

Performance and Risk Measurement Challenges For Hedge Funds: Empirical Considerations

Peter Blum¹, Michel M Dacorogna² and Lars Jaeger³

1. Risk and Risk Measures

Complexity and rapid change have made risk management one of the most challenging subjects within the hedge fund industry. Considered separately, both hedge funds and financial risk management have evolved dramatically in recent years, offering tools and strategies of increasing complexity. However, in combination the complexity multiplies. This chapter elaborates on some of the quantitative pitfalls and challenges of risk management as they apply to hedge funds. The reader should note, however, that there are hardly any particular risk management techniques that uniquely suit hedge funds. On the contrary, since hedge funds are part of the overall financial market and trade in well-known asset classes, we should apply to them the same measurement methods that are used by risk managers dealing with other types of investment strategies and instruments. The real challenge lies in the wide spectrum of instruments employed by hedge fund managers, as well as the dynamic nature of their trading strategies and the sparseness of reliable performance data. It is our intention in this chapter to show how advanced risk analysis techniques can be meaningfully applied to hedge funds. We shall thereby elaborate on some finer mathematical details, but try also to explain the concepts in clear words for the less mathematically trained reader.

After introducing various alternative measures of risk, we describe how they are typically evaluated in a portfolio of financial assets. In Section 3 we focus on the question of how the risk of various hedge fund strategies can be described in extreme situations, typically referred to as “tail events”. The challenge of risk assessment in extreme situations becomes even more challenging if considered in a portfolio context. It is well known by academics and practitioners that correlations between financial assets are not stable, and behave notably differently in periods of market distress. This is of ultimate importance to risk managers. Section 4 introduces some new concepts for measuring dependencies as they apply to hedge funds, and the final chapter touches on the issue of time dependencies and risk clustering for hedge funds.

Before diving into the specifics, we would like to define risk. Unfortunately the terms “risk” and “risk management” are used today in a variety of different contexts, adding to investors’ confusion. Within the realm of the financial markets, risk describes the uncertainty of the future outcome of a current decision or situation. Different outcomes materialize with different probabilities, and the range of all possible outcomes and their probability of happening are mathematically described by a probability distribution. The uncertain outcome is typically described by some random variable X , the probability distribution of which is specified through its cumulative distribution function $F(x)$, interpreted as

$$F(x) = \Pr(X \leq x)$$

¹ ETH, Department of Mathematics, 8092 Zürich, Switzerland and Converium Ltd,
Email: peter.blum@converium.com

² Converium Ltd, General Guisan Quai 26, 8022 Zürich, Switzerland,
Email: michel.dacorogna@converium.com

³ Partners Group, Zugerstrasse 57, 6341 Baar-Zug, Switzerland,
Email: lars.jaeger@partnersgroup.net

i.e. for each possible outcome x , $F(x)$ gives the probability of the actual outcome, X being smaller than or equal to x . Quantitative risk management concentrates on analyzing the probability distribution of the various risk factor returns.

It is generally not practicable to analyze the full distribution function $F(x)$. Rather, $F(x)$ has to be boiled down to a few meaningful numbers. To this end, various statistical measures have been developed. Most commonly known are the mean, the median, and the standard deviation. The latter is typically used to assess the dispersion of outcomes around the mean, which is considered as the most basic measure of risk, and is usually equated to the volatility of the market. However, the standard deviation assumes that return realizations are symmetrically distributed around the mean. Further, it measures only the average deviation from the mean, and fails to capture the extreme risks in the tails of the distribution, i.e. the values of $F(x)$ for values of x which lie far away from the mean.

As a risk manager, one is particularly interested in the most adverse outcomes. In this context, one has to concentrate on the extreme quantiles of the distribution. For this purpose one usually considers the 1% or 0.4% quantile. This quantity is called Value-at-Risk (VaR), and describes the maximum loss that cannot be exceeded in 99% or 99.6% of all cases. The main reason for the success of the VaR concept in recent years is the fact that it translates a function as complex as the probability distribution in a monetary amount easily understandable by non-technical people.

VaR is today the most widely used quantitative analysis tool in the financial risk management community. The magic of VaR is that it introduces a uniform measuring system for various risks in a global portfolio by providing a method for comparing risk across different instruments and asset classes. The two important parameters that a risk manager has to define for VaR are the time period and the confidence level (e.g. 99% for a one-day horizon).

Table 1: Empirical estimate of various risk measures for a set of financial instruments. The estimate is based on the logarithmic returns of 10 years of daily prices⁴ (1.1.1993 – 31.12.2002). The data is ordered by increasing volatility (standard deviation).

	Standard Deviation	VaR(99%)	ES(98.75%)	VaR(99.6%)	ES(99%)
MSCI World Sovereign Index	0.37%	-0.90%	-1.08%	-1.10%	-1.27%
Foreign Exchange (USD/GBP)	0.52%	-1.37%	-1.66%	-1.78%	-1.75%
Daily Hedge Fund Index	0.77%	-1.95%	-2.32%	-2.41%	-2.44%
Dow Jones Industrial	1.08%	-2.92%	-3.77%	-3.76%	-4.04%
Brazilian Stock Index	2.96%	-7.85%	-9.59%	-10.13%	-10.20%

One of the drawbacks of VaR is that it measures only one specified point on the distribution function, thereby neglecting valuable information about the rest of the distribution. In particular, one is interested in what happens in the 1% or 0.4% of cases when the loss exceeds the VaR value; in other words one wants to know: "How bad is bad?" A risk measure which addresses this issue is Expected Shortfall (ES, also known as Conditional VaR or Tail VaR). ES quantifies the risk of extreme loss in those events which exceed the threshold given by VaR. It is simply the average outcome of X given that X is beyond VaR:

$$ES(\alpha) = E[X|X \geq VaR(\alpha)],$$

⁴ The hedge fund index is a collection of managed futures and long/short equity funds that provide daily NAV's; it was made available by PartnersGroup. Data for the other instruments comes from Bloomberg Data License.

(the symbol E denotes the mathematical expectation operator). ES is numerically more stable than VaR because it is the result of an averaging procedure, and does not rely only on one point on the probability distribution.

Another drawback of VaR is that it can have bad aggregation properties with respect to sub-portfolios. In certain situations the VaR of a portfolio consisting of two sub-portfolios can be larger than the sum of the single VaR values of the two sub-portfolios. This would neglect the benefits of diversification. In other words, VaR itself does not qualify as a “coherent risk measure” (see Artzner *et al.* 1999).

The goal is to have a risk measure that has the following desirable properties (coherence):

1. Scalability (twice the risk should give a double measure),
2. Correct ranking of risks (bigger risks get bigger measures),
3. Accounting for diversification (aggregated risks should have a lower measure).

It can be shown that, on top of the other desirable properties described, Expected Shortfall also qualifies as a coherent measure of risk.

In Table 1 we show values for the various described risk measures estimated empirically on the logarithmic returns of 10 years of daily prices of various asset classes. One can observe that the expected shortfall gives a more conservative estimate of the risk than the VaR. The threshold at which the risk is measured is important. If one chooses to be less conservative but retain the advantages of a coherent measure, the Expected Shortfall can be used with a lower threshold, say 98.75%.

