The Tail that Wags the Hedge Fund Dog

Craig French

May 2008

"Special study of the outliers may prove very rewarding."
-- F.J. Anscombe (1973)

Unconditional and Conditional Regressions

It has become somewhat *de rigueur* in the investment community to adopt the following stance: "If I'm paying 2 and 20 for these hedge fund returns, I don't want the returns to be beta in disguise. I can get beta to risk factors a lot cheaper than that. I only want to pay these steep fees for true alpha." One problem with such a position is it utilizes a static definition of beta. However, betas are dynamic. In this paper we will explore a way to evaluate the dynamic nature of hedge fund risk factor exposures, and we will see that, in some cases, a high fee structure such as 2% management fee plus 20% incentive fee may be well worth the price of hedge funds, both for alpha and for the downside protection which can be provided by the dynamic beta.

There are many risk factors an investor might be interested in; risk system vendors such as BARRA, Northfield and others can provide and analyze dozens and dozens. We might be concerned with equity risk, credit risk, interest rate risk, volatility risk, currency risk, liquidity risk, etc. Normally the way the analyst estimates a portfolio's exposure to such risks is to run ordinary least-squares ("OLS") linear regressions using the periodic (possibly excess over a risk-free rate) returns of some proxy for the risk factor (e.g., the S&P500 Index as a proxy for U.S. equity risk) as the independent variable and the contemporaneous periodic returns of a portfolio asset (e.g., a hedge fund) as the dependent variable; the slope of the best-fit OLS line is interpreted as the risk exposure, or "beta", of the asset to the risk factor.

The regression approach is convenient, since each beta estimate has an associated p-value and t-statistic which can give the analyst a level of confidence with which to assess the beta estimates, higher t-statistics, for example, indicating "better" beta estimates. Also, the alpha estimates, often interpreted in practice as "expected excess return", and their associated confidence indicators, are derived directly from the regression.

When we utilize the entire vector of asset returns, and the contemporaneous vector of risk factor returns in such a regression, I will call this an "unconditional" regression. Unconditional regressions are frequently employed in the asset management industry. When we utilize only a subset of the vector of risk factor returns, I will call this a "conditional" regression. This definition is nonstandard, since "conditional" in econometrics usually refers to conditioning on some additional factor.

Notice that employment of non-contemporaneous (lagged or leading) risk factor returns in the regression, which might be done to account for stale pricing, would make the regression conditional, since it would be conditioned on a time-shift. Also notice that, if we were to apply a return decay scalar to dampen the effect of older observations, a practice employed for example, in the RiskMetrics system, such an approach would also fit my definition of "conditional", because, since less weight is given to some of the observations, the regression is "conditioned" on an *effective* time-shift. Stepwise regression, commonly used for variable selection, is *per se* neither unconditional nor conditional – this type of regression may be run in either setting. However, lags, dampers and variable selection are not the issues we will focus on in this paper; rather, we are going to examine conditioning on the tails of the distribution of risk factor returns.

In order to evaluate hedge fund tail risk, we will employ three methods: linear regressions (using Markowitz terminology, we are tempted to call the approach "piecewise linear" regression), polynomial regressions, and extreme value theory (EVT).

A Portfolio of Insurance Policies

Within the usual conception of a portfolio as a set of risk exposures to be parsed into systematic risks (betas) and idiosyncratic risks (alphas), let's focus specifically on the systematic risks and examine their conditional distribution characteristics. *Consider a portfolio of hedge funds as a portfolio of insurance policies against a set of risk factors.* Now there are two ways to think about the value of an insurance policy:

First, we may estimate the *unconditional* value of the policy as :

$$V = P(\text{damage})*IP + P(\text{no damage})*PREM - DED \qquad (1)$$

Where V is the policy value
IP is the present value of the Insurance Payout
P is the unconditional probability, $0 \leq P \leq 1$, and P(damage) + P(no damage)=1
PREM is the present value of the premium payments
DED is the present value of the deductible amount

Equation 1 tells us what we should be willing to pay for the insurance policy. This is the analog of our unconditional beta estimate – unconditional beta can be interpreted as our best overall estimate of the "value" of a hedge fund with respect to its exposure to a risk factor.

Second, we may estimate the *conditional* value of the policy as:

$$V = IP \qquad (2)$$

Equation 2 tells us how much the insurance company will pay us when the loss we are insuring against actually occurs. It is derived by setting P(damage) to 1 and P(no damage) to 0 in equation 1. This is the analog of our conditional beta estimate –

conditional beta can be interpreted as our best estimate of the value of a hedge fund with respect to its exposure to the left tail of a risk factor's return distribution.

Nonlinear Tail Risk

Suppose we want to understand how a hedge fund (or a set of hedge funds) could protect us against a US equity market crash. Would the unconditional beta estimate be useful to us? Not necessarily. The following is an example of the net returns of an anonymous hedge fund which existed for the six years ended December 2007, and its statistical relation with the S&P 500 Index, our proxy for the US equity market. Figure 1 shows the distribution of returns for the fund and the S&P over the sample period. The hedge fund appears to have a fairly similar left tail, a flatter peak, and a fatter right tail, relative to the index.

Figure 1

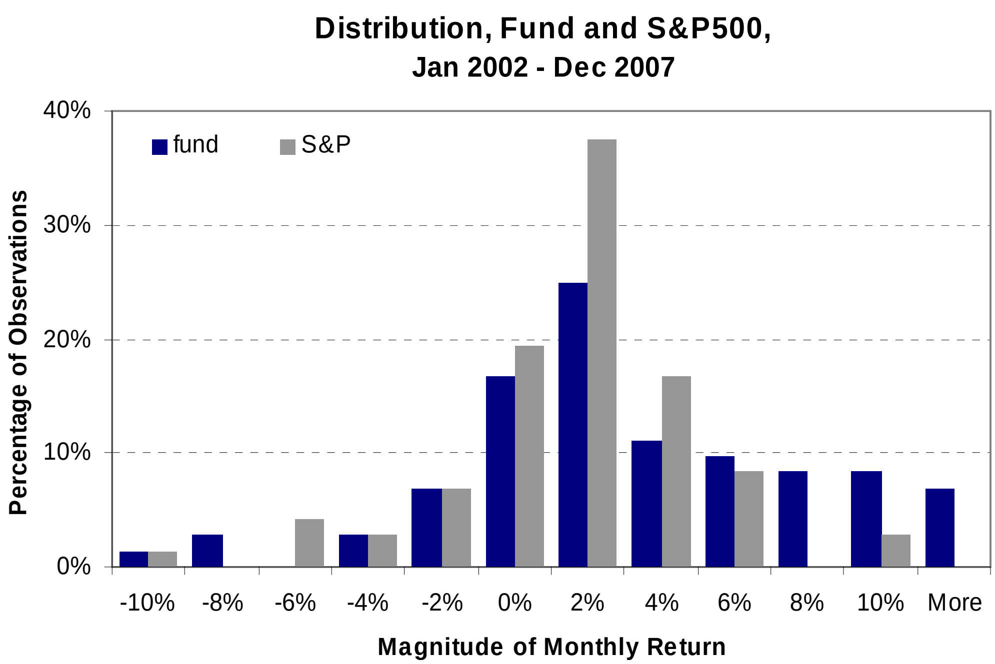

One problem with making observations about the distribution of returns is the sequence is ignored; time series characteristics of the return patterns are washed out of the picture. To evaluate time series characteristics, it is common in the industry to run rolling estimates of correlation, as shown in Figure 2. Figure 2 depicts a trend that could be "worrisome" to an investor concerned about the fund's ability to hedge against US equity risk, since the trend of all the estimates is upwards – the 12-month correlation, while volatile, has risen from -0.63 to +0.52; the 24-month correlation has risen from +0.12 to +0.53; and the 36-month correlation has risen from +0.17 to +0.53. Standard MPT interpretation of these trends would lead the investor to conclude that the fund has become less of a "diversifier" for US equity risk over time, and therefore less attractive as a hedge. This approach is better than simply observing the distribution of returns, but it continues to suffer from taking an unconditional approach – it ignores the condition of the market. We can improve our understanding of the hedge properties of the fund by modifying our correlation analysis to be conditioned on the market return. The reason we should do this is the investor ideally wants a nonlinear exposure (an option-like payoff), meaning when the market does well we prefer high correlations and when the market does poorly we prefer low correlations.

