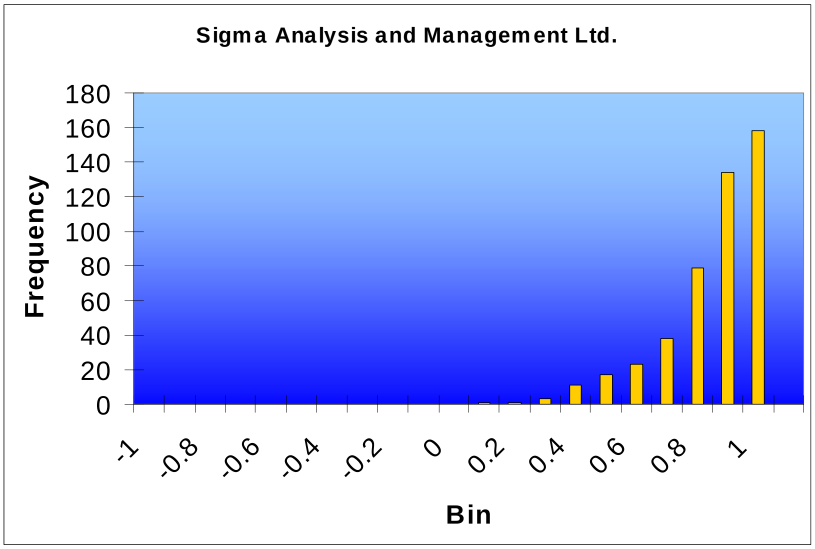

Chart 1.3

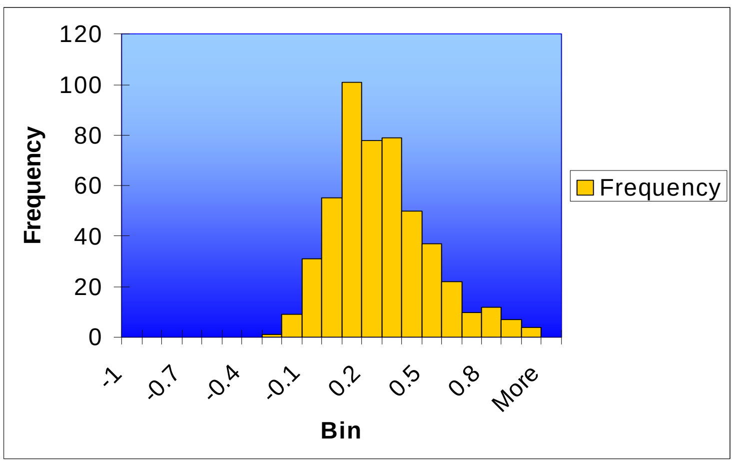

Chart 1.4

1.1 Hedge Fund Products

Today, hedge fund investments are usually done through three different, but related, products: fund-of-funds, hedge fund indices, and hedge fund structures.

Funds-of-funds are simply portfolios of hedge fund investments. Management companies set up investment funds that typically sell shares in the pooled fund, and invest the funds' assets in a portfolio of hedge funds. Very often, the investment is done after leverage: the fund-of-funds manager borrows money from a bank (a loan of 100% of the fund's assets is very common), and invests the total in a portfolio of hedge funds. In this manner, the bank earns a spread over LIBOR in the loan, and the investor earns the return on the overall investment, after bank's fees and manager's fees are paid. This type of investments owes its popularity to the low interest rate environment we have witnessed over the last decade, and the high average return obtained by hedge funds. When leveraged, fund-of-fund products give rise to very interesting issues in risk management, as we will see below.

Hedge fund indices grew very quickly in 2004. We need to make here a very important distinction, as there are two types of hedge fund indices. First, there are the *Industrial Indices*, averages of hedge fund returns (weighted in one way or another), which are used to track the performance of the industry. The best-known are the ones by Hedge Fund Research, Morgan-Stanley, S&P, CSFB and Vanhedge. They are typically published with sub-indices corresponding to each of the major trading styles. Although they can be useful to get a general idea of the evolution of the industry, because of the low correlation that exists amongst hedge funds, they are of necessity quite different between each other. The following chart outlines the correlations among them

CSFB Hedge Investible Index	1.00				
MSCI Composite Equal-Weighed Index	0.47	1.00			
MSCI Composite Asset-Weighed Index	0.51	0.98	1.00		
S&P Hedge Fund Index	0.41	0.90	0.94	1.00	
Van Global Hedge Fund Index	0.80	0.48	0.51	0.41	1.00

Chart 1.5

Following the explosion in hedge fund allocations in 2003, most of the corporations that sponsored a hedge fund *industrial* index launched a hedge fund *investable index*.

It is important to note that, unlike equity or bond indices, there is no relationship between the reference or *industrial* index, and the investable one. The following chart summarizes the differences in year-to-June 1 2005 return performance for the HFR Indices:

Index	Investable	Industrial
Convertible Arbitrage Index	-8.88%	-5.60%
Equity Hedge Index	-2.51%	-1.83%
Equity Market Neutral Index	0.94%	1.43%
Event Driven Index	-1.12%	-0.89%
Macro Index	-2.17%	-0.30%
Merger Arbitrage Index	1.92%	-0.41%
Relative Value Arbitrage Index	-2.11%	-0.13%

Year-to-date results as of June 1 2005: HFR indices

Chart 1.6

Perhaps the most striking fact in this chart is that 2005 has not started all that well for hedge funds. But beyond this, one might wonder why the differences between the investable and industrial indices. The reason is at the heart of hedge fund activity; first, investing in a hedge fund is not like purchasing stock in a publicly traded company: the fund may be closed to investments, the investment may have to be locked for months or years, and redemptions may require advance notice. Investable indices are constructed with a number of features that are equity-like, and not hedge-fund-like (such as liquidity, unlimited capacity for new investments, or even investment transparency), and hence introduce a self-selection bias, which results typically in poorer performance.

Most of the assets in the hedge fund sector, and in particular those assets allocated over the last few years, made their way into the hedge funds through investment structures. Examples of them are:

❖ CPPI (Constant Proportion Portfolio Insurance) or capital protected notes; they provide access to the upside of hedge fund returns, while protecting the original investment.

❖ Leveraged investments; they maximize the investment return by borrowing against the fund's assets.

❖ Collateralized Fund Obligations (CFO); they securitize then loan part of a leverage structure, which is then sold to bond investors.

Those structures exist to satisfy regulations (change a US-based hedge fund investment into a domestic product in Europe, for instance), to obtain higher return targets, for insurance or capital protection, or in general to transfer some of the risks of the investment from the investors to third parties.

2. Risks in Hedge Funds

Hedge funds are uncorrelated to traditional markets, and internally uncorrelated also.

The efficient frontier is a concept introduced by Markowitz; an efficient portfolio is one whose return cannot be improved without increasing the risk, or its risk cannot be lowered without hurting the return.

Investors should try to have portfolios that are efficient. In a risk/return 2-d graph, those portfolios appear all on a curve, called the efficient frontier.

The following chart shows a well-known result, which is that from the Markowitz viewpoint, investments in hedge funds improve the expected return by about 10%, and reduce the volatility by about 6%. But, as we shall soon see, it is not clear what this means.

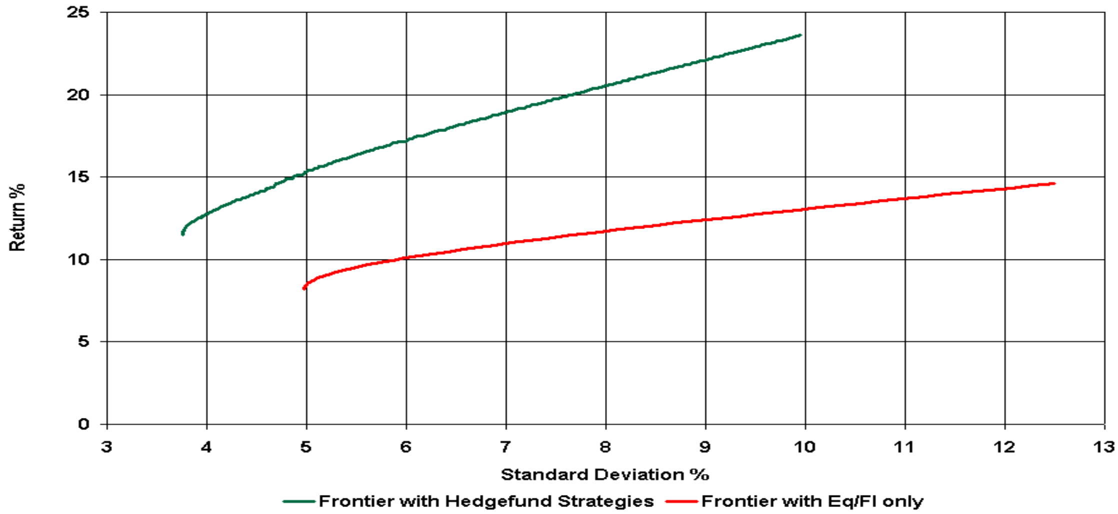

Chart 2.1

2.1 Risk in Hedge Funds

Hedge funds make *uncorrelated* returns because they take *different* risks. Analyzing their risks is not just a good idea, it is the beginning of any investment operation.

The following is a list (not an exhaustive one) of risks categories specific to hedge funds:

❖ Lack of transparency. Hedge funds are businesses, and as such, they often choose not to disclose their most precious asset: their strategy. It is typically impossible to get a hedge fund to report the positions it holds in its investment portfolio. However, they do offer a subscription document or offering memorandum, which is a legal document that binds the manager to a certain set of activities, and therefore the manager has limits to what she can do. Audited financial statements are typically available, and they should be consulted prior to any investment.

❖ Fraud. We mentioned above that the manger is bound to limit their activity to certain legitimate activities. Fraud will occur when they don’t. Fraud can also occur when they misquote performance or valuations.

❖ Counterparty risk. Although not specific to hedge funds, they are specially sensitive to this risk type because of the unregulated and specialized nature of their transactions. Counterparty risk (or credit risk) refers to losses that the fund can incur into when the counterparty to some of its financial transactions does not honor their obligations (default). This term also refers to situations when, without default or bankruptcy, the counterparty undergoes a credit downgrade, hence affecting the market value of the securities in the fund.

❖ Portfolio liquidity and redemption orders. Hedge funds often restrict fund redemptions to quarterly terms (or more), and usually with some advance notice. Liquidation orders, therefore, can take time to process. Moreover, funds can choose to suspend redemption orders when they estimate that liquidation would be detrimental to the remaining investors in the fund: a hedge fund forced to sell securities to meet redemption orders is an easy prey to its counterparties or competitors in the financial markets.

❖ Capacity risk. Hedge funds make money on fees; while management fees are considered generous by investment standards (1% or more), it is with the performance fees that most funds make their money. Hence, they are encouraged to close the fund to new investments if they see no new opportunities for return and sense that they have reached capacity. Those who don't limit their fund-raising activity, may raise assets beyond their natural capacity, which may lead to a decrease in their future returns.

❖ Style drift. Individual hedge funds are normally selected to be part of an investment portfolio for a good reason, which is usually the particular trading style they employ; when a fund changes its style and adopts another one, it may create imbalances inside the portfolio; if sufficiently many funds undergo style drift over a period of time, the risk pattern of the portfolio can be changed drastically and give rise to unintended risk concentration.

❖ Data. Imagine a hedge fund that trades between New York and Tokyo. When it calculates the daily value of the assets, does it use NY closing time of 4 pm EST? or Tokyo closing time? Or does it use NY close for some positions and Tokyo for others? Questions as simple as this (and of course much more complex) can introduce huge differences in valuation of the firm's assets and hence of the price settlement when new investors join the fund or the price paid to investors exiting the fund. It can also lead to a smoothing effect in the fund's performance numbers.