Data Issues

Quantitative risk management requires the availability of data of sufficient quality and quantity. In its absence, meaningful statistical analysis and the calibration of mathematical models become impossible. This requirement for data is, to some extent, in conflict with the nature of many hedge funds:

- Hedge funds are not subject to frequent (e.g. daily) mark-to-market valuations. Although many hedge funds trade instruments that can be and are individually marked to market on a daily basis (stocks, bonds, futures, options, FX, etc.), regular NAV calculations on the fund level are usually not performed, as redemption periods and the accounting practices of hedge funds make them unnecessary. Often insufficient back-office resources and unwillingness by the manager prevents more frequent valuations of hedge funds. Secondary markets for closed-end hedge funds – to the extent that they exist at all – are usually not very liquid, so frequent quotes are not available.
- For legitimate and less legitimate reasons many hedge fund managers are reluctant to disclose any information beyond the absolutely essential.
- Many hedge funds have not been in place for long.
- Hedge fund index data is subject to a number of biases (e.g. survivorship or selection bias) that are difficult to deal with, and which can seriously distort the meaningfulness of performance indices⁵.

⁵ The properties of and problems with hedge fund (index) data are described in depth in Lhabitant 2002 and Jaeger 2002. The latter author makes a particular case for the requirements of institutional investors.

For these reasons hedge fund performance data is generally limited in quantity, and all too often in quality as well. Performance is generally reported on only a monthly or even quarterly basis, and histories are seldom longer than a few years. The hedge fund index included in Table 1, for which ten years of daily data are considered, is therefore the exception rather than the rule. In the sequel of this text, we will use data from the CSFB/Tremont hedge fund indices, for which there have been monthly returns quotes since January 1994. Our choice was motivated, among other things, by the fact that these indices are asset-weighted and investible. Results for similar indices (e.g. the HFR) are not different in any meaningful way.

The number of available data points (e.g. nine years of monthly data = 108 data points for the Tremont indices) poses some restrictions on the selection of quantitative methods, since many of these methods require large amounts of data in order to lead to significant results. The methods presented here were chosen with these limitations in mind. Another drawback is that, due to the non-existing or rather illiquid secondary markets for hedge funds, there are basically no traded hedge fund derivatives, so one cannot derive return or volatility forecasts based on implied market views, as is common practice in traditional markets.

2. Evaluation of Risk Measures

While VaR and even expected shortfall are conceptually and intuitively quite simple, the technical details of their calculation can be rather involved, depending on the heterogeneity of the portfolio and the distributional assumptions made. Here follows a short overview that applies to all risk measures.

The three main elements of standard risk measure calculation are:

- Mapping the portfolio positions to risk factors.
- Calculating the risk factor covariance matrix based on historical prices.
- Determining the model for the risk and calculating the risk measure.

Mapping Risk Factors

Risk factor mapping is equivalent to decomposing individual securities in a portfolio, and categorizing the risks they present into those over which the risk manager has control (exposure) and components that are exogenous (risk factors). Prices of the thousands of financial securities available worldwide are influenced by common risk factors. A risk factor is a variable that directly affects the value of a security. Individual securities are usually influenced by a variety of different risk factors. Examples of risk factors are interest rates, the broad equity market as represented by stock indices, FX rates, and credit spread curves.

The dependency of a security on a specific risk factor (the security's "sensitivity") is expressed in a "pricing function" and can take a variety of different forms (linear as well as non-linear). The pricing function usually depends on the particular VaR method used (see below). The concept of risk factors and sensitivities is not a new one. Sensitivities to certain factors have been used for many years for the management of yield curve risks, as well as in option theory. Examples are the "value of a basis point move" or the "Greeks". A simple (here linear) model of this type, which is often referred to as "Factor Pricing Model", could be formalized as:

$$R_t = \alpha + \sum_{n=1}^N \beta_n \cdot F_{n,t} + \varepsilon_t$$

Here, R_t is the return of the financial asset, α is a deterministic component (i.e. the risk-free rate of return), $F_{n,t}$ is random variables representing the outcomes of observable risk factors that affect the market as a whole (e.g. equity market portfolio returns, interest rates,

inflation), β_i is the sensitivity of R_t to $F_{n,t}$, and ε_i is the idiosyncratic risk of R_t , i.e. the specific risk not attributable to one of the risk factors. There are various degrees of freedom in the structure of such a pricing model: the number of risk factors can vary, and the dependence on the risk factors can be linear (as in the example) or non-linear (which appears particularly sensible in the realm of hedge funds). The pricing functions have to be correctly adjusted by calibrating the model parameters to the market prices of specific benchmark instruments.

Moreover, different kinds of constraints may apply to the parameters. If these constraints are aimed at eliminating arbitrage, then we are in the realm of the well-known Arbitrage Pricing Theory. Once such risk factor mapping is in place, there still remains the task of defining a suitable stochastic model for the risk factors to be able to estimate risk measures for the various assets R_t and the portfolio in total. This topic is covered in the following sections.

Calculation of the Covariance Matrix

The next step in calculating the portfolio risk consists of estimating the dependence between the risk factors. If one makes the standard assumption that security returns are multivariate normally distributed, the correlation matrix (together with the vector of expected returns) uniquely determines the distribution of portfolio returns, and therefore any risk measure for the portfolio. We briefly present the usual method of estimating the covariance matrix, but we shall see in the next section that the assumptions made thereby are very restrictive and should not be applied without careful examination of their validity, especially in the cases of extreme outcomes. There are various methods for calculating volatility (variances) and correlations (co-variances). The following variance and co-variance models are most commonly used (the RiskMetrics Technical document (1996) is still an excellent source for details⁶).

- Equally weighted moving average of squared returns and cross products over some specified time horizon.
- Exponential moving average of squared returns and cross products with specified decay parameter. This is the method chosen by RiskMetrics and many others. The decay parameter chosen is 0.94 for daily observations and 0.97 for monthly observations.
- The family of ARCH models: ARCH (autoregressive conditional heteroskedastic) and GARCH (General ARCH) models that were developed in the early 1980s have become a starting point for sophisticated volatility and correlation forecasting models⁷. The exponential moving average model can be seen as a special case of a GARCH model.

The main problem in estimating the covariance matrix is that the number of risk factors can be quite large. The higher the dimension of the matrix the more observations are needed to estimate the matrix to a sufficient degree of accuracy. Practitioners often do not consider this point enough, which can affect considerably the stability of the results. Another limitation of the covariance matrix is the fact that the dependence between risk factors is usually neither linear nor static. It is a known fact among practitioners that the dependence may increase during crises (see, for instance, the study of Hauksson *et al.* 2001 of the foreign exchange market).

⁶ For more information and a discussion of the quality of different volatility and correlation forecasts, the reader is referred to Carol Alexander's book: *Market Models; A Guide to Financial Data Analysis* (2001).

⁷ See the article by Tim Bollerslev *et al.*: "Arch modelling in finance," *Journal of Econometrics*, 52, p.5-59 (1992).

Calculation of Risk Measures

After attributing cash flows to risk factors and calculating the risk factors' co-variance matrix, risk measures can be calculated using several methods that differ mainly in respect to two factors:

1. Assumptions regarding security valuation as a function of risk factors. A determination must be made as to the sensitivity of security prices to changes in risk factors. The industry distinguishes between local valuation and full valuation methods.
2. The distributional assumptions made. The industry distinguishes between parametric (mostly Gaussian) distributions versus historical distributions. In Table 1, we estimated the risk measures from the historical distributions. We note that, with the use of a co-variance matrix, one implicitly makes the assumption that the multi-variate distribution of the risk factor set is elliptical (see discussion below).