Figure 2

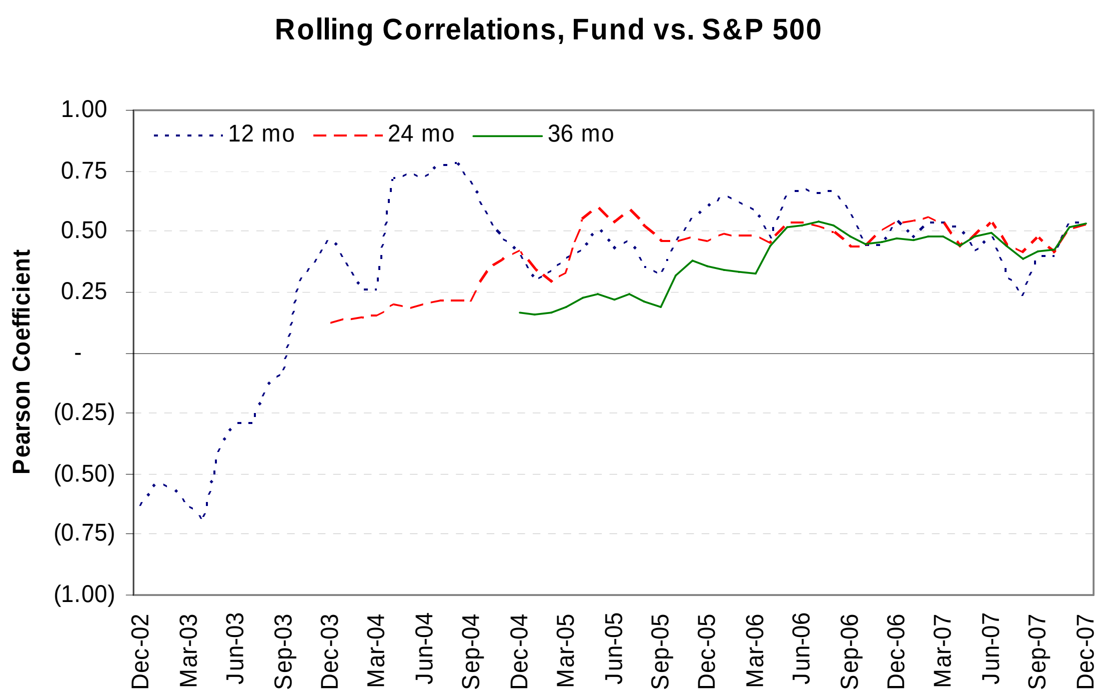

Figures 3a and 3b show a visual method I like to use in order to analyze the conditional correlation and the conditional beta of the fund with respect to risk factors (in this case, US equity risk proxied by the S&P 500 index), which can tell us whether there are (un)attractive nonlinear correlation characteristics we might have missed simply by analyzing the distribution and the rolling linear estimates.

Figure 3a

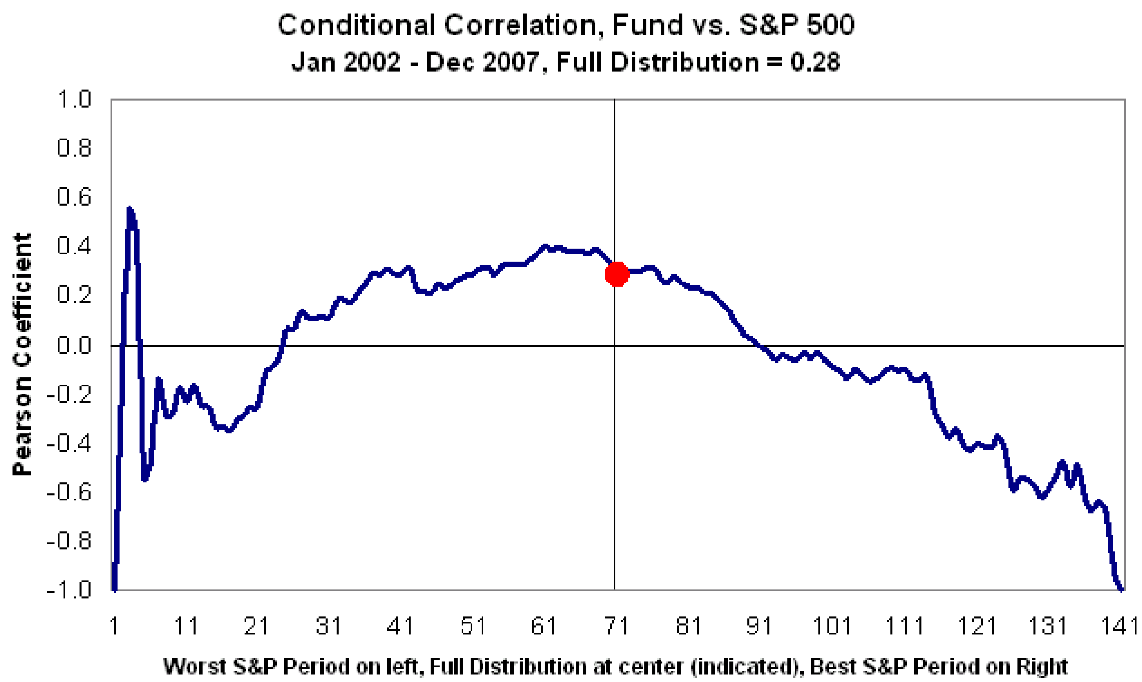

Figure 3b

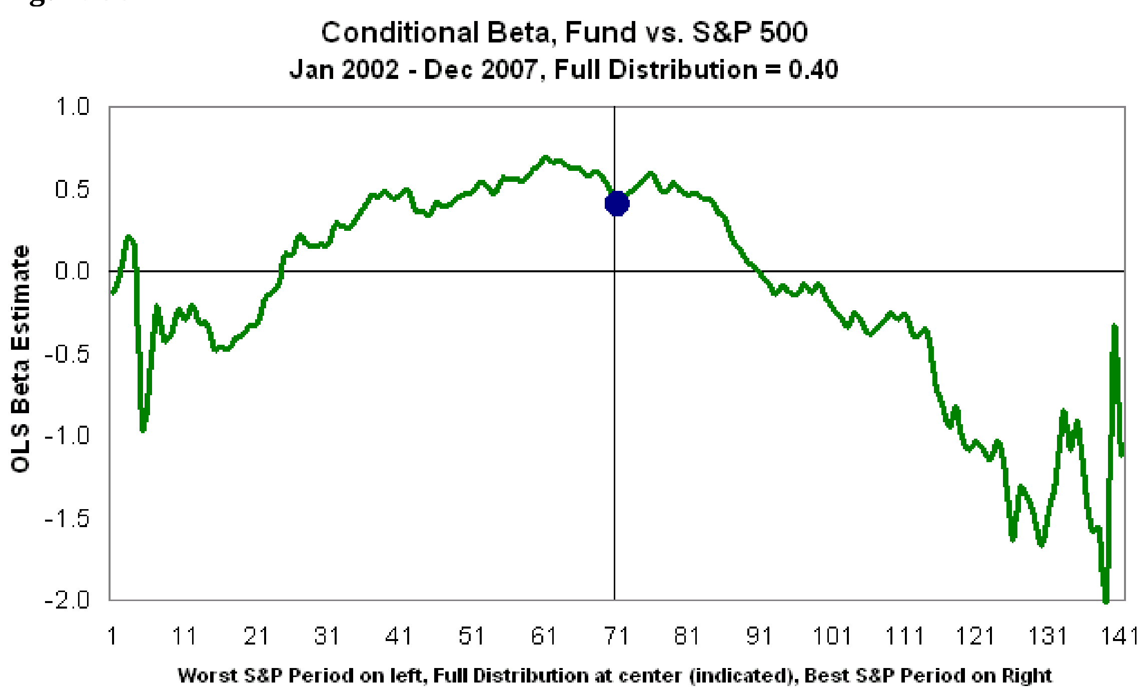

Readers familiar with options will perhaps interpret Figure 3a as a "short correlation straddle" (excusing the brief 2-period spike up to +0.55 in the third-worst S&P period), and Figure 3b as a "short beta straddle". The unconditional correlation coefficient of the entire 72-month (n=72) distribution is 0.28, highlighted by the point in the center of figure 3a, while the unconditional beta of the entire 72-month (n=72) distribution is 0.40, highlighted by the point in the center of figure 3b. As we move to the left from these center points we remove the next marginally best month of the S&P 500 index from the distribution, continuing to do so until we arrive at the left-most estimate of correlation/beta, conditional on the left tail of the S&P 500 index returns; here n=2. As we move to the right from the unconditional estimate, we remove the next marginally worst month of the S&P 500 index from the distribution, continuing to do so until we arrive at the right-most estimate of correlation/beta, conditional on the right tail of the S&P 500 index returns; here also, n=2. We see that the manager does provide diversification in nearly all markets, and generally that the higher or lower the market, the lower the manager's correlation becomes. Presumably, readers interested in capturing gains from US equities yet concerned about losses due to US equities would prefer lower correlations and betas during poor S&P periods (the left side of Figures 3a and 3b) and higher correlations during good S&P periods (the right side of Figures 3a and 3b). This fund delivers part of what the investor seeks -- lower correlations and betas as the S&P does poorly. However, the fund also delivers lower correlations and betas as the S&P does well. Conditionally, this fund appears to be a better diversifier than the rolling correlations in Figure 2 would indicate.