❖ Legal risk; in 2003, after some illegal activity involving mutual funds and some hedge funds, market timing activities (one small but profitable hedge fund style) became under general legal scrutiny. Investors in certain funds rushed for redemptions, driving the value of the assets remaining in the fund down dramatically. Changes in law affect all activities in life and in particular in the investment sector, but when they mix with highly complex, illiquid, investments such as the ones inside a hedge fund, the result can be dramatic. Tax laws can be particularly sensitive for certain hedge fund activity.

And of course, hedge funds are exposed to investment risks in general:

❖ Market Risk. The risk in reducing the value of the portfolio's positions due to changes in markets.

❖ Credit Risk. The risk in reducing the value of the portfolio's assets due to changes in the credit quality of the counterparties.

❖ Liquidity Risk. The risk of losses because of travel-time delays of assets.

❖ Common factor risk: industry specific, geographical risk, etc.

❖ Operational Risk. Internal systems, people, physical events.

❖ Corporate event risk: earnings revisions, mergers, etc.

❖ Model risk.

❖ Legal and Regulatory Risk.

All these risks must be analyzed and conclusions about the investability of a hedge fund must be made prior to the investment, and periodically during the life of the investment. This is what we refer to as the due diligence process, and is key for fund of fund managers.

J. R Giraud, in the Operational Risk Management manual produced by EDHEC (see [Giraud]) proposed a polygonal visualization method for the determination of risks in hedge funds according to several categories. This method allows for a summarized, multidimensional risk view in a particular hedge fund; here we propose the septagonal version, which would look as follows:

BADXXX

Category	BADXXX Risk Rating
Operations	4
Portfolio Valuation	2.5
Concentration	4
Instruments	2.5
Liquidity	4
Management Experience	4
Conflicts	4

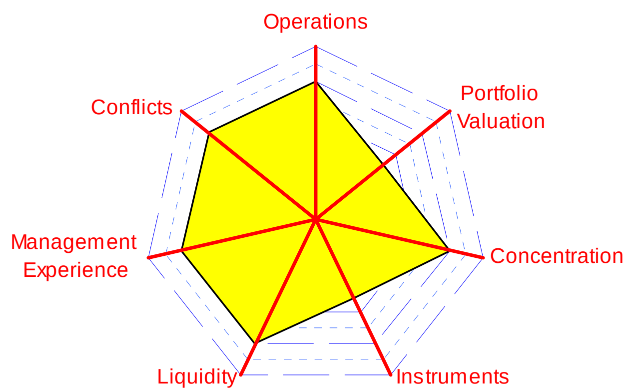

Chart 2.2

We will later see how this can be incorporated into a Markowitz portfolio theory in a CreditMetrics-flavored framework.

3. Risk and Performance Measurement

Hedge funds are appealing for two main reasons: they obtain higher returns than traditional investments, and they do so in a largely uncorrelated fashion. Because of this, performance measurement is also one of the key elements when understanding hedge fund behavior. However, business risk is substantial when investing in hedge funds, largely because of the unregulated nature of their operations; as a consequence, due diligence becomes a key exercise in determining risks of hedge fund portfolios.

3.1 Return

When looking at a hedge fund performance, there are two series of interest: the NAV (Net Asset Value) and the return. The NAV, which is really a number extracted from accounting statements and is (or should be) audited, is usually replaced by a theoretical analog, which is named with the acronym VAMI: it is the evolution of a hypothetical $1 (or $1,000) investment over time. Return for a hedge fund is usually measured on a **monthly** basis, and quoted on an **annualized** basis. If the series of monthly returns (in percentages) is given by numbers r_i, where the sub-index "i" denotes every consecutive month, and we have N observations, the average monthly return is given by

$$r = \frac{1}{N}\sum_i r_i$$

Because returns are expressed in percentages, one has to be careful, as the following -extreme- example shows: imagine a hedge fund with a monthly NAV given by $1, $2, $1, $2, $1, $2, etc. The corresponding monthly return series is given by 100%, -50%, 100%, -50%, 100%, -50%, etc. Its average return (say, after one year) is 25% monthly, which is an annualized return in excess of 300%, yet the hedge fund makes no real money. Ito's lemma (see [Hull]) is the tool we have at our disposal to translate average returns into performance numbers: in this particular case, Ito's lemma will tell us that we need to adjust the average return subtracting the quantity

$$\frac{1}{2}\sigma^2$$

which in our example will correct the average monthly return from 0.25 down to –0.30 (or -30%). The problem with this, of course, is that this correction it only works for Brownian motion: the limit of instantaneous performance with normally distributed increments; as we shall see later on, this is not an appropriate assumption for hedge fund performance. This already leaves us with an important open problem: how to decipher hedge fund return information. But things are about to get a lot worse.

In most situations, what we want is not so much to get the hedge fund return information, but the performance information for a portfolio of hedge funds. The accepted standard here is that, if we have hedge funds labeled by k and months labeled by i, monthly returns given by r^k_i, and percentage allocations to each hedge fund given by w_k, the return of the corresponding portfolio is the average of the returns of its constituents:

$$r_i^w = \sum w_k r_i^k$$

A quick calculation will show that this formula is not exact, only approximate. We will choose however to take it as good, because it allows us to set up Markowitz theory (see [Merton]) in the context of hedge funds and optimal hedge fund portfolio construction.

3.2 Volatility

Like returns, volatility is usually measured on a monthly basis, and quoted on an annual basis. Using our previous notation, the monthly volatility σ is estimated by the expression

$$\sigma^2 = \frac{1}{N}\sum_i (r_i - r)^2$$

One can also similarly estimate covariances and correlations, which measure the joint dependence of uncertain returns. They are applied to pairs of investments. If two investments have monthly return series given by numbers r_i^k and r_i^j respectively, with averages denoted by r^j and r^k, then

$$\sigma_{k,j}{}^{2} = \frac{1}{N}\sum_{i}(r_i^k - r^k)(r_i^j - r^j)$$

and their correlation is defined by

$$\rho_{k,j} = \frac{\sigma_{k,j}^2}{\sigma_k \sigma_j}$$

Because correlations and covariances are expressed in terms of pairs of investments, they are usually arranged in matrix form.

3.3 Efficient Frontier

The concept of the efficient frontier by Harry Markowitz (see [Markowitz], [Merton]) is one of the fundamental developments of modern portfolio theory. It is the result of a vast simplification of the portfolio concept, which is viewed as nothing but a supplier of return (given by its average return) at the cost of taking a risk (identified as the standard deviation). From this viewpoint, an equivalence relation is created between all possible investment portfolio and points in a plane. The subset of the plane where portfolios live is called *the feasible set*, one that can easily be verified to be convex, and the mathematical boundary of that two-dimensional set is denoted the efficient frontier. Efficient portfolios then include those whose return cannot be improved without increasing the risks, or risks cannot be diminished without hurting the returns. Efficient portfolios must be right at the efficient frontier.

According to this theory, rational investors should always invest in efficient portfolios. Markowitz therefore simplifies the investment problem from a two-dimensional decision process down to a one-dimensional one: the efficient frontier is a smooth curve where decisions are to be made, but the investor has the responsibility of making the final decision as to his risk level or return objective to select the appropriate portfolio on the curve.

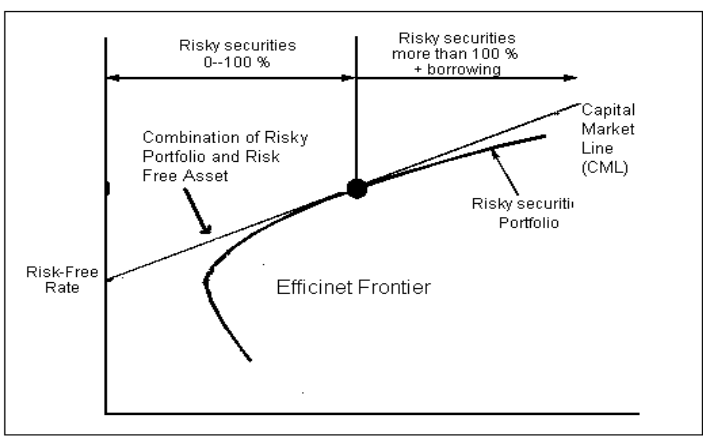

Markowitz Efficient Frontier and the Capital Asset Pricing Model

Chart 3.1

3.4 Sharpe's ratio

A way to bring return and risk into one number is by the Sharpe's ratio. In a certain sense, the Sharpe ratio will allow the investor to make an optimal choice on the efficient frontier: if Markowitz theory reduced the investment decision from a two-dimensional problem to a one-dimensional one, the Sharpe ratio will reduce it further from one-dimensional to a single point.

The Sharpe ratio uses the concept of the benchmark; in its simplest formulation, let's just say that our benchmark is the risk-free interest rate, which we will denote by c: then the Sharpe ratio is defined as

$$s = \frac{r - c}{\sigma}$$

It measures the average excess return per unit of risk. Portfolios with higher Sharpe ratios are considered to be better. The concept of the Sharpe ratio is one of the most useful in modern portfolio theory, but also perhaps the one that is most often misunderstood. The following fact is key to the understanding of the Sharpe ratio:

Imagine one is looking for the portfolio that has the best chance of optimizing its performance against a benchmark given by the risk-free interest rate. In other words, we are looking for the portfolio with returns Π that maximize the following quantity:

$$\Pr\{\Pi > c\}.$$

If we denote by μ and σ the expected return and volatility of our target portfolio Π, then *if its returns are normally distributed*, we clearly see that

$$\begin{aligned}\Pr\{\Pi > c\} &= \Pr\left\{\frac{\Pi - \mu}{\sigma} > \frac{c - \mu}{\sigma}\right\} \\ &= 1 - \phi\left(\tfrac{c-\mu}{\sigma}\right)\end{aligned}$$

where ϕ is used to denote the normal cumulative distribution function. Since this is an increasing function, we see that the portfolio that will maximize the probability of exceeding the benchmark is the same that will maximize the Sharpe ratio.

3.5 The normality assumption

In our discussion above, we saw that the Sharpe ratio is easier to understand when returns are normally distributed. Also, Markowitz theory, by virtue of the fact that it only looks are the average return and standard deviation of portfolio returns, is also making an implicit assumption that portfolio returns are normally distributed. This has prompted many authors and practitioners to claim that both theories are useless when viewing hedge fund investments, because of the non-normal nature of their returns, and a collection of alternative theories have been proposed.

We first analyze the normality assumption in some more detail:

Under the normal assumption, a portfolio with a 1% standard deviation will have annual returns which will vary no more than 1%, up or down, from its expected return, with a 65% probability.

If a higher degree of certainty about portfolio performance is desired, then one can say that the portfolio return will vary more than 2% from its expected return only about 1% of the time.

In the chart that follows, we can see the histogram of monthly returns of the CSFB Fixed Income Arbitrage Index. Assume, for the sake of the discussion, that this is a reflection of the return density of the typical hedge fund that arbitrages the debt markets. A simple inspection reveals that the distribution of returns is skewed to the left, indicating that losses occur more frequently and are deeper than predicted by a normal distribution. This phenomenon is usually referred to as a *fat tailed* distribution.