The three most popular methods to compute risk measures (including VaR or Expected Shortfall) are:

1. Parametric approach. The idea behind parametric models is to approximate the pricing function of every instrument, i.e. the relationship between each instrument and the risk factors, in a way that an analytical formula for the risk measures can be obtained. The simplest parametric approach is the delta method, also referred to as "variance-covariance based VaR"⁸.
2. Monte Carlo simulations⁹. The Monte Carlo method is a full valuation method based on simulating the behaviour of the underlying risk factors through a large number of draws produced by a random generator. Using given pricing functions, the values of the portfolio positions are calculated from the simulated values of each risk factor. The method accounts fully for any non-linearity of the relationship between instrument and risk factors, as the positions in the portfolio are fully re-valued under each of the random scenarios. Every random draw of risk factor values leads to a new portfolio valuation. A high number of iterations (several thousand) provide a simulated return distribution of the portfolio, from which the VaR or the Expected Shortfall values can be determined (a number of numerical techniques exist for making simulation more efficient and more accurate). The underlying distribution of the randomly generated values of the risk factors can essentially be chosen freely, and – in particular – one is not constrained to the Gaussian distribution since analytical tractability does not matter in the Monte Carlo set up. However, the use of non-Gaussian distributions can involve some mathematical problems¹⁰ in terms of simulating the dependency structures of the risk factors correctly, as outlined in Section 4 below.
3. Historical simulation. Instead of simulating return distributions, they can be determined by looking into the past. The historical method is also a full valuation method and relies on the (unconditional) historical distribution of returns by applying past asset returns to the present holdings in the portfolio. The values of the portfolio positions are then fully evaluated for each set of historical returns. This method has the advantage that no explicit assumptions about the underlying return distribution have to be made. However, the problem with this method is that it relies on historical price behavior, which may not be relevant in the current conditions, or may form a too small sample to assess all the

⁸ This method is often called the "RiskMetrics VaR" method, as this was the original method introduced by RiskMetrics in the early 1990s. However, today RiskMetrics also offers other parametric VaR approaches as well as Monte Carlo based and historical simulation VaR.

⁹ In "Beyond Value of Risk" (chapter 5) Kevin Dowd presents a good discussion about the different aspects of Monte Carlo approaches.

¹⁰ See the excellent paper by P. Embrechts et al., "Correlation and Dependence in Risk Management; Properties and Pitfalls".

possible outcomes, particularly in the tails of the distribution. Moreover, this method is not sensible (and workarounds are difficult) if significant inter-temporal dependence is present in the returns.

Even in the case of the Monte Carlo simulation, it is still commonplace among practitioners to assume that risk factor returns follow a normal ("Gaussian") distribution with the mean and standard deviation observed in the historical data¹¹. While this is convenient from a computational point of view, it bears the danger of underestimating extreme risks and related tail-based risk measures as Value-at-Risk. As long as normally distributed draws are used (as is often the case in practice), even the Monte Carlo method does not address the issue of abnormally distributed asset returns. However, the use of Monte Carlo simulations is a starting point for modeling fat-tails and non-linear dependencies, and is thus gaining favor among practitioners willing to go beyond the Gaussian model¹². The dependence and correlation structure of non-normal multivariate distributions is still the subject of intense mathematical research.

In Table 2 we compare the VaR estimates obtained from the full distribution of historical returns ("empirical") with VaR estimates for the same risk levels obtained by making the assumption of normality. We can clearly see that the Gaussian model systematically underestimates the actual risk. Even worse, the degree of underestimation becomes higher the further out we stretch in to the tails. Section 3 of this chapter will treat the problem of properly modeling extreme events in more detail.

Table 2: Value-at-Risk computed with the empirical distribution and under the assumption of a Gaussian distribution (sample information).

	VaR 99% (Gaussian)	VaR 99% (Empirical)	VaR 99.6% (Gaussian)	VaR 99.6% (Empirical)
MSCI World Sovereign Index	-0.85%	-0.90%	-0.97%	-1.10%
Foreign Exchange (USD/GBP)	-1.21%	-1.37%	-1.38%	-1.78%
Daily Hedge Fund Index	-1.71%	-1.95%	-1.96%	-2.41%
Dow Jones Industrial	-2.48%	-2.92%	-2.83%	-3.76%
Brazilian Stock Index	-6.59%	-7.85%	-7.56%	-10.13%

3. Risk in Extreme Situations

One of the crucial points for a risk manager is to quantify risks in periods of tension or crisis on the global capital markets. It is precisely during these periods when risk management should prove its value. Hedge funds have shown in the past that they can serve as valuable portfolio diversifying investments in times of market crises, but equally have exacerbated losses in certain instances (e.g. during the crisis of summer 1998). In this section, we want to present some techniques that have the potential to help assess extreme risks and see whether hedge funds differentiate themselves from the crowd behavior of financial assets in distressed capital markets. There are two elements that need to be examined: the extreme risks of single hedge fund strategies (i.e. amount of probability that is in the tails of the distribution) and a possible change in the dependence between hedge fund risks and other investments during periods of market turmoil.

To assess purely the presence of fat tails, one can simply estimate the kurtosis of the logarithmic returns of a time series empirically, using returns measured at different time horizons, and compare the values obtained for the performance of hedge funds with those

¹¹ Notice that the Gaussian distribution is fully determined by its mean and standard deviation.

¹² The insurance industry, increasingly confronted with the problem of non-linear dependencies and fat-tailed distributions, has developed a Monte Carlo method called Dynamic Financial Analysis (DFA) in order to cope with this problem. For an introduction to this technique, see Blum and Dacorogna 2003.

from the underlying markets traded. This way one can obtain a first indication of whether hedge funds have been able to cope with extreme risks. The kurtosis is related to the fourth moment of the distribution, and can be estimated empirically from historical data. The convergence of the fourth moment is not guaranteed for financial data (see Dacorogna et al. 2001), but this does not prevent us from computing this quantity in order to obtain an initial idea of how heavy-tailed the distribution can be.

Table 3: Basic descriptive statistics for some hedge fund indices compared to classical financial markets (monthly logarithmic returns, January 1994 – December 2002).

Financial Instruments	Mean (μ)	Standard Deviation (σ)	Skewness	Excess Kurtosis
Tremont Hedge Fund Index	0.89%	2.56%	0.10	1.39
Tremont Convertible Arbitrage	0.82%	1.40%	-1.62	4.12
Tremont Dedicated Short Bias	0.20%	5.31%	0.84	1.96
Tremont Emerging Markets	0.68%	5.26%	-0.54	3.67
Tremont Equity Market Neutral	0.84%	0.89%	0.12	0.18
Tremont Event Driven	0.85%	1.81%	-3.32	21.18
Tremont Fixed Income Arbitrage	0.60%	1.35%	-1.14	13.94
Tremont Global Macro	1.17%	3.67%	-0.02	1.59
Tremont Long/Short Equity	0.97%	3.32%	0.24	2.91
Tremont Managed Futures	0.57%	3.46%	0.04	0.84
MSCI World Equity Index	0.35%	4.29%	-0.59	0.35
MSCI European Equity Index	0.37%	5.46%	-1.24	4.25
S&P 500	0.70%	4.68%	-0.58	0.17
Lehman US Bond Index	0.74%	2.43%	-0.13	0.20
SSB Bond Index	0.50%	1.83%	0.44	0.47


The kurtosis measures the amount of probability mass that is concentrated in the tails. Empirically the excess kurtosis (excess compared to the normal distribution whose kurtosis is three) can be estimated as follows:

$$Kurt = \left\{ \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^n \left(\frac{x_i - \mu}{\sigma} \right)^4 \right\} - \frac{3(n-1)^2}{(n-2)(n-3)}$$

where μ is the mean and σ the standard deviation of the data. Positive values of the excess kurtosis signal a heavy tailed distribution.