It is well known that, when the relationship of variables is nonlinear, the correlation coefficient is a poor summary statistic, because it cannot replace the examination of the individual data. Furthermore, OLS regressions can be completely misleading in such cases. Figure 4 illustrates Anscombe's (1973) quartet, in which each of the four sets of data pairs exhibit a correlation of 0.82 and have the same OLS regression coefficients. Clearly there is a complete failure of the correlation and regression coefficients to distinguish between the differing relationships in the four panels of Figure 4.

Figure 4

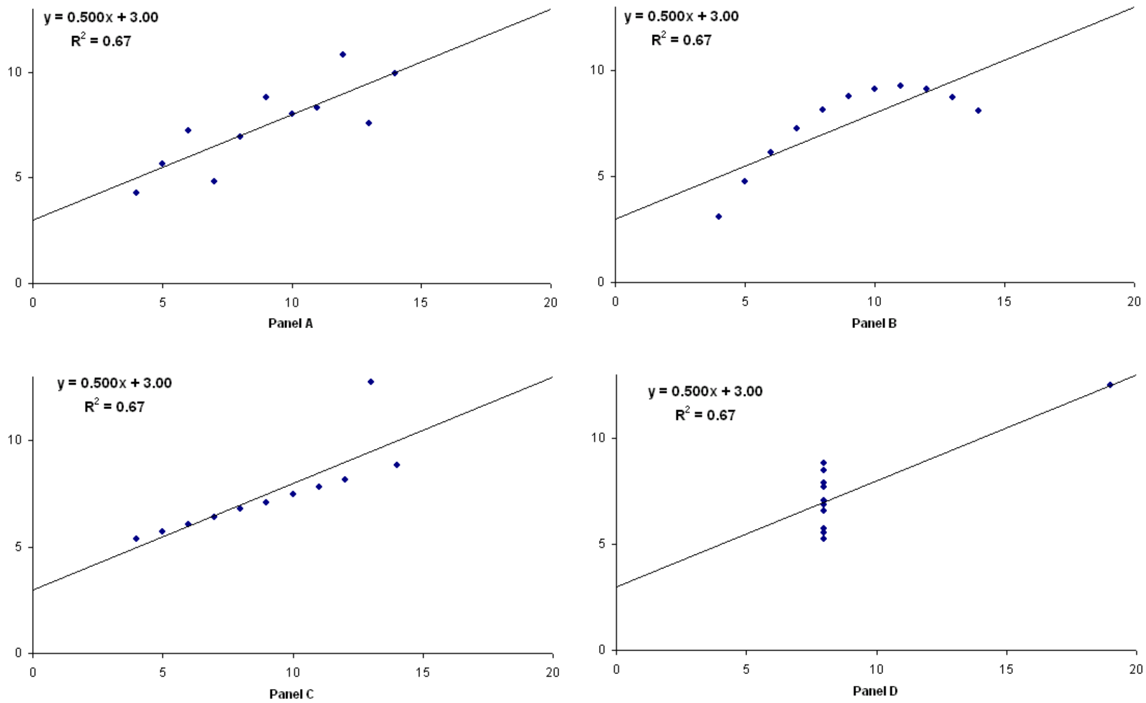

Because of this deficiency of linear estimates, we directly examine the fund returns relative to the S&P returns directly in Figure 5. This chart is sorted on the S&P returns, and a visual inspection of Figure 5 reconciles easily with the results in Figures 3 and 1, particularly the negative extreme tail correlation at both ends of the distribution.

Figure 5

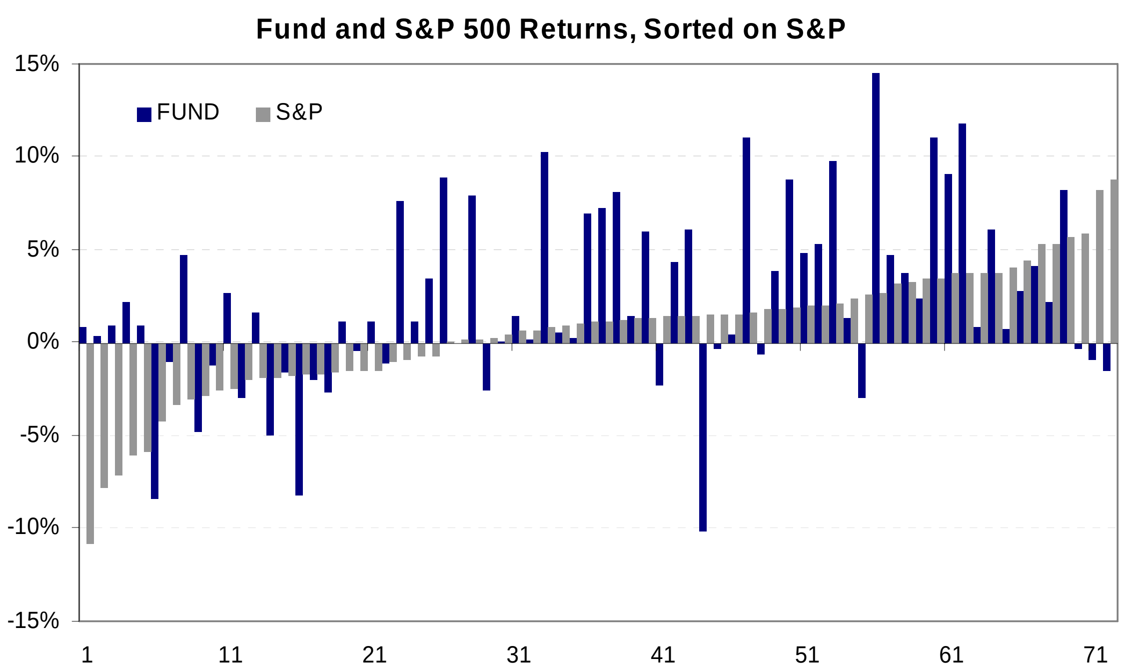

We can also view the relationship of the fund and the S&P500 index using regression scatterplots rather than simple correlations, which are a component of the OLS equation. In figure 6 we estimate the unconditional beta of the fund to the S&P500 Index. The beta estimate of +0.40 might lead the analyst to assume that if the US equity market were to drop, say, 10% in the future, that the fund would drop 4%. However, looking at the data we see a different, nonlinear pattern. In this history, the S&P500 index dipped below 10% once, returning just better than -11% in September 2002. In September 2002, the fund returned nearly +1%. Because this fund delivers nonlinear payoffs, the unconditional beta estimate can be misleading, particularly in tail events.