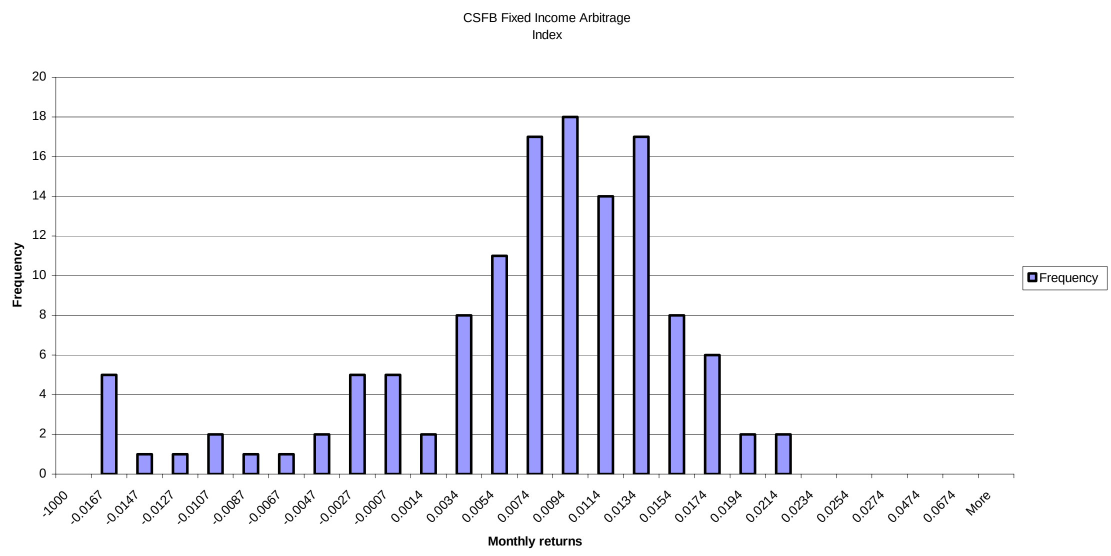

Chart 3.2

The informed reader would have seen the tragic events of 1998 as responsible for the large losses in this chart. The Russian default, coupled with the possibility of the impeachment of then President Clinton, drove the markets into a credit and liquidity crunch that resulted in the liquidation of Long Term Capital Management (a large hedge fund), and moreover evidenced a lack of credit risk management in the banking sector, and is largely responsible for BIS-II, the credit risk resolutions of the Bank of International Settlements. Incidentally, this is another piece of evidence of the psychological ubiquity of hedge fund activity.

Graphs such as this one have lead many people to consider ad-hoc variants of the usual Markowitz theory, which we now briefly summarize.

Gain deviation: it measures the deviation of portfolio returns from its expected return, taking into account only gains. In other words, portfolio losses are not taken into account with calculating the deviation. *Loss deviation* is the corresponding deviation when losses only are taken into account in calculating portfolio deviations. Both of these are used when one is trying to get a feeling as to the asymmetry of the gain/loss distribution.

Sortino ratio was introduced by Sortino and van der Meer in 1991, and is the substitute of the Sharpe ratio when one looks only at the loss deviation, instead of looking at the combined standard deviation. A common belief is that by not punishing unusual gains, like the Sharpe ratio does indirectly, one maximizes the upside while maintaining the downside.

Moments. One of the criticisms of the use of volatilities and correlations as risk measures is the presence of extreme events in portfolio returns, which will go unnoticed in those calculations. From a certain viewpoint, volatilities and correlations are measures inherited from normal distributions, according to which events such as the ones in 1987, 1995, 1998, etc. should have never occurred. One attempt to capture “tail events” is by introducing higher moments to measure large deviations: higher moments (centralized and normalized) are defined as follows:

$$\mu_k = \sqrt[k]{\frac{1}{n}\sum_{i=1}^{n}\left(\frac{r_i - r}{\sigma}\right)^k}$$

According to this definition, a normal distribution has moments given by 0, 1, 0, 3, etc. The third moment in this framework is called skew, and the fourth is called kurtosis; however, because the normal distribution has kurtosis 3, the quantity which is often used is called *excess kurtosis*, and it is simply the kurtosis minus 3. With this notation, the established belief amongst most investment managers is that positive skew is good, especially when kurtosis is big. As we are about to see, this is nonsense.

Do moments capture tail behavior?

Common wisdom is that they do indeed, as large events will dominate the sum. Extreme events, due to the larger power, affect the higher moments more than the lower moments. The conclusion then would be that using the information coming from higher moments we can reconstruct the return probability distribution more accurately than using means and standard deviations.

The reality is quite different when dealing with hedge funds, for a number of reasons. Most importantly, one is confronted with the problem of limited access to data; many funds have been operating for only a couple of years (hedge funds, as opposed to mutual funds, will close to new investments when they reach capacity, which can occur fairly quickly.) When data is scarce, large events introduce distortion in the higher moments; the distortion is greater the higher the moment and the shorter the return history.

The chart below shows the evolution of the skew estimation of a collection of hedge funds using a rolling historical window: there we can see that the estimate for skew has a sudden jump when one takes into account *just one* more month of performance: July 2002.

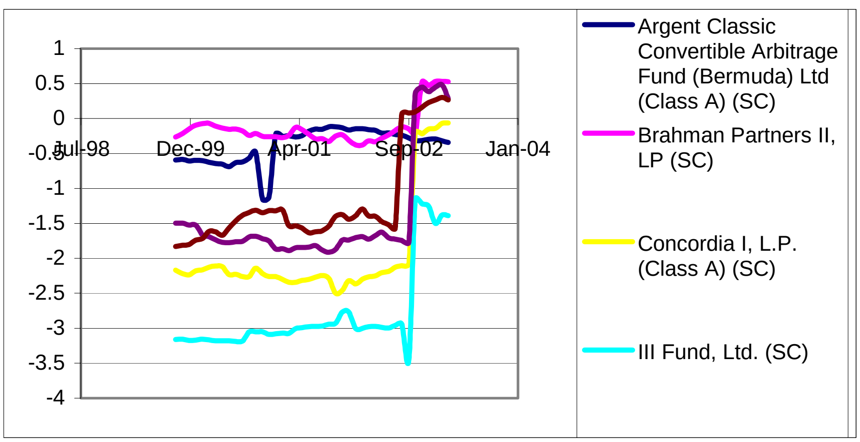

Chart 3.3

The biggest problems with moments, however, come when we realize that we are dealing with investments, not just statistics. Even if we could measure moments, the common belief that positive skew is good, and negative skew is bad, is totally unfounded when applied to investment returns. We already touched on this indirectly earlier when we were defining the average return, and we considered the series \$1, \$2, \$1, \$2, \$1, \$2, etc. In the following chart we present the NAV evolution of an investment that, if its return statistics is everything we look at, gives us a very attractive answer from the commonly accepted view: its average monthly return is a wonderful 7% (which annualizes to about 100%!), and a very positive skew: +30%. Its kurtosis, by the way, is near 0. Yet the investment loses money, and lots of it.

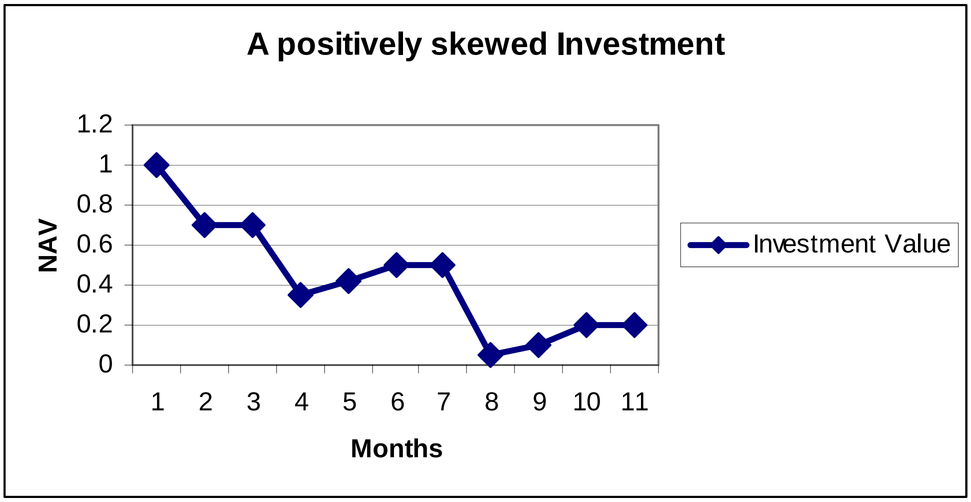

Chart 3.4

To summarize, two things are wrong with moments; first, they are hard to observe (un-robust, lack of good estimators, lack of sufficient data, etc.) Second, investment returns and return statistics do not get along.

W. Shadwick and C. Keating introduced the concept of "Omega" a few years ago, as the replacement of the Sharpe ratio when returns are not normally distributed. His aim was to capture the "fat tail" behavior of fund returns. His proposition was that, once the "fat tail" behavior has been captured, one only needs to optimize investment portfolios to *maximize the upside, while controlling the downside.*

Omega is defined as follows:

$$\Omega(r) = \frac{\int_r^{\infty} (1 - F(x))\ dx}{\int_{-\infty}^{r} F(x)\ dx}$$
$$= \frac{\int_r^{\infty} (x - r)\,\rho(x)\ dx}{\int_{-\infty}^{r} (r - x)\ \rho(x)\ dx}$$
$$= \frac{\mathrm{E}\left[\max(R - r, 0)\right]}{\mathrm{E}\left[\ \max(r - R, 0)\right]}$$
$$= \frac{\text{Dembo's reward}}{\text{Dembo's regret}}$$

It is of interest to contrast this with the papers by Dembo and his co-athors. As a statistical observable, Omega presents intriguing possibilities: it measures the proportion of over-performance relative to under-performance, and as such it is conceptually clear. But remember that statistical observables for returns do not, in general, mix well with investments performance. In fact, the Omega of our positively skewed investment earlier is larger than 1 (for most benchmarks), despite the fact that it loses money.

The largest misconception about the Omega, as used by many investment managers, is that it is a "curve", not a "number"; indeed, Omega is most useful when considered a function of the benchmark return r, as it is equivalent to the cumulative return distribution and expresses the statistics in investment terms. However, as a single number, it carries less information about the tails of the distribution as, for instance, the standard deviation; indeed, as a single number Omega is a ration of truncated first moments, so its information content about tails is less than the variance, a second moment.

The role of the benchmark is key to understand and use Omega: if one is interested in risk management, then the benchmark should be chosen to be large negative; the larger the Omega in that case, the safer the investment. However, a large Omega when r is negative does not provide any insight as to the probability of obtaining target returns in line with moderate portfolio objectives. Once again, we have to take into account the values of Omega for various values of the benchmark, and somehow balance one against the others: we lack an Omega "Benchmark utility theory".

For hedge fund portfolio optimization, it is not clear which benchmark to use in the optimization process, but we see that the use of small benchmarks would be attractive for risk management, but would give up portfolio returns, whereas the use of medium-sized benchmarks will be useful in portfolio management but will ignore tail effects.

4. Correlation risk

Hedge funds are uncorrelated to traditional markets, so they constitute excellent diversification strategies. At least, at first sight. At any rate, diversification is the main property that hedge funds bring to investment portfolios. Situations in which this diversification disappears should therefore be considered to be grave.

We discussed earlier in this paper the normality assumption for hedge fund returns. When we look at returns of many funds simultaneously, normality can be lost for two reasons: the marginal distributions are not normal (as we discussed) or the dependence structure is not determined uniquely by the correlation. The purpose of this section is to elaborate on the latter, presenting a methodology to analyze *correlation risk.*

When we presented the performance histogram of the CSFB Fixed Income Arbitrage Index, we noted that there are too many large losses that cannot be explained by a normal distribution.

To understand other type of events that cannot be reconciled with a normal distribution, imagine you use one hundred normal coins: each gives you heads or tails with a 50% probability, and you flip all of them simultaneously several times. Imagine also that, when you do that, you observe that approximately 10% of the times, all of the 100 coins give tails at the same time. In this case, there is nothing abnormal about each coin's output: what is non-normal is their dependence structure.