In Table 3 we report empirical estimates of the first four moments of the return distribution of various hedge fund strategies and traditional markets. We have added to the kurtosis the skewness that gives us the degree of asymmetry of the distribution. The equity indices are known to exhibit a negative skewness, which is clearly visible in the table for the MSCI and S&P 500 indices. It is interesting to note that one can learn a lot from such simple statistics. Some of the hedge fund indices in the table present considerably higher kurtosis than stock or bond indices, while their values for the standard deviations are lower than or equal to the other indices. We note that the computation of basic moments provides an idea of whether conventional portfolio optimization techniques can be applied. These techniques require that the returns do not exhibit significant skewness or kurtosis (see Lhabitant 2002). These assumptions are obviously violated, as shown in Table 3, e.g. for the Tremont Convertible Arbitrage Index.

Figure 1: Comparison of empirical probability densities of monthly returns of two hedge fund indices (on the left) and of one hedge fund index and a stock index (on the right) (sample information).



To illustrate further the message given by the values computed in Table 3 we draw, in Figure 1, the full empirical probability densities of monthly returns for two of the hedge fund indices, Convertible Arbitrage and Equity Market Neutral, and compare them with the MSCI World Equity Index. One can see clearly the asymmetry of the distribution and, on the left, the fat left tail for the Convertible Arbitrage index. One can further see on the right plot that the distribution for the MSCI is much wider than for the Equity Market Neutral strategy.

Tail Analysis

The analysis we proposed in the previous section looks at the entire probability distribution of the returns. In risk management one is particularly interested in the probability of large adverse movements. As long as we are only interested in the extreme events, we do not need to consider models that cover the full range of possible outcomes, as the Gaussian model does. Instead we can restrict our attention to dedicated methods for the analysis of extreme events, i.e. the analysis of the tails of the probability distribution, which can be described by a function $F(x)$ where x takes only values greater than some specified threshold u . Powerful methods for this tail analysis come from the realm of Extreme Value Theory (EVT), which, in recent years, has become fairly popular in various areas of quantitative risk management. McNeil 1999 provides a concise and solid introduction to the use of EVT in risk management, whereas Embrechts *et al.* 1999 constitutes a comprehensive reference.

One of the fundamental theorems of EVT states that for a broad class of probability distributions $F(x)$ (including most of those popular in finance), the tail behavior above a sufficiently high threshold u falls into one of three classes:

1. *Fréchet* or *heavy-tailed* class: the tail of $1 - F(x)$ is essentially proportional to a power function $x^{-\alpha}$ for some $\alpha > 0$. This means that the tail decays slowly and the probability of extreme outcomes is relatively high. The parameter α that essentially governs the tail behavior is called tail index. The closer α is to 0, the higher is the tendency for extreme outcomes.
2. *Gumbel* or *thin-tailed* class: the tail $1 - F(x)$ is essential proportional to an exponential function. This means that the tail decays quickly and the probability of extreme

outcomes is relatively low, though not nil. The thin-tailed case corresponds to the limit of the heavy-tailed case for the tail index α tending to infinity.

3. *Weibull or short-tailed* class: the tail is zero above some finite endpoint. If it comes to modeling the returns from financial assets, truncated distributions are generally not considered, since one cannot define a reasonable endpoint.

Notice that there is a linkage between the definition of "heavy tails" in terms of the kurtosis and in terms of the tail index α . Another important theorem of EVT says that if a distribution has tail index α , then the n^{th} moment of the distribution is infinite for $n > \alpha$. So, if a distribution has, say, $\alpha = 3.5$ (a value quite often seen in returns distributions of financial assets), then its skewness (3rd moment) is finite, whereas its kurtosis (4th moment) does not converge to a finite value. If, on the other hand, we measure a very high excess kurtosis in a sample of returns, we can take this as an indication that a heavy tail is present.

Instead of investigating the tail of $F(x)$ itself, one can also investigate the excess distribution of the return variable X above the threshold u . This is the conditional distribution of $X - u$ given that X is greater than u , i.e.

$$F_u(y) = \Pr(X - u \leq y \mid X > u)$$

The original distribution $F(x)$ for $x \geq u$ can then be recovered via:

$$F(x) = (1 - F(u))F_u(x - u) + F(u)$$

Indeed, yet another main theorem of EVT states that, for some reasonably high threshold, u , $F_u(y)$ can be approximated to deliberate accuracy by the *Generalized Pareto Distribution (GPD)*, which is defined as:

$$G_{\xi, \beta}(y) = \begin{cases} 1 - (1 + \xi y / \beta)^{-1/\xi} & \mid \xi \neq 0 \\ 1 - \exp(-x / \beta) & \mid \xi = 0 \end{cases}$$

While $\beta > 0$ is a mere scale parameter, ξ governs the shape of the distribution. It can be verified from the above definition that the $\xi > 0$ corresponds to the heavy-tailed case, $\xi = 0$ is the thin-tailed case, and $\xi < 0$ the short-tailed case. In the first case, the relation $\xi = 1/\alpha$ holds for the shape parameter which explains the name *tail index* for α . For the reader interested in the finer mathematical details, we note that it is easy to verify that the different values of ξ (resp. α) give rise to the different classes of tail shapes introduced above. Hence, tail analysis essentially boils down to estimating the shape parameter ξ . Assuming the GPD model, a variety of methods is available, including the well known Maximum Likelihood technique. Methods exist for cases in which observations show serial correlations, besides those methods applicable to the basic case of uncorrelated observations (see McNeil 1999 for an overview and Embrechts *et al.* 1999 for full details). As an alternative to this parametric approach, non-parametric approaches to tail index estimation are available. The most popular is known as the Hill estimator (see also Embrechts *et al.* 1999). All these approaches come with confidence intervals for the estimates obtained, allowing the statistician to judge whether some estimated ξ are actually significantly greater than zero, indicating a heavy tail.

Easy though it may look, practical tail estimation suffers from a number of problems. The most basic one is the selection of a reasonable threshold u , on which the estimated tail index is often heavily dependent. Moreover, the amount of data available in the tail is often very low, leading to broad confidence intervals and only slightly significant estimates. The latter

problem applies particularly to the realm of hedge funds, as described above. Practical tail estimation is therefore rarely a straightforward process in practice. It usually involves some trial and error and good judgment. The good news, however, is that powerful tools and algorithms are available today (see again the references stated in this section).

Once we feel sufficiently confident with the estimated tail model, we can use it to estimate tail-related risk measures such as VaR and Expected Shortfall. We demonstrate the case for the GPD model, and assume no serial correlation in the data. Substituting the GPD for the excess distribution $F_u(y)$ in the above representation and recalling the definition of Value-at-Risk, we obtain the following easy-to-compute formula:

$$VaR_q(X) = u + \frac{\beta}{\xi} \left(\left(\frac{n}{N_u} (1-q) \right)^{-\xi} - 1 \right)$$

where n is the total number of observations and N_u is the number of observations exceeding the threshold u and the other parameters are as defined above. Recalling the definition of the Expected Shortfall, we furthermore obtain:

$$ES_q(X) = \frac{VaR_q(X)}{1-\xi} + \frac{\beta - \xi u}{1-\xi}$$

The benefits of estimating tail-related risk measures by using a model instead of just historical data are twofold. First, we can estimate at quantiles q beyond what is covered by available data. Moreover, even if we are within quantile ranges still covered by historical data, applying the model yields smoother estimates, which is a considerable advantage given the low amount of data usually out in the tail.

Table 4: GPD model estimates for Tremont hedge fund indices and traditional market indices (sample information).