Figure 6

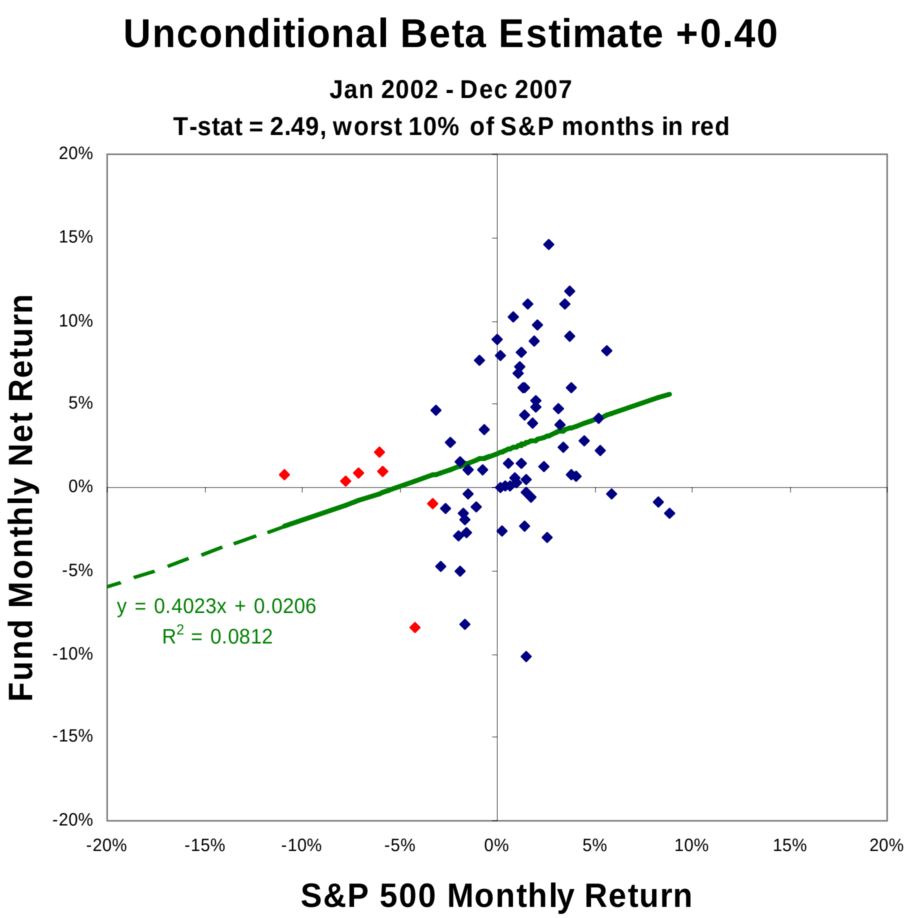

In Figure 7 we estimate a conditional beta, running an OLS regression not on the entire distribution, but only on the left tail of the return distribution of the S&P500 index. This conditional beta estimate of -0.68 has the opposite sign as the unconditional beta, and although its slope is even greater, we must discount our belief in its accuracy because the statistical significance is weaker than that of the unconditional beta, because the data in this tail region are sparse. Nonetheless the strong difference between our unconditional and conditional beta estimates can help us to understand the nonlinear nature of the fund's exposures to the US equity market.

Figure 7

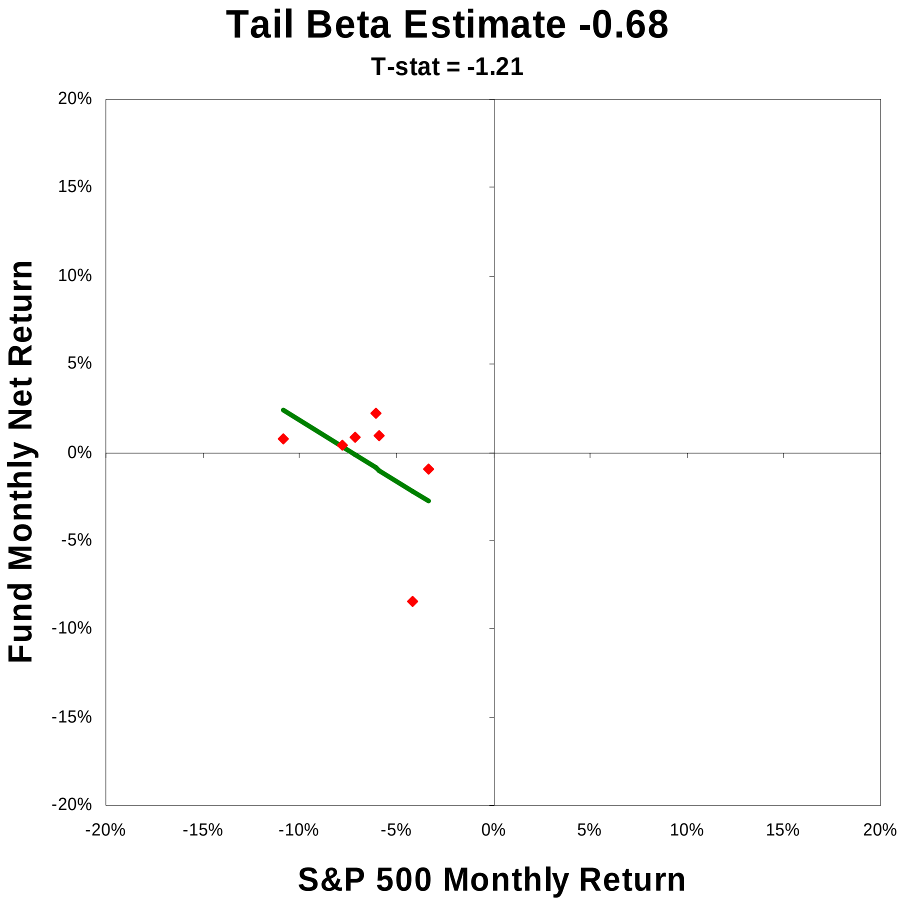

A quick glance at the fit of the piecewise regression in Figure 7 is enough to inform us that the linear OLS approach is not necessarily a good model of tail fit. The low coefficient of determination of 0.23 and the associated t-statistic of -1.21 confirms this conclusion. We can potentially overcome some of the deficiencies of correlation and OLS regression methods by fitting a polynomial, as in the quadratic fit depicted in Figure 8. Our unconditional regression in Figure 6 would lead us to expect the hedge fund to lose 5.99% when the US equity market is down 20%; our conditional regression in Figure 7 would lead us to expect the hedge fund to make 8.63%; however since the linear fit is so poor perhaps we should really expect the quadratic fit to be a better indicator, at a loss of 19.78%. Unfortunately the fit of the quadratic tail estimate is only marginally better than the linear fit, with a coefficient of determination of only 0.33.

Figure 8

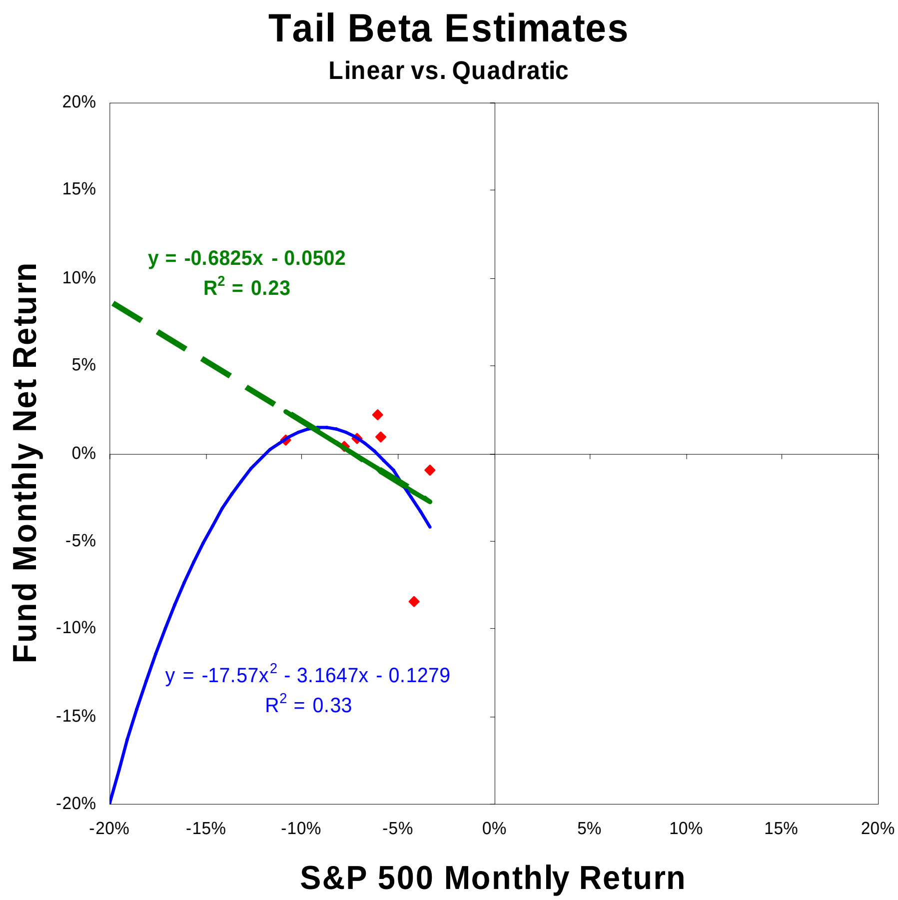

Extreme Value Theory

Moving on from the low power of regressions, we explore a method known as extreme value theory (EVT) in order to estimate loss severity distributions, a technique used in the insurance industry. Employment of such insurance technology is appealing since in this exercise we are considering the hedge fund to be an insurance policy against poor factor returns.

The fundamental result of EVT is the Fischer-Tippett (1928) theorem, which can be thought of as an extreme version of the central limit theorem. F-T describes the limiting behavior of normalized sample maxima. This theorem suggests fitting the generalized extreme value (GEV) distribution to sample maxima, and has been long employed in other fields, for example hydrology, and more recently, insurance. Gumbel (1958) is a standard reference. We observe that our tail-loss data are maxima (maximum losses, which one might otherwise consider minima), and proceed to fit the GEV distribution to our data.

Define the distribution function of the GEV as:

$H_{\xi}(x) = \{exp(-(1+\xi x)^{-1/\xi})$ if $\xi \neq 0$; $exp(-e^{-x})$ if $\xi = 0\}$

Where ξ is the shape parameter (indicating the fatness of the tail) and x satisfies $1+\xi x > 0$.