Something similar happens with hedge fund events; in addition to their own individual performance and possibly a fat-tailed behavior, there are months in which too many funds exhibit the same behavior; in May 2005, the vast majority of hedge funds lost money; something similar happened in April 2004, and of course we still remember 1998, despite the fact that their mutual correlations are around 20% on average. In other words, their mutual behavior is something that according to a normality assumption should only occur once in a hundred years.

We will investigate these facts from a regime switching perspective; to this end, we will distinguish between two states of the market: normal and distressed, and we will refer to the normal distribution as Gaussian, to avoid confusion between the two market states.

Mixtures of Gaussians are standard in the study of fat-tailed distributions in the univariate case. Multivariate Gaussians are relevant for our study because they offer a framework to understand how uncorrelated returns can become correlated under distressed market conditions.

An *n*-variate probability distribution is a mixture of two Gaussian distributions when the value of its density function *f(X)* at an n-dimensional vector *X* can be expressed as a linear combination of two Gaussian densities as follows:

$$\frac{pe^{-\frac{1}{2}(X-M_1)^t A_1 (X-M_1)}}{\sqrt{\det(2\pi A_1)}}+\frac{(1-p)e^{-\frac{1}{2}(X-M_2)^t A_2 (X-M_2)}}{\sqrt{\det(2\pi A_2)}}$$

Here, *M* denotes each of the two means in each of the two states, and *A* denotes each of the correlation matrices in each of the two states. This models a process where each event in a time series is normally distributed, but it follows one or the other distribution with probability *p* or (1-*p*), respectively (from here on, we shall refer to *p* as the *mixing coefficient*).

To work with such a distribution, it is necessary to calibrate it using market data, in order to find the parameters (means, variance-covariance matrices, mixing coefficients) that uniquely define it. A better procedure is to use maximum likelihood estimators, or Bayesian methods. However, the drawback of this approach is that the resulting distribution offers no insight as to what constitutes distress. For the purpose of this article, we will offer a simple alternative that explains market behavior, while remaining compatible with maximum likelihood estimators. In particular, we will use this to model “normal” situations, in contrast with situations of *market distress*. This is done as follows:

Define distress in terms of the *deviation* of a particular hedge fund from its mean in either direction. More precisely, we say that a hedge fund is *in distress on a given month* when its performance r deviates from its mean μ by more then a given number n times the standard deviation: i.e. $| r - \mu | > n * \sigma$ (say n=2 for simplicity; n=2 is also meaningful as events more than two standard deviations away from the mean should be practically impossible according to a Gaussian rule). Continuing with our literary license, we will call a certain month *distressed* when sufficiently many funds are in distress. One would find that the hedge fund universe is in distress about 25% of the time.

Note also that our definition of distress allows us to account for exceptionally bad months as well as exceptionally good months (such as May 2003). The opportunistic investment reader may be discouraged by this, as it appears to be something that may result in confusing good and bad performance. Our ultimate objective, however, is very different: we will just want to be able to identify risk convexity, which only manifests itself in large event situations (positive or negative), hence both type of events, good and bad, are going to be very useful in determining risk commonality throughout the hedge fund universe.

To understand the value of this approach, consider the following exercise: take a large collection of hedge funds (say 30) and measure their correlation during the normal times and during distressed times separately; what we get is that the correlation during normal times is given by

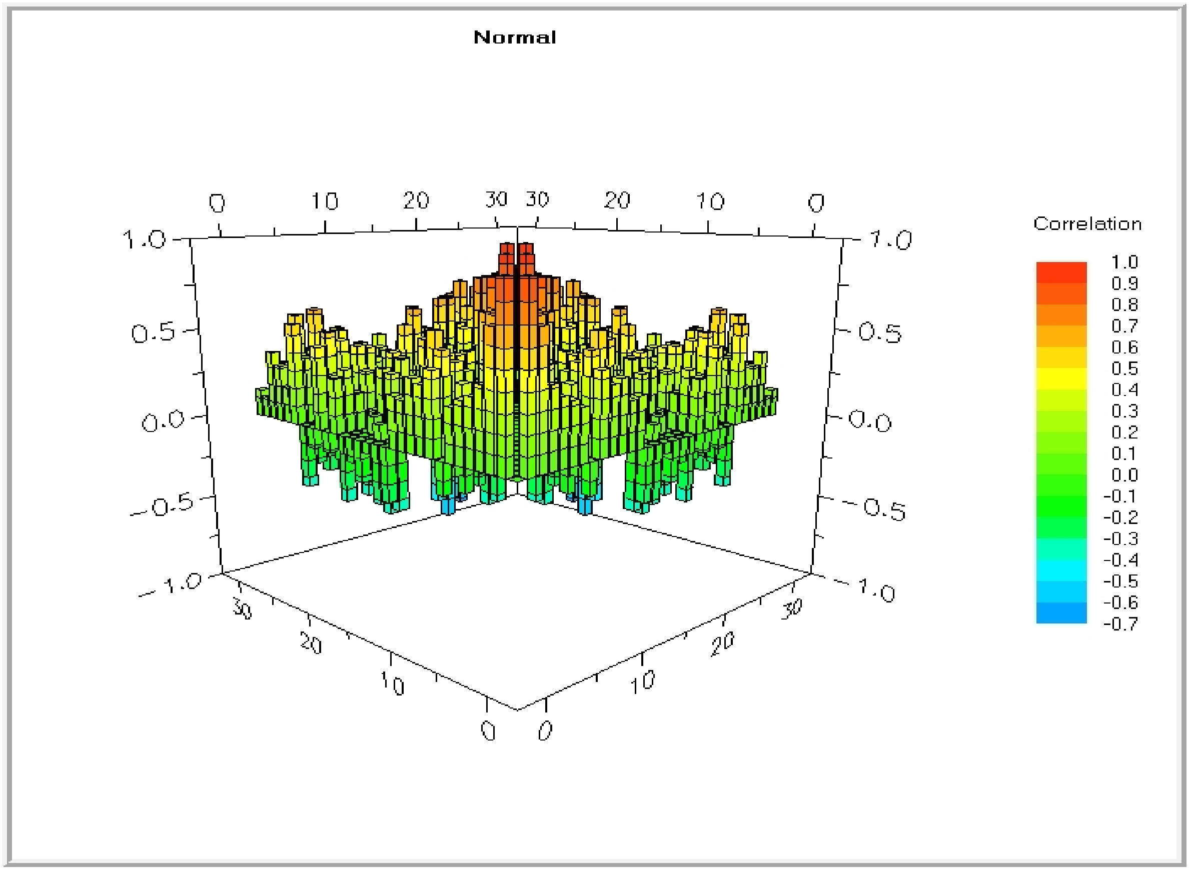

Chart 4.1

Here we are adopting a pictorial approach to correlation matrix display: to each number in the square array, we attach a high red pile when it is close to +1,and a low blue pile when it is close to –1. This allows us to visualize a matrix easily.

When we look at the distressed correlation matrix, we get instead

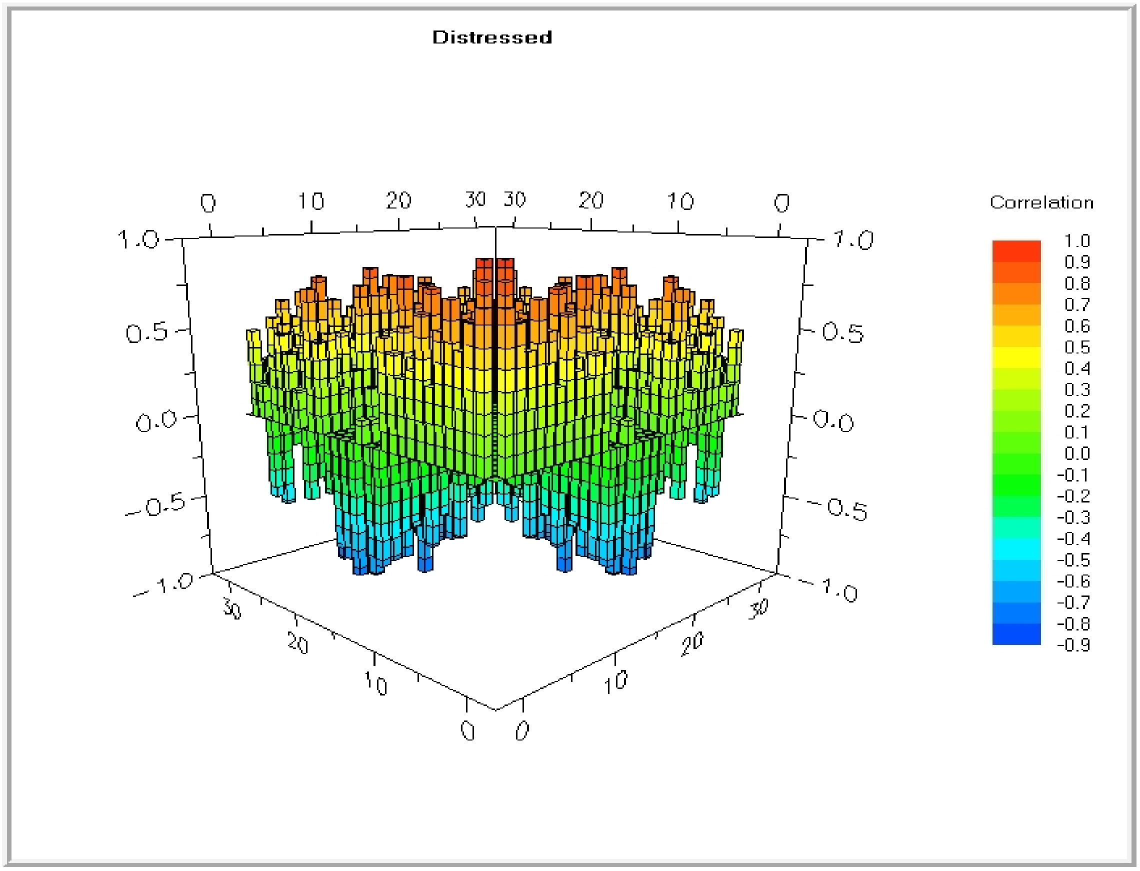

Chart 4.2

What we see here is very interesting: many of the correlations are driven higher, close to +100% during times of market distress; but a handful of them actually become lower, identifying in this manner hedge fund which help with diversification in the worst of times.

We can take this pictorial approach a little further, turning each element in the matrix into a pixel with the same coloring rules; this allows us to summarize both matrices by a combined 2-dimensional aerial view, which we can use to quickly display the correlation properties of even larger collections of hedge funds (we took 100 in the example below):

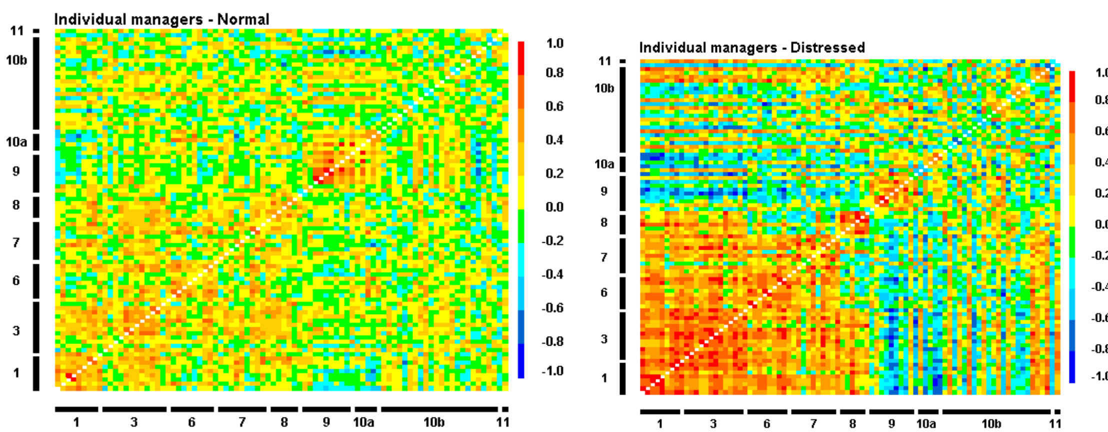

Chart 4.3

The conclusions are the same in both situations: while the correlation numbers are fairly low during normal market conditions, they spike to +1 during market distress, except that a handful of managers actually gives us even lower correlation during market distress; for the cautious investor, these managers bring a type of investment insurance with them that, if their individual returns and risks are appropriate, may make them very desirable for a balanced portfolio.