	Excess Kurtosis	Shape parameter ξ	90% conf. interval for ξ	VaR 95% empirical	VaR 95% GPD model	VaR 99.6% GPD model
Hedge Fund Index	1.39	-0.2968	[-0.47,-0.15]	6.05%	5.88%	8.44%
Convertible Arbitrage	4.12	0.0828	[-0.17,0.35]	3.24%	2.99%	5.30%
Dedicated Short Bias	1.96	0.2814	[-0.08,0.69]	8.80%	9.37%	18.63%
Emerging Markets	3.67	0.2181	[-0.26,0.70]	9.80%	10.24%	22.14%
Equity Market Neutral	0.18	-0.2606	[-0.43,-0.07]	2.14%	2.38%	3.28%
Event Driven	21.18	0.3105	[-0.10,0.72]	3.09%	3.37%	7.15%
Fixed Income Arbitrage	13.94	0.3759	[0.06,0.69]	2.01%	2.41%	6.32%
Global Macro	1.59	0.1110	[-0.13,0.33]	9.33%	9.84%	15.89%
Long/Short Equity	2.91	0.1735	[-0.16,0.57]	7.07%	6.97%	14.90%
Tremont Managed Futures	0.84	-0.4736	[-0.88,-0.07]	7.85%	7.60%	9.97%
MSCI World Equity Index	0.35	-0.1828	[-0.36,-0.03]	8.51%	8.54%	12.54%
MSCI EU Equity Index	4.25	0.3146	[-0.07,0.70]	10.05%	10.24%	25.28%
S&P 500	0.17	0.0250	[-0.29,0.30]	8.22%	8.64%	13.51%
Lehman US Bond Index	0.20	-0.1083	[-0.37,0.16]	4.95%	4.88%	7.39%
SSB Bond Index	0.47	0.0624	[-0.39,0.40]	3.60%	3.76%	6.39%

In Table 4 we report the results of a tail study based on the basic GPD model introduced above. The database consists of monthly absolute returns of the Tremont hedge fund-style indices and some traditional market indices for the time period from January 1994 to December 2002. This makes 108 data points per index, which is rather few in absolute terms, but a typical situation in view of the short histories and low reporting frequencies prevalent in the hedge fund industry. The second data column reports maximum likelihood estimates for the shape parameter ξ , complemented by related 90%-confidence intervals in the next column. Selection of a suitable threshold employed graphical evaluation techniques

suggested in Embrechts *et al.* 1999, and by repeating the estimation across a range of possible thresholds. We notice that the confidence intervals are rather wide, which is not astonishing given the low amounts of data at hand. For some of the indices we obtained negative values for ξ , indicating that the returns distributions follow short-tailed law¹³. For the other indices, the results suggest values of ξ above zero, i.e. heavy-tailed returns distributions, with particular significance for the Fixed Income Arbitrage index. Notice also that the excess kurtosis estimates correspond to the tail shape parameter estimates; there is only low excess kurtosis for those indices where the estimated ξ suggests thin tails, whereas there is high excess kurtosis in the case of highly positive values of ξ . In the last three columns, we report Value-at-Risk estimates for the 95% and 99.6% confidence levels. For the former we have estimates from both the empirical distribution of the data and the GPD model. We see that the two estimates correspond relatively well with each other, with the GPD-based estimates being generally more conservative, except for the cases where the estimated negative ξ suggests a capped distribution. The 99.6% level is beyond the range covered by data, so we have to rely on the model entirely. Notice also that the series with the heaviest tail (Tremont Fixed Income Arbitrage) has the lowest outcomes on an absolute scale. Indeed, "extreme values" in this context must always be thought of as extreme with respect to the "usual" behavior of the risk factor, and not on an absolute scale. This is further stressed by Table 5, where we restate the VaR estimates as multiples of the standard deviation above the mean, i.e. as the n in the equation $VaR = \mu + n\sigma$, hereby relating VaR to the classical volatility measures and removing the impact of location and scale. We can clearly see that VaR increases much more dramatically between the 95% level (not yet really in the tail) and the 99.6% level (far out in the tail).

Table 5: Value-at-Risk expressed in terms of mean and standard deviation (sample information). A Gaussian factor would be 1.65 for 95% and 2.65 for 99.6%.

	μ	σ	ξ	VaR 95%	VaR 99.6%
Hedge Fund Index	0.89%	2.56%	-0.2968	1.95	2.95
Convertible Arbitrage	0.82%	1.40%	0.0828	1.55	3.20
Dedicated Short Bias	0.20%	5.31%	0.2814	1.73	3.47
Emerging Markets	0.68%	5.26%	0.2181	1.82	4.05
Equity Market Neutral	0.84%	0.89%	-0.2606	1.73	2.74
Event Driven	0.85%	1.81%	0.3105	1.39	3.48
Fixed Income Arbitrage	0.60%	1.35%	0.3759	1.34	4.24
Global Macro	1.17%	3.67%	0.1110	1.82	4.01
Long/Short Equity	0.97%	3.32%	0.1735	1.81	4.20
Managed Futures	0.57%	3.46%	-0.4736	2.03	2.71
MSCI World Equity Index	0.35%	4.29%	-0.1828	1.91	2.84
MSCI Europe Equity Index	0.36%	5.46%	0.3146	1.81	4.56
S&P 500	0.70%	4.68%	0.0250	1.70	2.73
Lehman US Bond Index	0.74%	2.43%	-0.1083	1.70	2.74
SSB Bond Index	0.50%	1.83%	0.0624	1.65	3.09

The above example suggests that instruments with heavy-tailed return distributions reside in both traditional and hedge fund markets, even on time aggregations as high as monthly periods. Prudent risk management requires the ability to test and account for this phenomenon. The approach presented here is an easy-to-implement tool, which is in practice well-proven in settings with low amounts of data and low inter-temporal dependence. For further very useful methods applicable in more involved settings, and related information, we refer to Embrechts *et al.* 1999 and McNeil 1999.

¹³ It is not customary in finance to use capped (i.e. short-tailed) distributions to model return distributions, since there is always a potential – possibly very small – for extreme moves. Given a negative estimate of ξ , one would usually select a thin-tailed model ($\xi=0$), e.g. the Gaussian one.

4. Dependence and Diversification

Contemporaneous Dependence and Correlation

One of the crucial problems a risk manager faces is to measure and assess dependences and correlations – and hence quantify the potential for diversification in the portfolio context. Of particular interest is whether the perceived diversification benefits hold in times of crises in the capital markets. It is a well-known fact among financial practitioners that dependence structures between different asset classes may change in times of extraordinary market stress: dependence increases or even changes in nature. Negative correlation may turn positive, or vice versa. These questions are particularly interesting in the realm of hedge funds, since a main argument in their favor is related to the claim of their superior diversification potential with respect to traditional asset classes. Investors often rely on the same standard methods used for traditional assets to explore the extent to which hedge funds provide diversification in the global portfolio, both in normal markets and in stressed markets.

Here again a good empirical study is hampered by the lack of data, particularly when it comes to assessing dependence in extreme situations. It is already difficult to assess the tail properties of one-dimensional return distribution functions. However, the problem is much more serious when one tries to explore extreme behaviors in a multi-dimensional (portfolio) space. In this section we present some methods to go beyond the standard approach and assess the real diversification potential offered by hedge funds, especially in the tail regions of return distributions.

The standard tool to investigate the dependence structure of hedge funds on traditional asset classes is to estimate their historical correlation to the returns of the traditional markets (see also Section 2). In Table 6 we show the correlation of various hedge fund-style indices with some traditional markets and – for comparison's sake – the correlations of the traditional markets with each other. We see that a number of the hedge fund indices present an appreciable diversification (in terms of low or even negative correlation) when calculated with respect to traditional markets. Those numbers provide a first idea about the potential for diversification that is achievable by investing in hedge funds. But the question remains, how stable and reliable is this dependence in various market circumstances¹⁴?