There are three special cases of the GEV which define classes of distributions:

1) when $\xi > 0$ the GEV falls within the Fréchet distribution class with the shape parameter $\alpha=1/\xi$; this class of distributions exhibits tails that decay like power functions, and includes the Cauchy, Fréchet, Pareto, and Gosset's t-distribution. These are known as "long-tailed" distributions.
2) when $\xi = 0$ the GEV resides in the Gumbel distribution class, a class including the exponential, gamma, Gumbel, lognormal, and normal distributions. These are known as "moderate-tailed" distributions and are popular in the insurance field for estimating loss severity distributions.
3) when $\xi < 0$ the GEV is a member of the Weibull distribution class. With the shape parameter $\alpha = -1/\xi$; this class of distributions includes the beta, uniform and Weibull distributions. These are known as "short-tailed" distributions.

Because we are not "worried" about a short left tail, the third class is not interesting for our purpose, so we will investigate the other classes by first attempting to fit the data to the Gumbel distribution class, and then fitting the data to the Fréchet class. Figure 9 shows our test for whether the risk factor, the S&P500 index, fits a Gumbel distribution. The location parameter is 2.09% and the scale parameter is -0.027, with a coefficient of determination exceeding 97%. Testing for goodness-of-fit, we obtain a Kolmogorov-Smirnov (K-S) statistic of 0.11 with a p-value of 0.32 which is not rejected at the 95th level of confidence. The Anderson-Darling (A-D) statistic, which gives more weight to the tails than the K-S test, is 0.90, also not rejected at the 95th level of confidence. Finally, we obtain a Chi-squared statistic of 8.33 with a p-value of 0.21, which is also not rejected at the 95th level of confidence. Given the strong linearity and statistical significance, we determine that the Gumbel distribution provides an adequate model of the US Equity market returns over the 2002-2007 period.

Figure 9

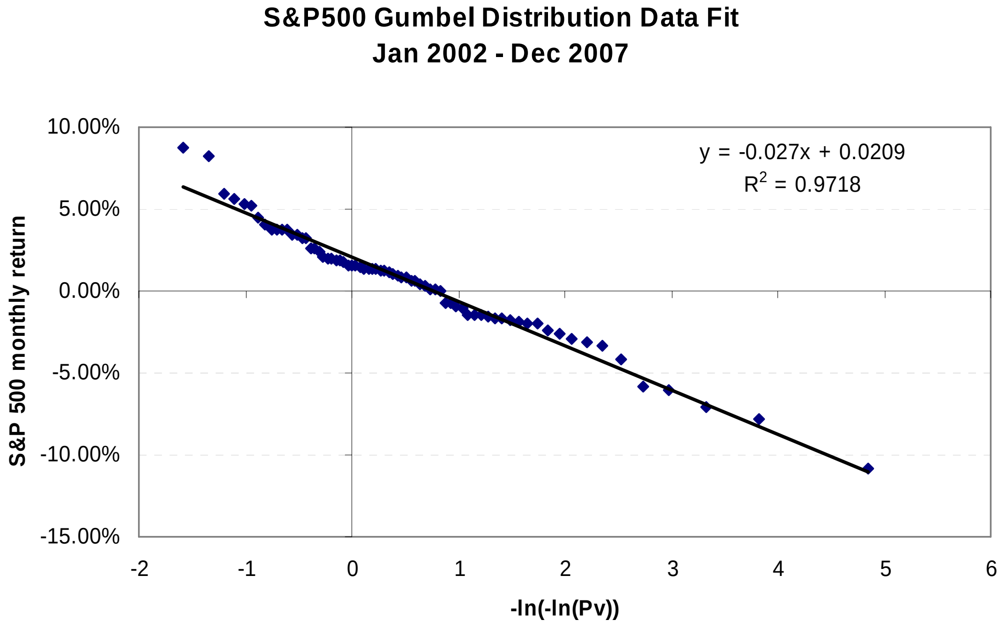

Next we test for whether the hedge fund fits, unconditionally, a Gumbel distribution. Figure 10 depicts this test, with a location parameter of 4.39% and the scale parameter of -3.70%, and a coefficient of determination just under 92%. The fund appears less linear than the S&P 500 index, exhibiting a distinct kink close to the zero point. Testing for goodness-of-fit, we obtain a K-S statistic of 0.18 with a p-value of 0.01 which is rejected at the 95th level of confidence. The A-D statistic, which gives more weight to the tails than the K-S test, is 3.07, also rejected at the 95th level of confidence. Finally, we obtain a Chi-squared statistic of 13.29 with a p-value of 0.02, which is also rejected at the 95th level of confidence. Given the evident piecewise linearity, we determine that the Gumbel distribution provides a less-than adequate model of the hedge fund's returns over the 2002-2007 period.

Figure 10

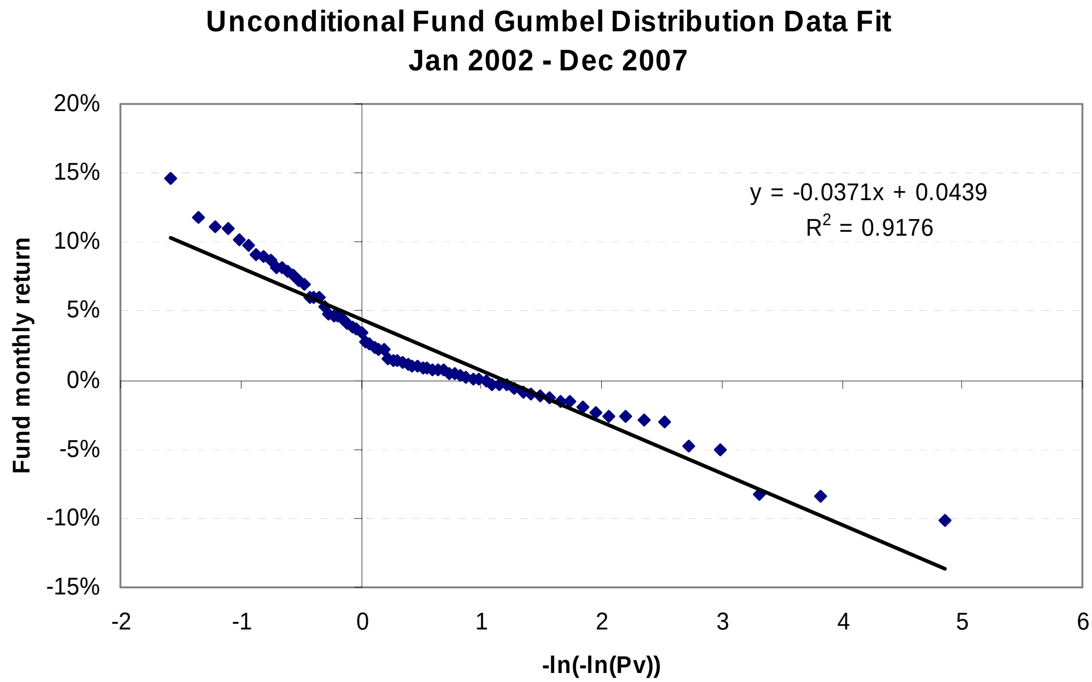

Now we test whether the hedge fund, conditioned on the left 10% tail of S&P500 returns, fits a Gumbel distribution. Figure 11 depicts this test, with a location parameter of 13bp and the scale parameter of -1.40%, and a coefficient of determination just under 21%. Conditioned on the left tail of the risk factor, the fund appears far from linear; in fact, fitting a quartic (4^{th} degree polynomial) to the conditional data increases the coefficient of determination from the linear case of just 0.21 to a healthy 0.97. Based on the nonlinearity, we determine that the Gumbel distribution provides a poor model of the hedge fund's returns conditioned on the left tail of the US equity market returns over the 2002-2007 period.

Although the returns of the S&P 500 index fit a Gumbel distribution pretty well, the hedge fund's returns do not – especially when conditioned on the left tail of the S&P 500 return distribution. Therefore we move from the Gumbel to the Fréchet class of GEV distributions.