What all of this is pointing to is a need for a Markowitz theory that takes into account regime switching. We refer the interested reader to the article by Bukley, Saunders and Seco for a mathematical treatment of these issues.

Later on we will see other applications of this point of view, when we discover a dual cluster in risk exposures in the equity arbitrage hedge funds.

5. Hedge Fund strategies and risks

There are many hedge fund strategies, and new ones are created every year. In order not to abuse the patience of the reader, in this section we will focus on 4 of the traditional ones in detail: convertible arbitrage, fixed income arbitrage, equity hedged and commodity trading. For the reader interested in other strategies, such as merger arbitrage, risk arbitrage, distressed securities and global macro, we refer to the text by Jaeger, which describes them in some detail. But perhaps the most interesting aspects of hedge fund styles is the new trading strategies which are being developed; energy trading, leasing arbitrage, commodity factoring, catastrophic bond and weather derivatives are some of them. Unfortunately (or fortunately, depending on the point of view), almost any attempt at general description of these strategies will either fail or not hold for a very long time, as they are evolving very rapidly.

5.1 Convertible arbitrage

The convertible arbitrage strategy uses convertible bonds; they can be hedged in several ways, normally by shorting the underlying common stock. The quantitative valuation is overlaid with credit and fundamental analysis to reduce the risk further and increase potential returns. Often, emphasis is placed on growth companies with volatile stocks, paying little or no dividend, with stable to improving credits and below investment grade bond ratings.

Convertible arbitrageurs construct long portfolios of convertible securities and hedge by selling short the underlying stock of each security. Convertible securities include convertible bonds, convertible preferred stock, and warrants.

The price of the convertible declines less rapidly than the underlying stock in a falling equity market and mirror the price of the stock more closely in a rising equity market.

Arbitrage opportunities are identified using valuation models that locate a “cheap” convertible security: the market value of the convertible is less then the expected value, given the market price of the underlying security, interest rates, credit quality, implied volatility, and time expiration or call probability.

Cash flows for these strategies are usually given by the bond yield, short interest rebate and a small outflow for short stock dividends. Trending stock price movement provides the opportunity for trading profits while holding the arbitrage position. Since hedge adjustments capitalize on stock price volatility, trading profits can be expected during normal market conditions of stock volatility and stable interest rates.

Consider the convertible arbitrage strategy example of a bond selling below par, at $80.00: it has a coupon of $4.00, a maturity date in ten years, and a conversion feature of 10 common shares prior to maturity. The current market price per share is $7.00. The client supplies the $80.00 to the investment manager, who purchases the bond, and immediately borrows ten common shares from a financial institution (at a yearly cost of 1% of the current market value of the shares), sells these shares for $70.00, and invests the $70.00 in T-bills, which yield 4% per year. The cost of selling these common shares and buying them back again after one year is also 1% of the current market value.

Scenario 1: Values of shares and bonds are unchanged:

	Today	1 yr later
Bonds	80	80
Stock	-70	-70
T-Bill	+70	+72.8
Coupon		4
Fee		-3.5
Total	$80	$83.3

Chart 5.1

Scenario set 2: In the next two examples, the share price has dropped to $6.00, and the bond price has dropped to either $73.00 or $70.00, depending on the reason for the drop in share market values. The net gain to the client is 7.87% and 4.12% respectively, again after deducting costs and fees.

	Today	1 yr later (a)	1 yr later (b)
Bonds	80	73	70
Stock	-70	-60	-60
T-Bill	+70	+72.8	72.8
Coupon		4	4
Fee		-3.5	-3.5
Total	$80	$86.3	$83.3

Chart 5.2

Scenario set 3: In the following three examples, the share price increased to $8.00, and the bond price increased either to $91.00, $88.00 or $85.00, depending on the expectations of investors, keeping in mind that we have one less year to maturity. The net gain to the client is 5.37% and 1% in the first two examples, with an unlikely net loss of 2.12% in the last example.

	Today	1 yr later(a)	1 yr later(b)	1 yr later(c)
Bonds	80	91	88	85
Stock	-70	-80	-80	-80
T-Bill	+70	+72.8	+72.8	+72.8
Coupon		4	4	4
Fee		-3.5	-3.5	-3.5
Total	$80	$84.3	$81.3	$78.3

Chart 5.3

The reader with a penchant for risk management will quickly see here the beginning of a scenario analysis of a convertible trade that could easily be turned into a Value-at-Risk calculation. Also, it is interesting to note that a correlation breakdown in the bond-equity relationship will also give rise to much unwanted losses.

It is of interest to note that convertible arbitrage is the hedge fund style which, in 2004 and 2005, is having the most difficult time obtaining returns; performance in the first halt of 2005 is down by over 6%; while not unusual for equity markets, this is very tragic performance for a hedge fund group.

5.2 Fixed Income Arbitrage

This is a denomination for what is really several different trading styles; their common denominator is that hedge funds of this type make their returns by trading yield curves. Beyond this, they could be trading credit spreads (high yield/credit arbitrage), mortgage related securities, currency/yield arbitrage, or simply treasury curve relative value trading.

Risks that this group could exhibit exposure to include:

- Leverage: the relative value obtained from a single position is typically small, so leverage is necessary to boost its performance to reasonable levels.

- Liquidity: by the mere nature of curve trading, the notional that the positions control are typically several orders of magnitude the equity capital. For large funds, this can present liquidity problems, as the underlying positions may be hard to unwind.

- Model risk: interest rate markets are difficult to model, and hedge funds operating in this sector are about the most mathematically sophisticated. Model risk can affect their operations.

❖ Credit: if the hedge fund performs high yield arbitrage trades, they might be dealing with bonds with of low credit rating or inadequate collateral.

Perhaps the best-known hedge fund of this type was Long Term Capital Management. Set up by John Meriwether in the early nineties, it reached several billion dollars in size by 1998. The Russian default and the treasury liquidity in the US during the summer of 1998 lead the fund to its demise. Today, the largest fund operating in this style is Madrid-based "Vega", a highly sophisticated $11B hedge fund that saw its origins in capital provided by Banco Santander, and is now spinning off other sub-hedge funds in areas as unrelated to their origin as energy trading.

The events if 1998 created a myth about this style which can be best summarized as "Fixed income is bad"; we will see however, that this is quite far from the truth. Let's start by taking a look at their performance in 1998; the HFRI (Hedge Fund Research Index) performance in 1998 is given by the following chart:

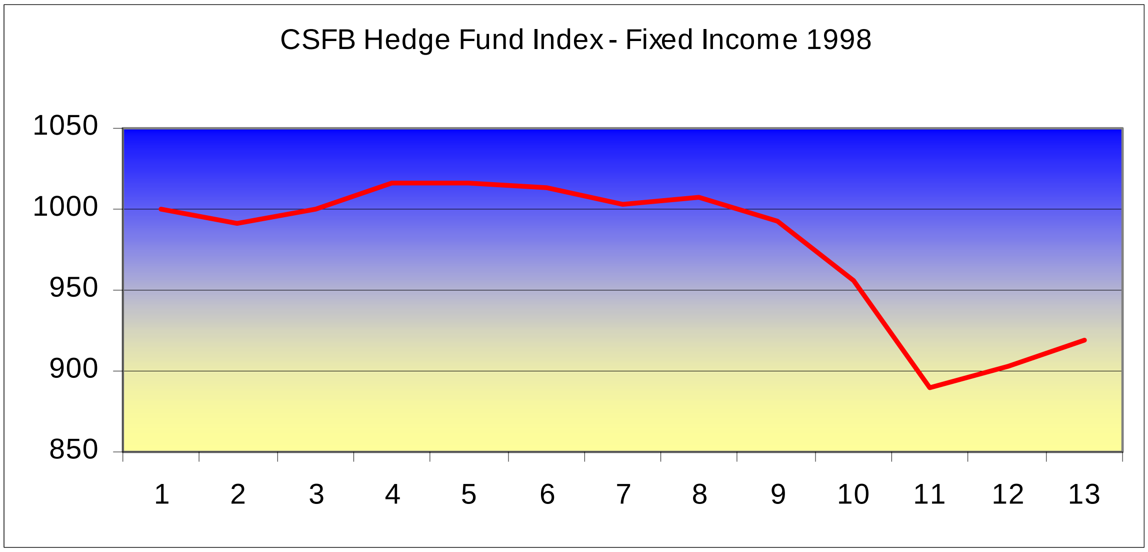

Chart 5.4

The index ended the year losing about 8% for its investors: this is a large loss in hedge fund land. By contrast, equity arbitrage ended the year making 17%, and equity market neutral made 13%: a very attractive return.

Let's now contrast the mythology of 1998 with the one created in 2003, which reads: "Hedge funds (equity hedged) are good"; this myth was created because of the following performance summary:

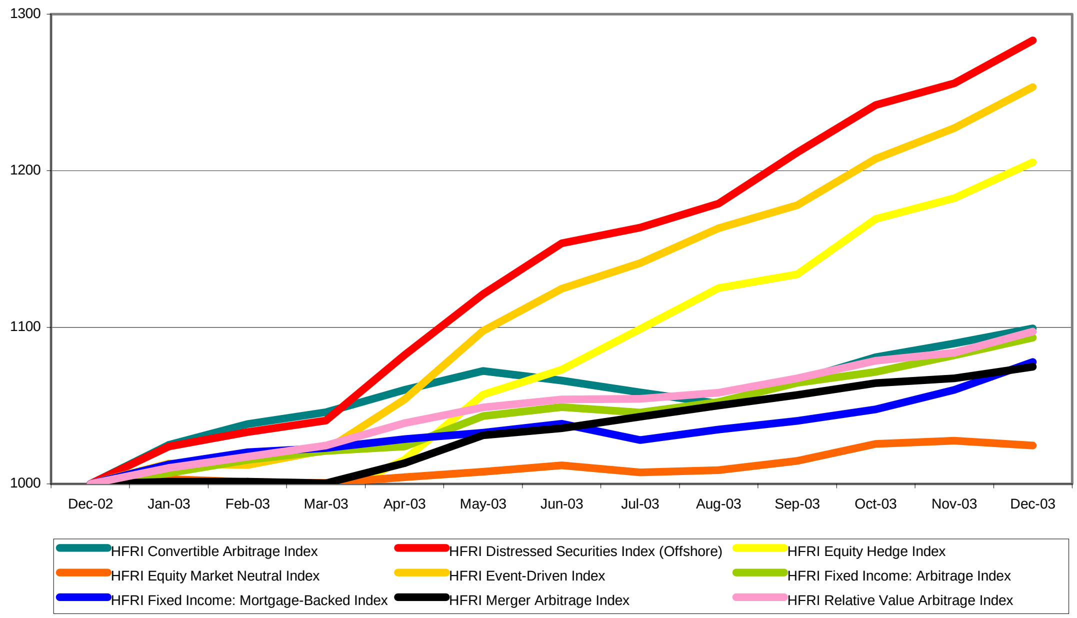

Chart 5.5

Equity hedge funds appeal to the retail investor: they buy and sell stocks, just like many retail investors tried to do in 1999 and failed in 2000; their return in 2003 clearly sounded attractive, and that style was single-handedly responsible for the huge growth of the industry that year. It is interesting to note that fixed income ended 2003 with a “modest” performance, 8%, compared with the equity styles obtaining double-digit returns.