Table 6: Correlation between hedge fund performance indices and major financial market indices. The numbers in italics are not significantly different from zero at the 95% confidence level.

	msci eu	s&p500	le us bd	ssb bd
msci eu	1.00	0.76	<i>-0.19</i>	<i>-0.01</i>
s&p500	0.76	1.00	<i>-0.05</i>	<i>-0.06</i>
Lehman Brothers us bond Index (le US bd)	<i>-0.19</i>	<i>-0.05</i>	1.00	0.45
Salomon Smith Barney bond Index (ssb bd)	<i>-0.01</i>	<i>-0.06</i>	0.45	1.00
Tremont Hedge Fund Index	0.38	0.47	<i>0.15</i>	<i>-0.17</i>
Tremont Convertible Arbitrage Index	<i>0.07</i>	<i>0.14</i>	<i>0.05</i>	-0.21
Tremont Dedicated Short Bias Index	-0.55	<i>-0.77</i>	<i>0.13</i>	<i>0.09</i>
Tremont Emerging Markets Index	0.41	0.47	<i>-0.15</i>	-0.27
Tremont Equity Market Neutral Index	0.21	0.41	<i>0.05</i>	<i>-0.03</i>
Tremont Event Driven Index	0.48	0.56	<i>-0.08</i>	-0.21
Tremont Fixed Income Arbitrage Index	<i>0.06</i>	<i>-0.03</i>	<i>0.05</i>	-0.21

¹⁴ One illustrative extension of the unconditional correlation estimate shown in Table 6 is the calculation of conditional correlations, e.g. correlations measured only in months of negative (positive) returns of some stock indices.

Tremont Global Macro Index	0.16	0.24	0.26	-0.20
Tremont Long Short Equity Index	0.49	0.58	0.03	0.00
Tremont Managed Futures Index	-0.25	-0.26	0.32	0.34

Dependence in Periods of Crises

Correlation measures the “average” dependence between returns of different instruments. In particular, correlation is a measure of linear dependence, i.e. it assumes that the dependence is constant across all possible return outcomes and in all time periods. As mentioned above, this assumption may not be true if periods of unusual stress are involved. Therefore, correlation cannot take into account miscellaneous anomalies. An additional problem with interpreting the dependence structure solely on the basis of linear correlations is that this can be done only with some (implicit) assumption about the distribution of returns. In particular, the presence of significant skewness or a tendency towards extreme outcomes seriously distorts the meaningfulness of correlations as a measure of dependence. A thorough discussion of the pitfalls associated with correlation, along with some remedies, is presented in Embrechts et al. 2002.

Few tools are available for exploring the effect of market turbulence on the dependence between various asset classes. Often enough, a simple two dimensional scatterplot representing the return of different assets on each dimension provides some indication about the presence of non-linearity and other anomalies in the dependence structure. On the left-hand side of Figure 2, we show two such scatterplots. The lower one indicates a relatively stable dependence structure, whereas the upper one allows for the conjecture that a low number of extreme returns might have a strong impact on the measured dependence.

This relatively simple exploration already gives a feeling for the types of dependence one is confronted with for particular hedge funds, but it is cumbersome to explore many of these graphs, and the criteria for discriminating linear from non-linear behaviors are difficult to formalize. We would like to introduce the reader to a powerful test method to examine the stability of correlation coefficients in periods of stress. This test (see Hauksson *et al.* 2001) is based on the examination of correlation that is conditional on the return outcome of both assets being jointly extreme. In general, two random variables that are linearly correlated are well represented by an elliptical distribution in two dimensions. Elliptical distributions have some nice features, including closure undertaking linear combinations and marginal distributions¹⁵. Elliptical distributions include the well-known multivariate normal distributions, but also some more heavy-tailed ones. While the tails of the normal distribution are too thin for most financial assets, it is still possible for other elliptic distributions to fit financial returns better. Here we shall be mostly concerned with how well elliptic distributions capture the dependence structure in the tails.

Recall that a multivariate random variable X is said to be elliptic if there exists a constant vector μ and a positive definite matrix Σ , such that the random variable $Y = \Sigma^{1/2}(X - \mu)$ is spherically distributed, i.e. its probability distribution is invariant under rotations. The matrix Σ is then a constant multiple of the covariance matrix and μ is the mean, assuming that both exist. Note that, if we define the sets¹⁶:

$$\Omega(\mu, \Sigma) = \left\{ x \mid (x - \mu)\Sigma^{-1}(x - \mu) \geq s \right\}, s \geq 0$$

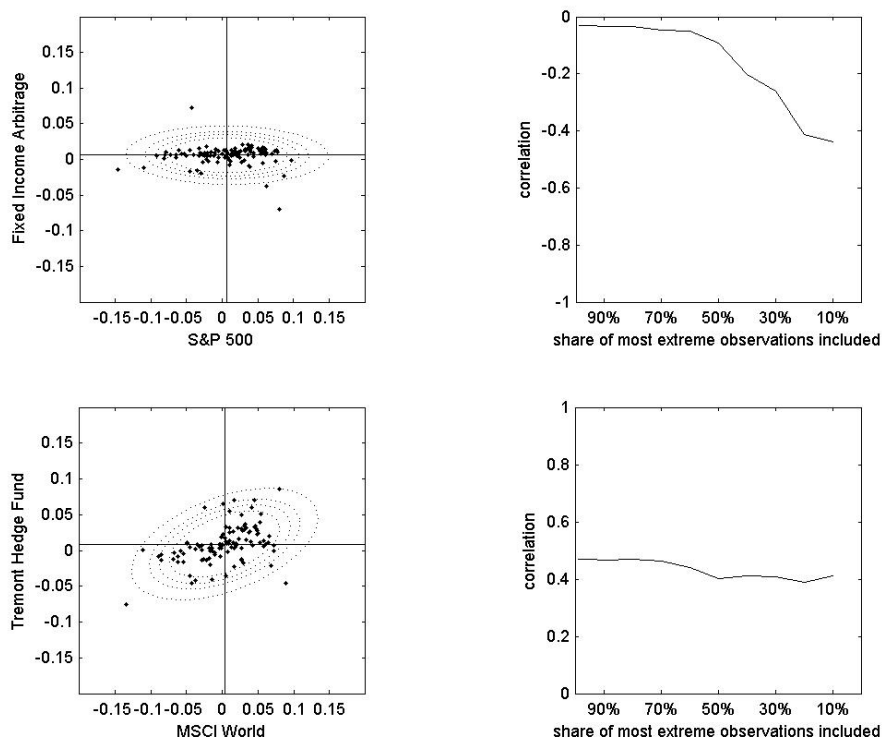
then the conditioned random variable $X|_{\Omega}$ is also elliptic, with the same mean μ and covariance matrix Σ . In particular X and $X|_{\Omega}$ both have the same correlation matrix. Therefore, if we estimate the correlation matrix of $X|_{\Omega}$ as a function of s , it should be

¹⁵ Closure means that a linear combination or a marginal remains an elliptical distribution.

¹⁶ Note that in two dimensions, these sets are ellipses; some of them are drawn in the two panels on the left-hand side of Figure 2.

constant. In practice we fit an elliptic distribution to the entire data set, construct the appropriate ellipses, and then compute the correlations for the data points lying outside of each ellipse (the size s of which is parameterized by the percentage of data points included in the ellipse). The right-hand side of Figure 2 shows the results of these computations for the two pairs of asset classes presented in the scatterplots on the left. For the upper pair we see – as conjectured before by looking at scatterplot – a significant change in the correlation if one takes into account only the more extreme outcomes. This is a clear sign that the simple linear correlation is not a good way of reflecting the dependence structure between these two asset classes. For the lower pair – as for the majority of a large number of other pairs of asset classes and hedge fund returns examined – the correlation is quite stable as a function of increasing s , indicating a stable dependence structure between the variables.