Figure 11

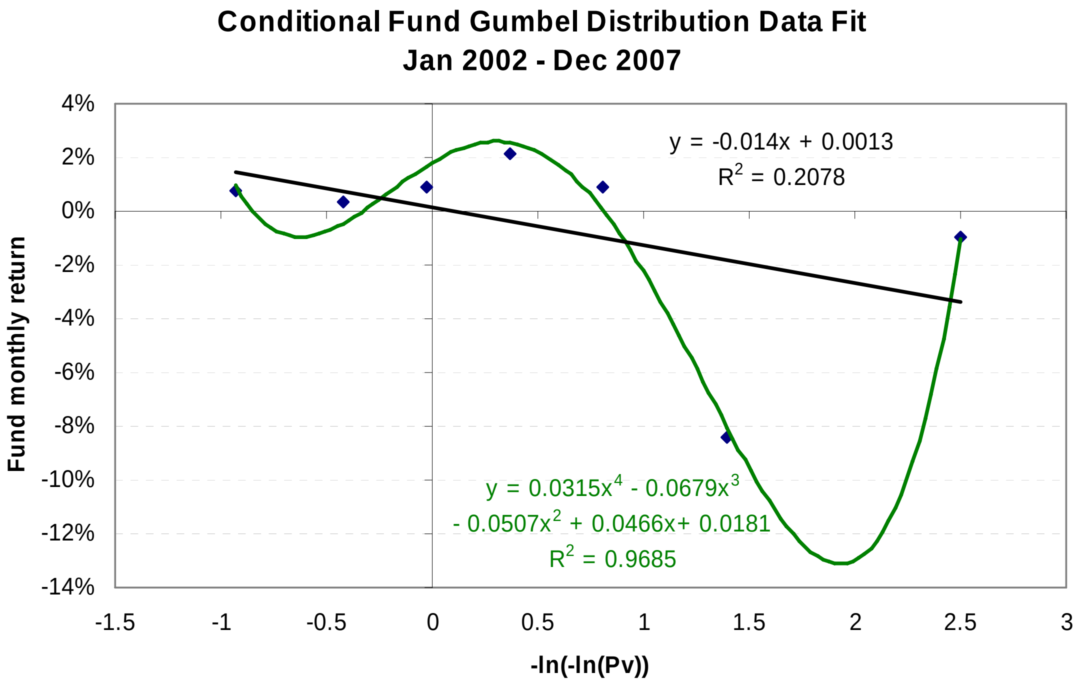

To test whether the hedge fund, conditioned on the left 10% tail of S&P500 index returns, fits a Fréchet distribution, we assume that the maximum domain of attraction (MDA) conditions hold, and invoke the Balkema-De Haan (1974) theorem, which states that under the MDA conditions, the generalized Pareto distribution (GPD) is the limiting distribution, and fit the GPD to our tail data using the method of L-moments. As before, we assume the data are IID, and we choose a threshold from the S&P dataset, at the worst 10% of observations and examine the tail events only. Unfortunately, both the A-D test and the K-S test reject the GPD distribution at the 95th level of confidence. We infer that, for this hedge fund the MDA conditions do not hold, the data are not IID, or both.

The best fit we can find, based on the A-D test, for the conditional hedge fund returns is to the Cauchy distribution (which is a long-tailed distribution within the Fréchet class). The Cauchy distribution is characterized by a kurtosis parameter α = 1, a skewness parameter $\beta = 0$, with variable scale parameter σ estimated at 0.00422 and shift parameter μ estimated at 0.00788 for the conditional hedge fund returns, with a K-S statistic of 0.26 and p-value of 0.64 and an A-D statistic of 0.66, this fit is statistically significant under both tests at the 95th level of confidence.

Goodness of Fit Testing

Given that we are concerned with the possibility of a long left tail, we now turn to the broad hedge fund universe to examine which hedge fund strategies might exhibit such characteristics. Figure 12 shows out goodness of fit results, for 12 HFRX indexes over the 120 month period from January 1998 through December 2007. The HFRX indexes are investable hedge fund indexes provided by Hedge Fund Research.[1] I fit each index return vector to nine distributions: three, the Generalized Pareto ("GPD"), Cauchy and Fréchet distributions, represent the long-tailed class of Generalized Extreme Value ("GEV") distributions; three, the Gamma, Normal and Lognormal distributions, represent the medium-tailed class of GEV distributions; and three, the Beta, Weibull and Uniform distributions, represent the short-tailed class of GEV distributions.

Figure 12

	EWS	EMN	MAC	MDI	EQH	GHF	ARI	CAI	EDI	MAI	RVA	DSI
GPD	1(k,a,c)	1(k,a,c)	1(k,a,c)	1(k,a,c)	8(k)	7(k)	8(k)	8(k)	7(k)	8(k)	8(k)	8(k)
Cauchy	6	6(c)	8	6(a)	6(c)	4(c)	2(k,a,c)	5(k,a,c)	6(c)	5(k,a,c)	6(c)	6(c)
Frechet	7	7(c)	7(c)	8	7	5	7	7	8(c)	7	7	5(c)
Gamma	3(k,a,c)	3(k,a,c)	3(k,a,c)	3(k,a,c)	1(k,a,c)	1(k,a,c)	3(k,c)	2(k,a,c)	3(k,a,c)	2(k,a,c)	2(k,a,c)	2(k,a,c)
Normal	5(k,a)	5(k,a,c)	6(k,a,c)	4(k,a,c)	5(c)	6(c)	1(k,a,c)	4(k,a,c)	4(k,a,c)	4(k,a,c)	5(k,c)	4(k,c)
Lognormal	4(k,a,c)	4(k,a,c)	4(k,a,c)	5(k,a)	3(k,a,c)	3(k,a,c)	4	6(k,a,c)	5(k,a,c)	6(k,a)	4(k,a,c)	3(k,a,c)
Beta	8(k)	8(k)	5(k,a,c)	7(k)	4(k,a,c)	8	5(k)	1(k,a,c)	1(k,a,c)	1(k,a,c)	1(k,a,c)	7(k)
Weibull	2(k,a,c)	2(k,a,c)	2(k,a,c)	2(a,c)	2(k,a,c)	2(k,a,c)	6	3(k,a,c)	2(k,a,c)	3(k,a,c)	3(k,a,c)	1(k,a,c)
Uniform	9(k)	9	9	9(k)	9	9	9(k)	9(k)	9(k)	9(k)	9	9

In figure 12, each distribution is ranked from 1 to 9 for every HFRX index, with 1 indicating the "best" fit, and 9 indicating the "worst" fit when we fit the distributions to the empirical HFRX data. We utilize three goodness-of-fit test statistics, the Kolmogorov-Smirnov ("K-S"), Anderson-Darling ("A-D"), and Pearson's Chi-Squared ("C-S"). Because K-S relies on a normally distributed sample CDF, the results in Figure 12 indicate the rank order based on the Anderson-Darling test, which gives more weight to the fit in the tails of the distribution. Test statistics which are significant at the 95th level of confidence are listed in the parentheses.

Four of the HFRX indexes strongly display long-tailed GEV characteristics: the Equal Weighted Strategies Index ("EWS"), Equity Market Neutral Index ("EMN"), Macro Index ("MAC"), and Market Directional Index ("MDI") all fit best with the Generalized Pareto distribution, with all three test statistics indicating statistical significance. Three of the indexes fit extremely well to the medium-tailed GEV distributions: the Equity Hedge Index ("EQH") and the Global Hedge Fund Index ("GHF") fit best to the Gamma distribution, while the Absolute Return Index ("ARI") fits best to the normal distribution, with all three test statistics indicating statistical significance. Finally, five of the HFRX indexes strongly display short-tailed GEV characteristics: the Convertible Arbitrage Index ("CAI"), Event Driven Index ("EDI"), Merger Arbitrage Index ("MAI"), and Relative Value Index ("RVA") all fit the Beta distribution best, with all three test statistics indicating statistical significance, and the Distressed Securities Index ("DSI") fits the Weibull distribution best, with all three test statistics indicating statistical significance.

[1] Refer to Appendix 1 for descriptions of the HFRX indexes.

Figure 13 shows the parameters for each of the best-fit results. We obtain our estimates of long-tai fit, for the Generalized Pareto Distribution, using the L-moment method; we obtain our estimates of medium-tail fit, for the Gamma distribution, using the Generalized Method of Moments, and for the normal distribution using Maximum Likelihood; we obtain our estimates of short-tail fit, for the Beta distribution, using the Generalized Method of Moments to obtain the initial parameter estimates, and then Maximum Likelihood to obtain the final estimates, and for the Weibull distribution using Least Squares.