Interestingly enough, 2004 presented the following picture by mid-year:

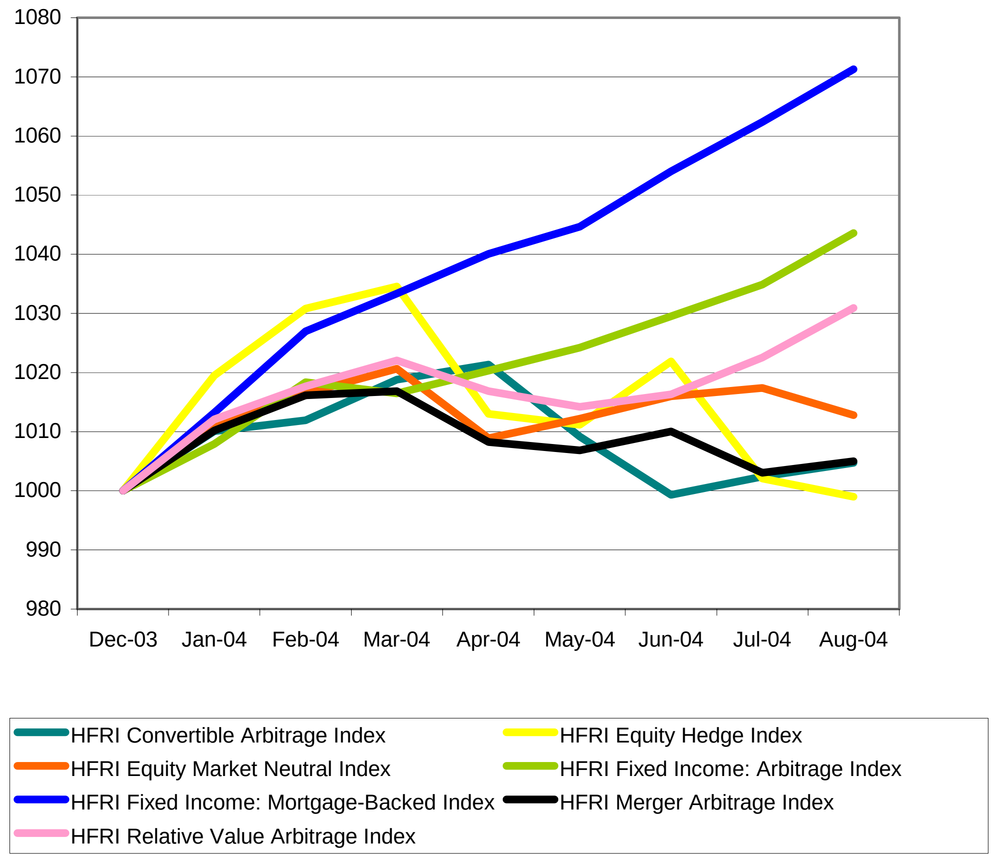

Chart 5.6

Since the public's attention was mostly with the equity hedge funds, this peformance gave rise to a new myth: "the party is over". But, as we can see in the chart, fixed income was still enjoying very attractive returns.

The lesson from this visual analysis is that fixed income and equity arbitrage are two trading styles which lead to low-correlated returns, which actually turn into uncorrelated precisely when one of them is having difficulty; in the language developed earlier, we should say that they are uncorrelated in times of distress, and provide mutual diversification to each other precisely in those times when diversification is hard to find.

5.3 Equity hedged/Equity market neutral

After comparing equity market neutral and equity hedge to fixed income strategies, let's now analyze this style in detail.

There are basically two types of managers in this broad sector; one of them trades in stocks only, but can take long or short positions. The objective is to hedge the long positions with the short ones, and obtain returns independent of market direction. The other one takes equity positions (long or short) and hedges the market exposure with futures, indices or derivatives, including options. These are the managers that are responsible for the term "hedge fund', although when the term was coined in the seventies they operated mostly inside the private banking group of the large banks.

As a trading style, they can be exposed to a variety of risks: Equity direction and volatilities, liquidity, operational risk, corporate events, etc. They can also be exposed to credit spreads if they employ large/small cap spread trades.

Let's focus on the detailed analysis of one of these exposures: exposure to equity volatilities and market direction. Both risks are actually related, as is well known and we summarize in the following linear regression chart:

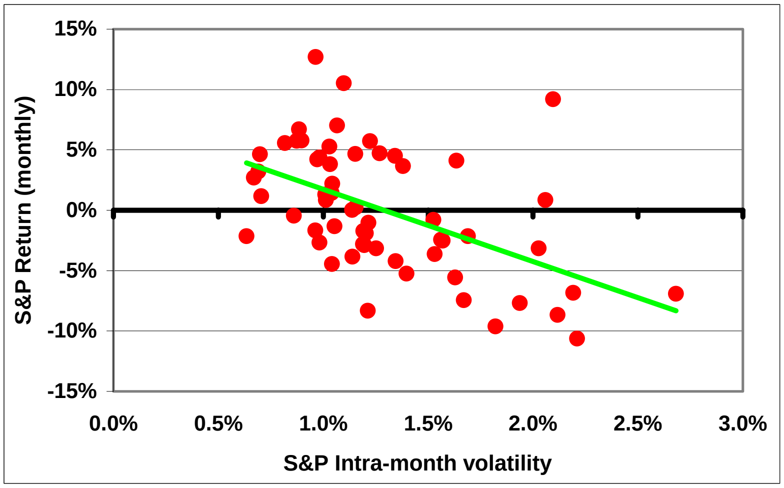

Chart 5.7

With this into account, the chart below shows that the overall correlation of these funds with US equities is biased to the positive side. However, it also shows that, if we slice this correlation picture by segregating bull and bear periods, we find that funds are more positively correlated to the equity markets when the markets are going down, and less when they are going up:

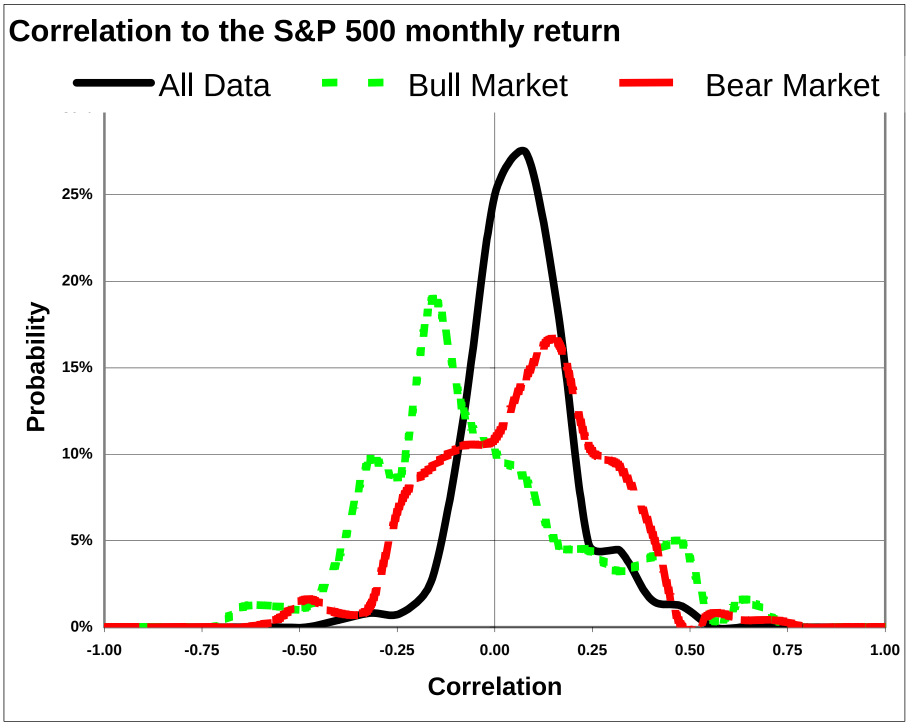

Chart 5.8

This is not particularly encouraging: it signifies an endemic problem with the group to deal with market volatility, and perhaps market liquidity.

However, perhaps the most interesting aspect of this style is that there is an internal duality; to this end, let's analyze the normal and distressed correlation picture inside it; what we obtain is the following correlation histogram:

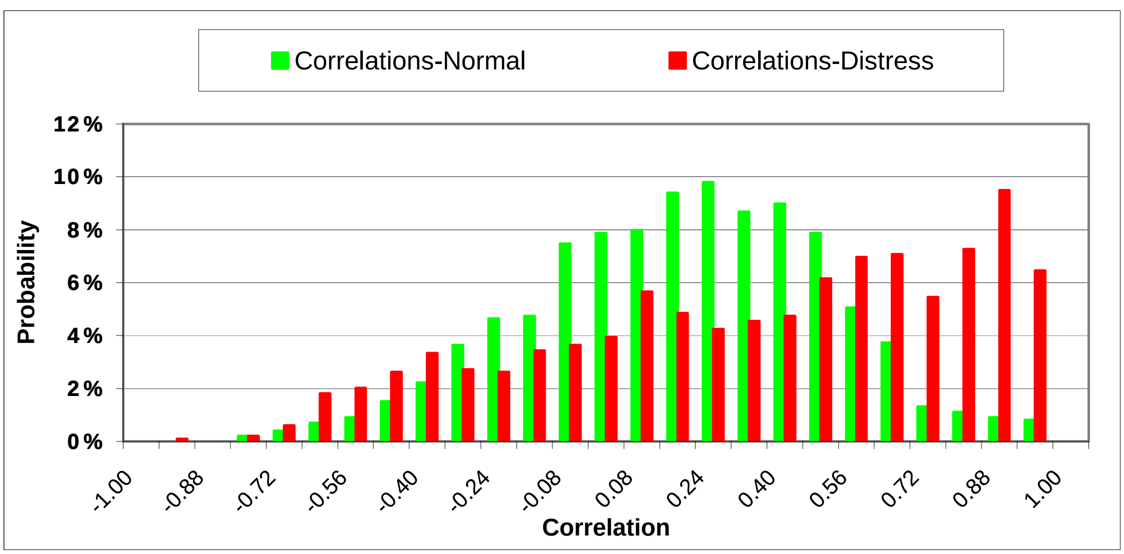

Chart 5.9

This means that, while the hedge fund style is internally uncorrelated during good periods, during market distress there is a correlation cluster near the 100% level. But note that there is also an increase in the negative correlation cluster during distress: the implication of this is major: there are some managers of this style that provide increased diversification against their peers during times of market volatility. This has very important implications for the construction of stable fund of hedge fund portfolios.

5.4 Managed Futures

Often associated with hedge funds, they share with them their economic independent returns, but offer a very different operational profile.

In the middle ages, buyers and sellers of commodities met annually at trading fairs to lock in future needs and prices. Middlemen provided banking and storage to facilitate trade. In the XVI century, rice trading in Japan developed to the point that an organized futures market was used by producers and consumers to transfer price risk to each other. Futures markets permit price certainty and increased economic activity. Futures function as an insurance market, providing price certainty to commercial entities. They are liquid, public markets that preclude special "inside" information. They are extremely regulated, and do not allow trades off-exchange. Exchanges create uniform product, so quality, delivery date and location of product are not variable. Investors concentrate on price trends. They also offer complete continuous disclosure of price, in 30 different countries.