Figure 2: Some examples of correlation under stress. The panels on the left-hand side show the scatterplots, the panels on the right-hand side present the computations of the correlation test introduced in the text.



Results obtained from carrying out this relatively simple exploration on a number of asset pairs show that the dependence between hedge fund indices and the traditional financial markets generally does not vary in stress situations, at least at this monthly level of time aggregation of returns. But we have also seen that there are indeed some cases where the correlation significantly varies out in the tails, and we have managed to quantify these changes in correlation by using our simple conditional under extreme correlation approach.

To conclude this section, we briefly touch on the methodology of copulas. Copulas represent the most general way of modeling dependence between random variables, and they are able to cater for a wide range of dependence structures, including but not limited to linear dependence. The copula methodology has recently received a lot of attention in financial risk management; (see the seminal contribution by Embrechts et al. 2002 for more information). Let us assume that we are given a vector $\mathbf{X} = (X_1, \dots, X_n)'$ of dependent risk factors; the joint distribution function of this vector is:

$$F(x_1, \dots, x_n) = \Pr(X_1 \leq x_1, \dots, X_n \leq x_n)$$

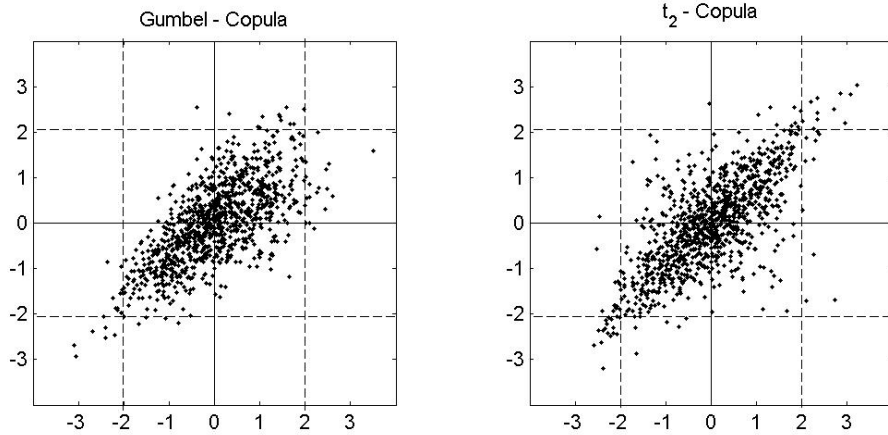
It can then be proven that, under fairly mild conditions, there exists a representation of the following kind:

$$F(x_1, \dots, x_n) = C(F_1(x_1), \dots, F_n(x_n)),$$

where the $F_i(x_i)$'s are the marginal distribution functions of the individual risk factors, and $C(u_1, \dots, u_n)$ is a uniquely defined function, the copula function. The beauty of this representation lies in the fact that the marginal behavior of the individual risk factors is separated from their dependence structure. This allows one to combine freely models for the single risk factors with those for the dependence structure, without the very restrictive assumptions imposed by linear correlation (i.e. the probability distribution being elliptic). This is a particular advantage in the realm of hedge funds where returns are – as we have seen – often significantly skewed and heavy-tailed, which causes linear correlation in many cases to be meaningful only to a limited extent. Moreover, the use of the copula method allows us to (stress-) test for a wide variety of different dependence structures.

Although we have not observed many departures from linear dependence in the study presented above, this does not mean “we have seen it all”. Our data sample might simply be too short to exhibit breaks in dependence. Thus, we rely on some simulated data to show the consequences of non-linear dependence, i.e. to illustrate to the reader the phenomenon of dependence beyond linear correlation. Many copula models are suitable for use in risk management (see e.g. Embrechts *et al.* 2002 or Kotz and Nadarajah 1999 for more detailed discussions). In Figure 3 we show data simulated with a Gumbel copula and with a t_2 -copula, where all marginal distributions are standard Gaussian. The t_2 -copula produces outcomes not dramatically different from linear correlation, except for a slightly increased tendency of marginally extreme outcomes to occur simultaneously. The Gumbel copula, however, shows a dependence pattern radically different from linear correlation: while the points to the upper right appear fairly uncorrelated, the points to the lower left show an increasing correlation the farther out one goes in the tail (“teardrop shape”).

Figure 3: Simulated bi-variate observations with a Gumbel and a t_2 copula; the marginal distributions are standard normal in both cases.



5. Time Dependence and Volatility Clustering

The previous section has dealt with the analysis of contemporaneous dependence between asset classes. For multi-factor risk management and for portfolio optimization, this is clearly the main dimension for dependence analysis, also in view of the high correlations that are actually present between major asset classes. Nevertheless, one should not neglect the dependence of asset returns over time. Mainstream thinking in finance generally asserts that asset returns on time aggregations as high as monthly, quarterly or yearly tend to be uncorrelated over time. There are, however, many statistical studies showing that this assertion does not always hold (see e.g. Campbell et al. 1997). Now, if independence over time does not even hold in general for highly liquid traditional markets, one should be even more suspicious of such complex investments as hedge funds. To explore time dependence, we rely on a study of the autocorrelation of the returns. It is a convenient way to discover time dependency. The autocorrelation function examines whether there is a linear dependency (i.e. correlation) between the current and past values of a variable X_t . It can be estimated by:

$$\rho_x(\tau) = \frac{\sum (x_{t-\tau} - \bar{x})(x_t - \bar{x})}{\sqrt{\sum (x_{t-\tau} - \bar{x})^2 \sum (x_t - \bar{x})^2}}$$

where x_t are observations from a stationary time series and τ is the time lag of interest. In Table 7 we show autocorrelation estimates at different time lags for monthly plain and absolute logarithmic returns of a selection of hedge fund indices and traditional markets. Looking first at the autocorrelations of plain logarithmic returns, we can see a fairly varied picture. For some of the indices, there is, indeed, significant autocorrelation, whereas for others there is none. The conclusion is that we must not blindly believe in the absence of time dependence, and some study thereof must be part of a prudent risk management effort. We also want to note that a positive conclusion to reach from the presence of significant autocorrelation is that – by using a time series model – we can to some extent forecast the next period's return by using past observations.

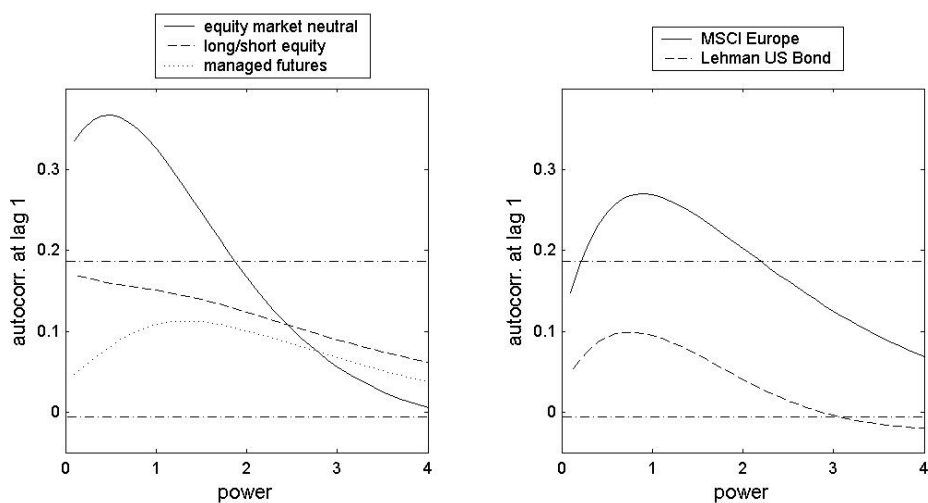
Table 7: Analysis of the autocorrelation coefficients of the monthly logarithmic returns ("plain") and their absolute value ("abs") of some traditional markets and hedge fund indices (sample information). Only values in italic typeset are significantly different from zero.