Figure 13

	EWS	EMN	MAC	MDI	EQH	GHF	ARI	CAI	EDI	MAI	RVA	DSI
Long Tail	GPD: k=-0.374 s=0.013 m= 9.967E-4	GPD: k=-0.232 s=0.010 m=- 7.875E-5	GPD: k=-0.221 s=0.029 m=- 2.647E-4	GPD: k=-0.408 s=0.028 m= 7.541E-4								
Medium Tail					Gamma: a=1.245 b=0.016	Gamma: a=1.195 b=0.0130	Normal: s=0.006 m=0.009					
Short Tail								Beta: a1=2.156 a2=10.502 a=-4.168E-4 b=0.065	Beta: a1=1.104 a2=3.013 a=1.697E-4 b=0.055	Beta: a1=1.848 a2=7.114 a=-2.333E-4 b=0.049	Beta: a1=2.370 a2=8.38E+6 a=-6.55E-4 b=38000	Weibull: a=1.174 b=0.01477

A comment on the goodness of fit testing: Because the HFRX indexes are indexes, by construction some tail characteristics of the constituent hedge funds are almost certain to be masked. Therefore, readers should use caution in their interpretation of Figure 12. Because of less than perfect positive correlations in the off-diagonals, index distribution characteristics are likely to appear "shorter tailed" than their constituents' return distributions actually are. In fact, to be conservative a good practice would be to find the highest row in the table for which all of the test statistics (K-S, A-D, and C-S) indicate statistical significance. So for example, although the HFRX Absolute Return index has a "best" fit to the normal distribution, it can also be fit with a similar level of statistical significance by the Cauchy distribution, a long-tailed, rather than a medium-tailed, distribution. For estimating tail risk, I would assume the Cauchy rather than the Normal to be conservative. Likewise, all of the indexes which are "best" fit by the short-tailed distributions (CAI, EDI, MAI, RVA and DSI) can also be fit with a similar level of statistical significance by at least two of the medium-tailed distributions, while two of them, CAI and MAI, can even be well-fit with the long-tailed Cauchy distribution. For estimation of tail risk in practice, Figure 14 is a modification of Figure 12, showing the "conservatively" appropriate distributions for the HFRX Indexes:

Figure 14

	EWS	EMN	MAC	MDI	EQH	GHF	ARI	CAI	EDI	MAI	RVA	DSI
GPD	1(k,a,c)	1(k,a,c)	1(k,a,c)	1(k,a,c)	8(k)	7(k)	8(k)	8(k)	7(k)	8(k)	8(k)	8(k)
Cauchy	6	6(c)	8	6(a)	6(c)	4(c)	2(k,a,c)	5(k,a,c)	6(c)	5(k,a,c)	6(c)	6(c)
Frechet	7	7(c)	7(c)	8	7	5	7	7	8(c)	7	7	5(c)
Gamma	3(k,a,c)	3(k,a,c)	3(k,a,c)	3(k,a,c)	1(k,a,c)	1(k,a,c)	3(k,c)	2(k,a,c)	3(k,a,c)	2(k,a,c)	2(k,a,c)	2(k,a,c)
Normal	5(k,a)	5(k,a,c)	6(k,a,c)	4(k,a,c)	5(c)	6(c)	1(k,a,c)	4(k,a,c)	4(k,a,c)	4(k,a,c)	5(k,c)	4(k,c)
Lognormal	4(k,a,c)	4(k,a,c)	4(k,a,c)	5(k,a)	3(k,a,c)	3(k,a,c)	4	6(k,a,c)	5(k,a,c)	6(k,a)	4(k,a,c)	3(k,a,c)
Beta	8(k)	8(k)	5(k,a,c)	7(k)	4(k,a,c)	8	5(k)	1(k,a,c)	1(k,a,c)	1(k,a,c)	1(k,a,c)	7(k)
Weibull	2(k,a,c)	2(k,a,c)	2(k,a,c)	2(a,c)	2(k,a,c)	2(k,a,c)	6	3(k,a,c)	2(k,a,c)	3(k,a,c)	3(k,a,c)	1(k,a,c)
Uniform	9(k)	9	9	9(k)	9	9	9(k)	9(k)	9(k)	9(k)	9	9

Even after making such an adjustment, best practice would be to examine the individual hedge funds rather than the indexes. For example, for the HFRX Event Driven Index, we first found a "best" fit in the short-tailed class of GEV (the Beta) distribution, and then we conservatively chose to use the Gamma Distribution, in the medium-tailed GEV category. While it might be tempting to assume that the Gamma distribution would, therefore, be appropriate to apply for tail-risk analysis of individual hedge funds within the "Event Driven" category, such an approach would be hazardous. For example, the anonymous hedge fund analyzed earlier in this paper (refer to Figs. 1-3, 5-8, 10 and 11), which best was fit by the Cauchy distribution, is classified as an Event-Driven hedge fund. The Cauchy is one of the long-tailed GEV distributions. It certainly would have been a mistake to assume that, because the fund is classified in a category, that the HFRX index results should be applied to the fund itself, since, in this case, even the Gamma is a shorter-tailed GEV distribution than is the Cauchy.

The Fractal Geometry of Hedge Funds

The Cauchy distribution is interesting for a number of reasons; first, falling within the Fréchet class of GEVs, it is long-tailed; second, its higher moments may not exist, or if they exist they may be infinite; third, these distributions exhibit self-similarity; finally, the Cauchy distribution was the first known non-Gaussian stable law. Bienayme (1853) pointed out that OLS regressions will be useless if the expected value of the sum (as in the Cauchy case) is infinite. Nearly half a century ago, Mandelbrot (1963) proposed that at least some asset prices follow such stable distributions. The Cauchy distribution (also known as the Lorentz distribution) is a member of the Levy skew alpha-stable distributions, which are called by Mandelbrot (1963) and Fama (1963) "L-stable" distributions. According to Gnedenko and Kolmogorov (1954), such distributions are the only possible limiting distributions for sums of IID random variables. Suppose that underlying securities follow IID random paths; if they have finite variance, portfolios of securities (hedge funds, for example) would have limiting distribution that are normal. However, since empirically we find hedge funds very rarely display normal distributions (and hence, the motivation for this paper), the underlying securities may instead have infinite variance and (if so) the limiting distribution of their sums (hedge funds) must be Levy skew alpha-stable. The empirical results, in the returns of the anonymous hedge fund and many of the HFRX indexes examined herein seem to support L-Stability. One implication of our empirical findings may be that Mandelbrot's Stable Paretian Hypothesis was correct even though it was cast aside by financial economists decades ago.

References

Anderson, T. and D. Darling, (1952), "Asymptotic theory of certain "goodness-of-fit" criteria based on stochastic processes". *Annals of Mathematical Statistics* Vol. 23, pp. 193 - 212.

Anscombe, F.J. (1973), "Graphs in Statistical Analysis". *The American Statistician* Vol. 27, No. 1, February, pp.17-21.

Balkema, A. and L. de Haan (1974), "Residual Life Time at Great Age". *Annals of Probability* Vol. 2, pp. 792-804.

Bienaymé, J. (1853), "Considérations à l'appui de la découverte de Laplace sur la loi de la probabilité dans la méthode des moindres carrés". *Comptes Rendus* Vol. 37, pp. 309-329.

Cauchy, A. (1853), "Sur les résultats les plus probables". *Comptes Rendus* Vol. 37, pp. 198-206.

Fama, E. (1963), "Mandelbrot and the Stable Paretian Hypothesis" *Journal of Business* Vol. 36, No. 4, October. pp. 420-429.

Fischer, R. and L. Tippett (1928), "Limiting Forms of the Frequency Distribution of the Largest or Smallest Member of a Sample". *Proceedings of the Cambridge Philosophical Society* Vol. 24, pp. 180-190.

Gnedenko, B. and A. Kolmogorov (1954), Limit Distributions for Sums of Independent Random Variables. K. Chung (trans). Reading: Addison Wesley.

Gumbel, E. (1958), Statistics of Extremes. New York: Columbia University Press.

HFR (2008) "HFRX Hedge Fund Indices: Construction and Methodology" Whitepaper. Chicago: Hedge Fund Research, Inc.

Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices" in Cootner, P. (ed.) The Random Character of Stock Market Prices. Cambridge: MIT Press. pp. 297-337.

Markowitz, H. (1959), Portfolio Selection. New Haven: Cowles Foundation.

Appendix

For complete descriptions of the methodology used in construction of the HFRX indexes, please refer to HFR (2008).

Brief descriptions, from https://www.hedgefundresearch.com/, of each of the HFRX indexes follow.

HFRX Global Hedge Fund Index - HFRXGL
The HFRX Global Hedge Fund Index is designed to be representative of the overall composition of the hedge fund universe. It is comprised of eight strategies; convertible arbitrage, distressed securities, equity hedge, equity market neutral, event driven, macro, merger arbitrage, and relative value arbitrage. The strategies are asset weighted based on the distribution of assets in the hedge fund industry.

HFRX Equal Weighted Strategies Index - HFRXEW
HFRX Equal Weighted Strategies Index is constructed from the same strategies as the HFRX Global Hedge Fund Index and every strategy is given equal weight. As a result, it offers a more balanced diversification and historically lower volatility.