Futures Protocols work as follows: investors provide initial margin or good faith deposit. Mark to market losses must be topped up every day. Customer receives gains of profitable positions daily. All traders participate as equals once margin is posted. Speculators have the same access as the most credit-worthy bank. The clearinghouse is the guarantor and counterparty of all trades, which are anonymous. Member firms only post their "net" position to the exchange. In addition, as futures contracts have no initial value (they represent a future obligation) they offer "free" leverage.

Futures prices are rarely a prediction of the future. They are a reflection of the current market, adjusted for interest rates, storage costs and seasonal factors. Grains, meats and other foods, energy, currencies, debt and equities all have futures markets. They often have very pronounced mean reversion properties (except equities). Price movements create supply / demand changes which in turn affect price. Contango is a carrying-charge market: a storage charge is built into the forward price. Backwardation is an inverse market with no carry charge because of high demand. They are easily characterized by the shape of the forward curve: backwardation leads to a decreasing, convex forward curve, while contango leads to an increasing, concave curve. The relationship between spot and future prices has significance in managed futures and passive futures profitability.

There are 3 differing skill sets in this universe:

- Trend followers: longer term traders, they exhibit large returns with large volatility and trend reversals can be very damaging to them.

- Short term and counter trend managers. Less volatile than trend followers, make their money in different markets (short trends, daily volatility, etc.)

- Specialty and fundamental managers. They are often driven by an economic view, often similar to the global macro hedge fund group.

Managed futures are not dependent on inflation or deflation, nor are they considered foreign property.

Their best characteristic is their totally different risk profile to other hedge funds. For example, the average manager had large gains in Aug. and Sept. 1998 and summer of 2002: in our language, they posses negative distress correlation with most other trading styles; they have no credit or liquidity risk; they operate on notional trades using cash exclusively for margin, leaving capital usually in cash or T-bills.

6. Manager default risk

Business risk in the hedge fund industry is synonymous with operational risk, and it refers to everything that isn't market risk. The best-known business risk event is the blow-up of Long Term Capital Management in 1998, since this event not only involved the investors in the fund, but perhaps was even more dramatic for the financial system as a whole, given the size and the number of players that were involved in it. The events started when Russia defaulted on its debt, and investors' confidence was shaken due to President Clinton's impeachment possibility; the result was a massive credit and liquidity crunch, which lead to the freezing of LTCM's assets (largely uncollateralized positions), inability to unwind positions, publication of the securities that were the source of the problems, and the market reaction to someone in need. This provided a nightmare scenario for the design of the BIS-II resolutions on credit risk management.

While not of the same size, events such as this one occur every year: Silvercreek (brought down by Enron's demise), Lancer (fraud), Beacon Hill (fraud on the part of the managers to attempt to survive massive refinancing orders in the mortgage market), or the Appalachian illegal dealings in mutual fund timing are a few examples of situations where investors lost money for reasons other than normal market fluctuations; a hedge fund "blows-up" typically refers to a situation where losses occur and law enforcement acts; the result is always forensic investigations, fund liquidation, and press coverage.

Modeling this risk in an investment portfolio is critical. All fund-of-fund managers contain a "due diligence team" that check that each of the underlying hedge fund investments are "sound"; equipped with a good nose and a magnifying glass, they visit the hedge funds, analyze their positions, follow trade tickets, ascertain independence of the administrators and true asset valuation processes, request audits, and sometimes hire private investigators to look for dark spots in the history of the firms and its key employees. This exercise will be repeated often, annually or perhaps quarterly depending on the taste or suspicions of the manager. If the result of all this is good, the investment manager approves the hedge fund, an investment may take place.

In order to get a sense as to how to go beyond the due diligence process, and indeed to build the due diligence process into the portfolio construction process of a hedge fund portfolio, consider the following example.

6.1 An example

Assume the entire hedge fund universe consists of 100 managers, and every year, one out of the 100 will have a blow-up event (fraud or any other), which will lead to a total loss of the investment in that one manager (no recovery rate). Assume also we have a $100 M portfolio invested in hedge funds.

Consider two alternatives:

- A portfolio that invests in all the money in one manager alone.
- A portfolio that invests equal amounts of money in every single manager ($1M per manager)

In deciding what to do, a risk averse investor will not go with the first choice; it is better to diversify and go with the second choice, knowing that the loss is guaranteed to be $1M (a diversification premium, or insurance cost), and knowing that the portfolio will produce positive returns as long as the average return for the other managers is in excess of 1% per year.

The portfolio that invests equal amounts of money in every single manager ($1M per manager) has the following characteristics:

- The probability of having one credit event inside the portfolio is 100%; it is certain that one of the funds will have a credit event, by assumption.

- The expected loss due to a credit event is equal to the loss that we know we will get for sure, which is $1M.

- The standard deviation of losses is 0; there is no uncertainty as to how much the portfolio will lose: it will lose $1M.

The portfolio that invests all the money in one manager alone has the following characteristics:

- The probability of having a credit event inside the portfolio is 1/100; it is smaller than the previous portfolio; in fact, this portfolio will have the smallest probability of a blow-up event.

- The expected loss due to credit events is equal to $1M. This expectation is the same as the original portfolio, and is in fact the same as the blow-up probability of any other portfolio.

- The standard deviation of credit losses is about $9M.

The conclusion of this example is clear: the standard deviation of losses is the key indicator of the business risk; the probability of blow-up, or the expected loss carries no information.

In the section to follow we will propose a methodology to incorporate a manager rating system into a portfolio construction process.

6.2 Malfunction frequency and severity

CreditMetrics is one of the two main methodologies for credit risk management. It is based on a rating system for each counterparty, from which a probability of default is derived.

For hedge fund portfolio management, an analog can be established, as follows:

First, we introduce a rating system for hedge funds. This rating system will aim at quantifying the due diligence process we described earlier.

Instead of the septagonal version we presented, let’s consider a simpler one consisting of only two dimensions: one for blow-up frequency, another for severity, as follows:

	A	B	C
A	Infrequent, small losses		Frequent, small losses
B			
C	Infrequent, large losses		Large, probable losses

Chart 6.1

According to this, ratings will be pairs of letters, such as AA for the best, and CC for the worst.

Next, we associate a frequency distribution to the vertical silos and a severity distribution to the horizontal ones, very much like CreditMetrics. Unlike CreditMetrics, we won't have history to calibrate these numbers, and they will be the responsibility of the due-diligence team to determine. Let's just say that, for a given fund, the blow-up frequency is given by a number p and the severity is given by a random variable X, of which we can generate samples denoted by $x_1, x_2,$etc. If the hedge fund has a return history given by $r_1, r_2, \cdots, r_n.$, we calibrate it to our favorite distribution and generate a sufficiently large future distribution of returns, which we will also denote by $r_1, r_2, \cdots, r_n.$ (except they could span one hundred years into the future, if we so wanted).

After this, we select pn of those r_i, and we simply replace them by x_i. The result is a modified future return distribution that takes into account possible manager malfunction. One then optimizes the portfolio construction using this generated data instead of the historical one.

The resulting portfolio will be one that takes into account market and blow-up risks simultaneously, and is an interesting way of extending Markowitz theory to this context.

7. Structured products

After the popularity that fund-of-hedge funds have gained in the investment world, structured products that use hedge funds as their underlying securities are beginning to be common.

7.1 Leveraged products

In a leveraged product, the fund provides capital, a bank provides a loan; the total of both is invested in a fund-of-funds structure. The NAV is calculated as the value of the underlying assets, minus the capital belonging to the bank and financing fees payable to the bank (typically LIBOR plus 1%). They can be very advantageous to the investor, if the return of the underlying funds is in excess of the financing cost. Since financing was very cheap in the low interest rate environment we have lived through, these structures have been extremely popular.

Leveraged structures are often issued as non-recourse loans: any losses in excess of the fund's assets are absorbed by the bank. In other words, the fund is the only collateral the bank has for the loan. The bank usually protects its downside by introducing a knock-out provision: it will de-lever the loan as the fund pool loses money; for a product with a leverage of 4-to-1 ($3 borrowed for each $1 invested), this knock-out provision can kick in when the fund loses as little as 5%.

7.2 Options on hedge fund baskets

In these structures, the fund purchases a call option on a synthetic fund-of-funds. The notional of the investment is usually large, a function of the volatility of the underlying hedge funds; the fund's assets are used to pay the premium, calculated within the Black-Scholes formalism, and the notional is implied by the option price. In this way, investors are exposed to the upside of the fund's performance with a notional investment which is orders of magnitude the size of their original investment (the fund's assets).

They work very similarly to the leveraged structures: in fact, we may view the leverage structure as a portfolio that replicates the option, at a slightly higher cost. In practice, their use is different; leveraged structures are used for investments that offer liquidity; for investments that offer little or no liquidity (for example, guaranteed notes or investment certificates), then the option will be the vehicle of choice, as it may be cheaper.

It is interesting to note that one must use the Black-Scholes framework here with extreme care, as the underlying investment in hedge funds does not satisfy many of the assumptions of the theory: we have no liquidity, no short-sales, and no Gaussian returns.

7.3 Notes and CPPI structures

These investments, a hybrid between a fund-of-funds and a leveraged structure, provide the investment returns of a fund-of-funds but in addition guarantee the initial investment at a later maturity time. Let's use the following imaginary example to understand their architecture:

Investors supply $100, and maturity for the note is 5 years from now. The bank then issues a 5-year zero coupon bond (a strip), which the bund buys for –say- $70; the remaining $30 assets are then used to pay for all fees first, and then purchase a call option with a $100 notional and strike $70 (equivalently, the fund obtains a non-recourse loan from a bank for another $70, and invests the whole $100 in a fund-of-funds). Upon maturity, the investors will collect the zero ($100), and perhaps some upside if the performance of the funds was in excess of the strike price.

There are good reasons why these structures have been very popular; first, some investors may feel that hedge funds are risky investments, but are attracted to their upside: the bank and the fund-of-funds may feel that the risk is not so high, and are ready to sell the risk at the fund's expense. Secondly, some jurisdictions allow these notes to be sold to any investor, while only qualified investors can invest in fund-of-funds (investors with high net worth).

In practice, these structures can be quite expensive, which will limit the upside of the investors.

7.4 Collateralized fund obligations (CFO)

These work in a similar manner to the leverage structure, except that the financing is done, not by the underwriting bank, but by issuing bonds to outside investors.

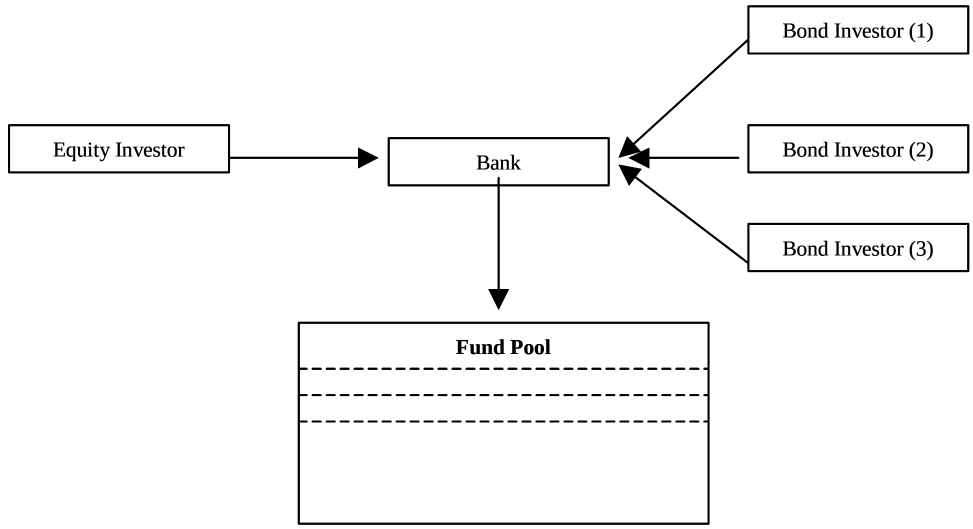

Chart 7.1

Here, bond investors are rated as a function of their seniority: when the structure is unwound, principal is repaid to the bond investors, and coupon payments (if any) are paid as specified in the structure. In the case that, upon maturity, the performance of the fund pool was not adequate, there may not be enough money to fulfill the obligations to the bond-holders; this constitutes default (bankruptcy), in which case senior debt is paid first (from the value of the assets contained in the fund pool), junior debt is paid next, etc. The debt that is paid last is called mezzanine, and each debt level is called a *tranche.* The reader with experience in the credit market will recognize this as the analog of a CDO (Collateralized Debt Obligation) where the debt portfolio here is replaced by the fund pool, with a major difference: CFO's have an equity tranche in addition to the bond tranches; it has less priority than any debt tranche, and picks up all the return of the fund pool after fees, interest rate and coupons are paid.