		Time lag τ				
		1 month	2 months	3 months	4 months	5 months
MSCI World Equity Index	plain	-0.04	-0.04	0.10	-0.11	0.02
	abs.	0.04	0.17	0.07	-0.02	0.00
MSCI Europe Equity Index	plain	-0.07	0.06	0.16	-0.07	-0.03
	abs.	<i>0.27</i>	<i>0.27</i>	<i>0.26</i>	0.08	0.13
S&P 500	plain	-0.04	-0.02	0.12	-0.09	0.07
	abs.	0.02	<i>0.19</i>	0.12	-0.10	0.07
Lehman US Bond Index	plain	0.15	-0.10	0.06	0.01	-0.05
	abs.	0.09	-0.09	0.06	0.07	0.06
SSB Bond Index	plain	<i>0.23</i>	0.00	0.02	<i>-0.19</i>	<i>-0.29</i>
	abs.	0.15	-0.06	-0.14	0.02	0.09
Tremont Hedge Fund Index	plain	0.11	0.04	-0.02	-0.09	0.06
	abs.	0.05	<i>0.26</i>	0.05	0.09	0.12
Tremont Convertible Arbitrage	plain	<i>0.56</i>	<i>0.42</i>	0.14	0.10	0.06
	abs.	<i>0.38</i>	<i>0.18</i>	-0.04	0.11	0.06
Tremont Dedicated Short Bias	plain	0.06	-0.07	-0.03	-0.11	-0.13
	abs.	-0.04	0.12	0.32	-0.02	0.06
Tremont Emerging Markets	plain	0.30	0.01	-0.02	-0.05	-0.09
	abs.	0.06	-0.10	0.06	-0.01	0.02
Tremont Equity Market Neutral	plain	<i>0.29</i>	<i>0.21</i>	0.10	0.06	0.15
	abs.	<i>0.33</i>	<i>0.30</i>	0.09	0.09	0.07
Tremont Event Driven	plain	<i>0.34</i>	0.15	0.02	-0.01	-0.04
	abs.	0.02	-0.04	-0.01	-0.03	0.08
Tremont Fixed Incom Arbitrage	plain	<i>0.35</i>	0.11	0.07	0.14	0.09
	abs.	<i>0.26</i>	0.14	0.14	0.05	0.01
Tremont Global Macro	plain	0.06	0.05	0.08	-0.09	<i>0.24</i>
	abs.	0.05	0.10	<i>0.18</i>	0.01	<i>0.24</i>
Tremont Long/Short Equity	plain	0.16	0.06	-0.04	-0.08	-0.17
	abs.	0.15	<i>0.37</i>	0.11	<i>0.29</i>	0.10
Tremont Managed Futures	plain	0.03	-0.09	0.02	-0.04	-0.05
	abs.	0.11	-0.17	0.05	0.06	-0.10

Another stylized fact, which is well-documented in daily and higher-frequency data, is the so-called volatility clustering. This means that periods of high volatility are more likely to be followed by periods of high volatility than by periods of low volatility. This of course could have important consequences on the risk of the investment. Volatility clustering can be detected by studying the autocorrelation of absolute returns (see Dacorogna et al. 2001). In Table 7, we show the autocorrelations for the absolute returns of the studied return time series, where we can see that they are significantly different from zero for some, but not all of the cases. This indicates the presence of volatility clustering for some hedge fund strategies, even on a monthly time aggregation. Here again, time dependence in the volatility allows the latter to be forecast to some extent.

Finally, we study the clustering of extreme values. To this end, we compute the autocorrelation of the absolute returns to a certain power ρ . Increasing the power of the absolute returns is equivalent to increasing the relative importance of extreme events in the statistics. In Figure 4 we see that for some of the investments studied in Table 7, the autocorrelation drops rapidly when $\rho > 1$. In other words, extreme events are less correlated with each other than average or small absolute returns. If this rapid decline were not present, we would have to be aware of a clustering of extreme events. From this study, it seems that the volatility clustering is mainly due to the average behavior, not the extreme events. This is consistent with studies on other types of financial assets (see Dacorogna et al. 2001).

Figure 4: Autocorrelations at lag 1 of powers of absolute returns, $|r|^p$ as a function of the exponent p . The horizontal dot-dashed lines are the 95% significance levels.



6. Conclusion

Despite the scarcity of the data sets available to us, it is not too difficult to identify the additional risks of extreme events unconsidered by conventional risk measurement techniques, and thus evaluate the real return and diversification potential of hedge funds in a global investment portfolio. We have presented two approaches to assess this risk, one by computing univariate risk measures for hedge fund strategies by looking at the occurrence of extreme outcomes in their return distributions, and the other by examining the generalized dependence structure of hedge fund strategies with traditional financial assets.

We conclude that, both in terms of their properties in the region of extreme risk as well as in terms of their diversification potential, hedge funds can serve as a fruitful investment in investors' portfolios, but not without proper risk management and assessment of past performance that goes beyond simple linear techniques to apply some methods derived from extreme value theory.

Bibliography

Alexander C., *Market Models; A Guide to Financial Data Analysis*, John Wiley & Sons Ltd, New York (2001).

Artzner P., Delbaen F., Eber J. and Heath D., "Coherent Risk Measures" ", *Mathematical Finance*, vol. **9** (3), pages 203-228 (1999).

Blum P. and Dacorogna M. M., "Dynamic Financial Analysis; Understanding Risk and Value Creation in Insurance", In Depth Brochure, Converium Ltd, (2003)
 [Downloadable from: <http://www.converium.com/>].

Bollerslev T., Chou R. and Kroner K., "ARCH Modelling in Finance", *Journal of Econometrics*, vol. **52**, pages 5-59 (1992).

Campbell J., Lo A., McKinlay C., *The Econometrics of Financial Markets*, Princeton University Press (1997).

Dacorogna M. M., Müller U. A., Pictet O. V. and De Vries C., "The Distribution of Extremal Foreign Exchange Rate Returns in Extremely Large Data Sets", *Extremes*, vol. 4(2), pages 105-127 (2001).

Dowd K., *Beyond Value at Risk*, Wiley Frontiers in Finance, John Wiley & Sons Ltd, New York (1998).

Embrechts P., Klüpperberg C. and Mikosch T., *Modelling Extremal Events for Insurance and Finance*, 2nd edition, Springer Verlag, Berlin (1999).

[See also <http://www.math.ethz.ch/~embrechts/> for additional materials.]

Embrechts P., Mc Neil A. and Straumann D., "Correlation and Dependence in Risk Management; Properties and Pitfalls", in *Risk Management: Value at Risk and Beyond*, edited by Dempster M., Cambridge University Press, pages 176-223 (2002).

Hauksson H. A., Dacorogna M. M., Domenig T., Müller U. A. and Samorodnitsky G., "Multivariate Extremes, Aggregation and Risk Estimation", *Quantitative Finance*, vol. **1**(1), pages 79-95 (2001).

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

Lhabitant F.-S., *Hedge Funds: Myths and Limits*, John Wiley & Sons (2002).

McNeil A., "Extreme Value Theory for Risk Managers", in *Internal Modelling and CAD II*, pages 93-113, RISK Books, 1999.

[See also <http://www.math.ethz.ch/~mcneil/> for additional materials.]

Kotz S., Nadarajah S., *Extreme Value Distributions*, Imperial College Press (1999).

RiskMetrics, "RiskMetrics – Technical Document", (1996)

[Downloadable from <http://www.riskmetrics.com/techdoc.html>].