HFRX Absolute Return Index - HFRXAR
The HFRX Absolute Return Index is designed to be representative of the overall composition of the hedge fund universe. Similar to HFRX Global, the strategies are asset weighted based on the distribution of assets in the hedge fund industry. However, as a component of the quantitative optimization and fund selection process, HFRXAR seeks to select constituents which are likely to provide stable performance regardless of market conditions, which are characteristically less volatile and less correlated to market benchmarks.

HFRX Market Directional Index - HFRXMD
The HFRX Market Directional Index is designed to be representative of the overall composition of the hedge fund universe. Similar to HFRX Global, the strategies are asset weighted based on the distribution of assets in the hedge fund industry. However, as a component of the quantitative optimization and fund selection process, HFRXMD seeks to select constituents which add value by participating in the direction of various financial markets; these characteristically have higher expected volatility than Absolute Return constituents.

HFRX Convertible Arbitrage Index - HFRXCA
Convertible Arbitrage involves taking long positions in convertible securities and hedging those positions by selling short the underlying common stock. A manager will, in an effort to capitalize on relative pricing inefficiencies, purchase long positions in convertible securities, generally convertible bonds, convertible preferred stock or warrants, and hedge a portion of the equity risk by selling short the underlying common stock. Timing may be linked to a specific event relative to the underlying company, or a belief that a relative mispricing exists between the corresponding securities. Convertible

securities and warrants are priced as a function of the price of the underlying stock, expected future volatility of returns, risk free interest rates, call provisions, supply and demand for specific issues and, in the case of convertible bonds, the issue-specific corporate/Treasury yield spread. Thus, there is ample room for relative misvaluations.

HFRX Distressed Securities Index - HFRXDS
Distressed Securities managers invest in, and may sell short, the securities of companies where the security's price has been, or is expected to be, affected by a distressed situation. Distressed securities managers invest primarily in securities and other obligations of companies that are encountering significant financial or business difficulties, including companies which (i) may be engaged in debt restructuring or other capital transactions of a similar nature while outside the jurisdiction of Federal bankruptcy law, (ii) are subject to the provisions of Federal bankruptcy law or (iii) are experiencing poor operating results as a result of unfavorable operating conditions, over-leveraged capital structure, catastrophic events, extraordinary write-offs or special competitive or product obsolescence problems. Managers will seek profit opportunities arising from inefficiencies in the market for such securities and other obligations.

Negative events, and the subsequent announcement of a proposed restructuring or reorganization to address the problem, may create a severe market imbalance as some holders attempt to sell their positions at a time when few investors are willing to purchase the securities or other obligations of the troubled company. If a manager believes that a market imbalance exists and the securities and other obligations of the troubled company may be purchased at prices below the value of such securities or other obligations under a reorganization or liquidation analysis, the manager may purchase the securities or other obligations of the company. Profits in this sector result from the market's lack of understanding of the true value of the deeply discounted securities. Results are generally not dependent on the direction of the markets, and have a low to moderate expected volatility.

HFRX Equity Hedge Index - HFRXEH
Equity Hedge, also known as long/short equity, combines core long holdings of equities with short sales of stock or stock index options. Equity hedge portfolios may be anywhere from net long to net short depending on market conditions. Equity hedge managers generally increase net long exposure in bull markets and decrease net long exposure or even are net short in a bear market. Generally, the short exposure is intended to generate an ongoing positive return in addition to acting as a hedge against a general stock market decline. Stock index put options are also often used as a hedge against market risk. Profits are made when long positions appreciate and stocks sold short depreciate. Conversely, losses are incurred when long positions depreciate and/or the value of stocks sold short appreciates. Equity hedge managers' source of return is similar to that of traditional stock pickers on the upside, but they use short selling and hedging to attempt to outperform the market on the downside.

HFRX Equity Market Neutral Index - HFRXEMN

Equity Market Neutral strategies strive to generate consistent returns in both up and down markets by selecting positions with a total net exposure of zero. Trading managers will hold a large number of long equity positions and an equal, or close to equal, dollar amount of offsetting short positions for a total net exposure close to zero. A zero net exposure is referred to as "dollar neutrality" and is a common characteristic of all equity market neutral managers. By taking long and short positions in equal amounts, the equity market neutral manager seeks to neutralize the effect that a systematic change will have on values of the stock market as a whole.

Some, but not all, equity market neutral managers will extend the concept of neutrality to risk factors or characteristics such as beta, industry, sector, investment style and market capitalization. In all equity market neutral portfolios, stocks expected to outperform the market are held long, and stocks expected to under perform the market are sold short. Returns are derived from the long/short spread, or the amount by which long positions outperform short positions.

HFRX Event Driven Index - HFRXED
Event Driven investment strategies or "corporate life cycle investing" involves investments in opportunities created by significant transactional events, such as spin-offs, mergers and acquisitions, industry consolidations, liquidations, reorganizations, bankruptcies, recapitalizations and share buybacks and other extraordinary corporate transactions. Event driven trading involves attempting to predict the outcome of a particular transaction as well as the optimal time at which to commit capital to it. The uncertainty about the outcome of these events creates investment opportunities for managers who can correctly anticipate their outcomes. As such, event driven trading embraces merger arbitrage, distressed securities, value-with-a-catalyst, and special situations investing.

Some event driven trading managers will utilize a core strategy and others will opportunistically make investments across the different types of events. Dedicated merger arbitrage and distressed securities managers are not included in the event driven index. Instruments include long and short common and preferred stocks, as well as debt securities, warrants, stubs, and options. Trading managers may also utilize derivatives such as index put options or put option spreads, to leverage returns and to hedge out interest rate and/or market risk. The success or failure of this type of strategy usually depends on whether the trading manager accurately predicts the outcome and timing of the transactional event. Event driven trading managers do not rely on market direction for results; however, major market declines, which would cause transactions to be repriced or break, may have a negative impact on the strategy.

HFRX Macro Index - HFRXM
Macro strategies attempt to identify extreme price valuations in stock markets, interest rates, foreign exchange rates and physical commodities, and make leveraged bets on the anticipated price movements in these markets. To identify extreme price valuations, trading managers generally employ a top-down global approach that concentrates on forecasting how global macroeconomic and political events affect the valuations of

financial instruments. These approaches may be systematic trends following models, or discretionary. The strategy has a broad investment mandate, with the ability to hold positions in practically any market with any instrument. Profits are made by correctly anticipating price movements in global markets and having the flexibility to use any suitable investment approach to take advantage of extreme price valuations. Trading managers may use a focused approach or diversify across approaches. Often, they will pursue a number of base strategies to augment their selective large directional bets.

HFRX Merger Arbitrage Index - HFRXMA
Merger Arbitrage, also known as risk arbitrage, involves investing in securities of companies that are the subject of some form of extraordinary corporate transaction, including acquisition or merger proposals, exchange offers, cash tender offers and leveraged buy-outs. These transactions will generally involve the exchange of securities for cash, other securities or a combination of cash and other securities. Typically, a manager purchases the stock of a company being acquired or merging with another company, and sells short the stock of the acquiring company. A manager engaged in merger arbitrage transactions will derive profit (or loss) by realizing the price differential between the price of the securities purchased and the value ultimately realized when the deal is consummated. The success of this strategy usually is dependent upon the proposed merger, tender offer or exchange offer being consummated.

When a tender or exchange offer or a proposal for a merger is publicly announced, the offer price or the value of the securities of the acquiring company to be received is typically greater than the current market price of the securities of the target company. Normally, the stock of an acquisition target appreciates while the acquiring company's stock decreases in value. If a manager determines that it is probable that the transaction will be consummated, it may purchase shares of the target company and in most instances, sell short the stock of the acquiring company. Managers may employ the use of equity options as a low-risk alternative to the outright purchase or sale of common stock. Many managers will hedge against market risk by purchasing S&P put options or put option spreads.

HFRX Relative Value Arbitrage Index - HFRXRVA
Relative Value Arbitrage is a multiple investment strategy approach. The overall emphasis is on making "spread trades" which derive returns from the relationship between two related securities rather than from the direction of the market. Generally, trading managers will take offsetting long and short positions in similar or related securities when their values, which are mathematically or historically interrelated, are temporarily distorted. Profits are derived when the skewed relationship between the securities returns to normal. In addition, relative value managers will decide which relative value strategies offer the best opportunities at any given time and weight that strategy accordingly in their overall portfolio. Relative value strategies may include forms of fixed income arbitrage, including mortgage-backed arbitrage, merger arbitrage, convertible arbitrage, statistical arbitrage, pairs trading, options and warrants trading, capital structure arbitrage, index rebalancing arbitrage and structured discount convertibles (which are more commonly known as Regulation D securities) arbitrage.