The interest in these structures is manifold. First, it is attractive to the equity investors (equity or fund share holders) because of the leverage: when the return of the fund pool is greater than the interest payments to the bondholders, they can obtain very attractive return levels; a notional level equal to three or four times their investment is not unusual. It is also attractive to the bond-holders since, for the AAA-rated tranche, it typically provides 50 bps above market yields. Other bond-holders also obtain higher yields than in debt markets, but perhaps most importantly, it offers credit diversification as the default events are assumed not to be correlated to normal credit markets. For the bank, it is a good way to securitize its leveraged structures and sell the investment risk to third parties.

8. Risk Management for Hedge Fund Structures:

The volatility of the underlying hedge fund portfolio continues to be the most important factor in the creation of the structure. However, there are a number of other risk factors that must be taken into consideration when constructing structures. Many of those risk factors are unique to the structure, and not to the underlying portfolio. Others are shared between the structure and the portfolio.

8.1 Case 1: The leveraged structure

Investors are trying to take advantage of low interest rates to leverage hedge fund portfolios. The mental calculation often used is:

- ❖ Expected hedge fund portfolio return: 10%
- ❖ Financing rate: 3-4%
- ❖ Expected return of a 4x leveraged structure: 30%!

The previous calculation, if true, is certainly attractive. What can go wrong?

Sensitivities allow us to understand how changes in current conditions will impact our expectations.

- ❖ How can portfolio volatility hurt? (vega-sensitivity)
- ❖ How can interest rate changes hurt? (rho-sensitivity)

Imagine the following imaginary situation: an investor leveraging 4 times a hedge fund portfolio as before, who expects an annual return of 30%. The portfolio was put together in 2004, avoiding hedge funds with declining returns. This yields an un-noticed correlation to interest rates of –say -50%. The portfolio's annualized volatility is 5% (which is quite low). The leveraged structure carries the usual knock-out option which partially unwinds the structure if the portfolio drops 5%.

If interest rates climb 2% in one year, then there is a 50% probability that: First, annualized portfolio returns drop to 0%!. Second, with small probability (but still significant for the calculation of value-at-risk), on some month you can lose more than 5%, and the structure unwinds.

The conclusion is that structures have risk exposures that we must understand separately of, although perhaps in relationship to, the risks of the underlying hedge fund investments.

8.2 Case 2: CFO correlation sensitivity.

The internal leverage within a CFO makes it more vulnerable to deviations from historical scenarios. Sensitivity analysis is crucial to get a complete risk picture. A key sensitivity that must be used is the increase in distress situations. Moreover, the underlying fund must be designed with objectives in mind to make the CFO work well as a structure, not just the fund portfolio. The rating of the underlying bond issues must also take into account potential deviations from historical means of distress components of future markets (see [Moody's Technical Document]).

In the chart to follow, we present the probability of default of the different components of a mythical CFO; as follows:

- ❖ We used the S&P CTA Index as the underlying fund pool
- ❖ We used a $100 equity investment, with three bond issues of $100 each, called BA (the most senior), BB and BC (the mezzanine).
- ❖ A bond issue defaults if the fund loses, in one year, an amount equal to the equity plus at least as the notional on the bonds with less seniority than it.
- ❖ By extension, we will abuse the language by saying that equity *defaults* when it loses money.

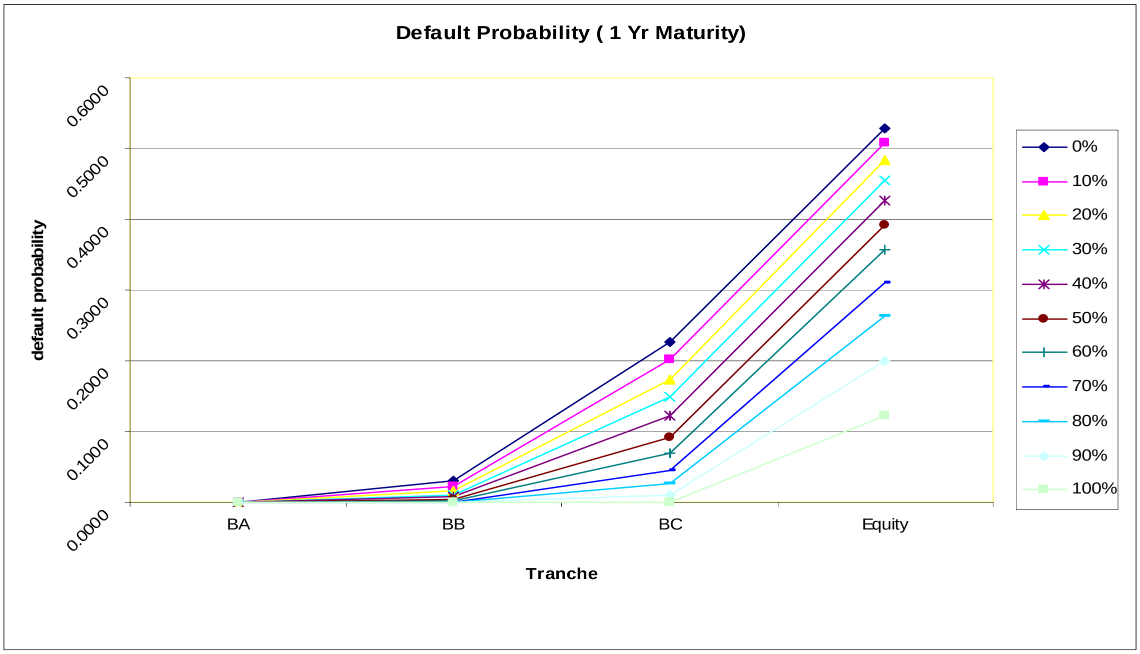

Chart 8.1

What this chart shows is that, if the future deviates from the past by increasing its levels of market distress, we will see a deterioration of the ratings of the bond tranches. As in our previous example, this structure also has it own idiosyncratic risk that must be analyzed separately from the risk of the underlying hedge fund portfolio.

9. Conclusions

Hedge funds are investment businesses that turn risk into return. Their trading strategy, the method to map risk into return, can be very varied and give rise to very different risk profiles. This diversity in their risk profile leads to an internally diversifiable investment class, which although not an asset class, has attracted the attention of investors over the last decade and is establishing itself as a righteous owner of a percentage allocation in every investor's portfolio.

Analyzing risks in hedge funds is not just a good idea: it is a necessary step to understand their return stream. We have seen a rough classification of risks specific to the hedge fund industry, and others that are shared with the traditional investment sector. We must also understand risk sensitivities, how the funds will perform under distress situations, and most importantly, whether their internal diversification holds throughout the worst of times. We have seen that some managers can maintain their diversification benefits and others don't. This must be taken into account when analyzing individual managers risks and returns to make a multifactor portfolio theory, one that looks at risks beyond standard deviation.

We have also looked at investment structures that use hedge funds, which are becoming increasingly popular, and seen that they share some of the risks of hedge fund portfolios, but also introduce new risks of their own. Interest rate risk is one to take seriously into account, especially after the recent history of low financing rates, and correlation breakdown can dramatically affect the default probabilities of CFO tranches.

Acknowledgements

This paper was written with lots of help from the members of Sigma Analysis and Management and the RiskLab of the University of Toronto. Thanks in particular to Ian Buckley, Gustavo Comezaña, Ben Djerroud, Marcos Escobar, Max Reydman, Dave Rudd, David Saunders, Tony Wong and Rafael Zamora.

References

To maintain a certain literary flow, references have been kept to a minimum inside the main body of the paper. As the popularity of this topic may attract the interest of readers not familiar with the hedge fund sector, I have also included in the references a number of books and articles that may be of interest to some readers.

Buckley, I., Saunders, D., Seco, L. *Portfolio optimization when asset sreturns have the Gaussian mixture distribution*, European Journal of Optimization, to appear.

Dembo, R. and Rosen, D. *The practice of portfolio replication*, Annals of Operations Research 85:267-284 (1999)

Dembo, R. *Risk/reward scenario optimization,* Research paper, Algorithmics Inc. (1997)

Fung, W. and Hsieh, D. *Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds*, Review of Financial Studies, 10 (1997), 275-302.

Fung, W. and Hsieh, D. *Survivorship Bias and Investment Style in the Returns of CTAs*, Journal of Portfolio Management, 24 (1997), 30-41.

Fung, W. and Hsieh, D. *Is Mean-Variance Analysis Applicable to Hedge Funds?* Economic Letters, 62 (1999), 53-58.

Fung, W. and Hsieh, D. *A Primer on Hedge Funds* Journal of Empirical Finance, 6 (1999), 309-331.

Fung, W. and Hsieh, D. *Performance Characteristics of Hedge Funds and CTA Funds: Natural Versus Spurious Biases,* Journal of Quantitative and Financial Analysis, 35 (2000), 291-307.

Fung, W. and Hsieh, D. *Measuring the Market Impact of Hedge Funds*, Journal of Empirical Finance, 7 (2000), 1-36.

Fung, W. and Hsieh, D. *The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers* Review of Financial Studies, 14 (2001), 313-341.

Fung, W. and Hsieh, D. *Benchmarks of Hedge Fund Performance: Information Content and Measurement Biases*, Financial Analyst Journal, 58 (2002), 22-34.

Fung, W. and Hsieh, D. *Asset-based Style Factors for Hedge Funds*, Financial Analyst Journal, forthcoming.

Giraud, J. R. *The Management of Hedge Funds' Operational Risks*, EDHEC Risk and asset management research center (2004).

Horwitz, R. *Hedge Fund risk fundamentals*, Bloomberg Press (2004).

Hull, J. *Options, futures and other derivatives*, Prentice Hall (2005).

Kestin, R. and Jaeger, L. *The benefits of alternative investment strategies in the institutional portfolio*, preprint, Swiss Alternative Investment Strategies Group, www.saisGroup.com

Jaeger, L. *Managing Risk in Alternative Investment Strategies*, Prentice Hall (2002).

Jaeger, R. *All about hedge funds*, McGraw-Hill (2003)

Markowitz, H. *Portfolio Selection*, Journal of Finance 7(1):11-91 (1952)

Merton, *Continuous Time Finance,* Blackwell Publishers (1992).

Shadwick, W. and Keating, C. *Omega, a universal performance measure*, Journal of portfolio measurement, Spring (2002).

Sharpe, W. *Adjusting for Risk in Portfolio Performance Measurement,* Journal of Portfolio Management, Winter 1975, pp. 29-34.

Sortino, F. and van der Meer, R. *Downside risk – capturing what's at stake in investment situations*, Journal of Portfolio Risk, 17 (1991).

Witt, G., Fu, Y. and Leahy, J. *Moody's Approach to Rating Collateralized Funds of Hedge Fund Obligation,* www.moodys.com