Crisis and Non-Crisis Risk in Financial Markets: A Unified Approach to Risk Management

Robert H. Litzenberger and David M. Modest*

December 5, 2007

1 Introduction

The Hunt Brothers silver crisis of 1979/80; the U.S. savings and loan crisis in the 1980s; the Mexican Default and the Latin American Debt Crisis starting in 1982; the failure of Continental Illinois in 1984; the Bank of New York systems failure resulting in a $24 billion overnight overdraft at the Federal Reserve Bank of New York in 1985; the Stock Market Crash of 1987; the equity market and property price collapse in Japan and the bankruptcy of Drexel Burnham in 1990; the Salomon Brothers treasury scandal in 1991; the Metalgesellshaft heating oil trading losses in 1993; the U.S. and European bond market crashes of 1994; the Orange County derivatives losses in 1994; the Mexican devaluation of the peso and the beginning of the Tequila crisis in 1994; the Barings failure and Daiwa trading scandal in 1995; the Sumitomo copper metal trading scandal in 1996; the Asia Crisis of 1997; the Russia and Long-Term Capital Management Crises in 1998; the dramatic stock market drop in the wake of 9/11; the Enron bankruptcy in 2001; the Allied Irish Bank trading losses in 2002; the Refco bankruptcy in 2005; the rapid demise of the hedge fund Amaranth in 2006; the sub-prime, credit, liquidity, and quantitative equity crises of 2007: the litany of financial crises and economic losses caused by failed financial institutions during

*Robert Litzenberger is the Edward Hopkinson Jr. Professor Emeritus of Investment Banking at The Wharton School, University of Pennsylvania and an Executive Director at Azimuth Asset Management where he manages a portfolio of hedge funds. He is a former Partner at Goldman Sachs and global head of Goldman Sach's firm-wide risk management. David Modest is a Managing Director at J.P. Morgan Chase in their Proprietary Positioning Business. He is a former Managing Director at Azimuth Trust, partner at Long Term Capital Management (LTCM) and professor at the Haas School of Business, University of California. Litzenberger was part of the Goldman Sachs team that examined LTCM's positions prior to the consortium investment in 1998. This paper reflects our practical experience in financial markets including the events of 1998, and our years as academics. The risk management framework discussed in the paper was developed during our time together at Azimuth Trust in 2003. We are grateful and indebted to all of our Azimuth Trust colleagues for their help in developing the model especially Alex Shapiro and Paul Toldalagi. We would like to thank Amy Litzenberger for editorial assistance, and René Stulz and Darrell Duffie for helpful comments.

the last quarter century has given a major impetus to the design, development, and implementation of robust enterprise-wide risk management systems.[1]

Value-at-risk (*VaR*) is currently the most popular risk metric used by global financial institutions to report their firm-wide risk exposure.[2] *VaR* is an estimate of the loss threshold such that at a designated confidence interval, $1-\alpha$, the probability of a loss greater than the threshold, over a specified horizon, is equal to α (e.g. 1% or 5%). There are two main alternative methods used for computing *VaR*: a parametric approach and a non-parametric approach. The former is based on the estimated standard deviation of the current portfolio and a parametric assumption about the distribution generating future returns. The commonly used assumption of normality simplifies the analysis since the sum of normally distributed random variables is normal, and hence the procedure works equally well with individual securities and portfolios. However, this approach does not reflect the empirical observation that returns have fat left tails. The non-parametric approach takes the current portfolio and generates a history of what the profit and loss for this portfolio would have been over a specified past period. To compute the appropriate *VaR*, one then reads off the relevant percentile from the constructed hypothetical historical P&L distribution. In general, the non-parametric approach is also unable accurately to reflect fat tails because of the relatively short data histories used. Due to these limitations, both approaches are often supplemented with stress tests and scenario analyses.

VaR approaches are also deficient in that future price movements often depend on the context of the positions. For instance, consider two stocks involved in a merger arbitrage deal. The future price behavior of a position long the target and short the acquirer is likely to have little to do with how the isolated past return series behaved prior to the merger announcement. One additional shortcoming of the *VaR* framework is that it does not reflect the actual magnitude of the losses in the lower tail. Expected Tail Loss (ETL) is therefore a better measure of downside risk than *VaR* since it accounts for the distribution of losses in the lower tail. It measures the expected loss conditional on the loss being greater than the specified α loss-threshold.[3] The use of ETL coupled with the assumption of normally distributed returns merely applies a multiplier to the standard deviation to generate the risk metric. Hence, its value as a risk measure hinges on the non-normality of returns and is especially beneficial in the presence of fat left hand tails.

In response to large losses on proprietary European bond positions in 1994, Goldman Sachs developed an enterprise-wide internal risk management system

[1] In their classic book, Kindleberger and Aliber [10] explore scores of global financial crises that have occurred over the last four centuries. Each has its own distinctive flavor, yet they share important traits. They are all marked by dramatic losses of financial wealth and many are accompanied by the ruin of major financial institutions. Despite their regular occurrence, they often appear decades apart and have affected a disparate collection of assets including stocks, real estate, tulips, canals, commodities, bonds, railroads, precious metals, office buildings, foreign exchange and golf courses.

[2] Jorion [8] provides a comprehensive discussion of VaR.

[3] ETL also has the advantage of being a coherent measure of risk in the sense of Artzner, Delbaen, Eber and Heath [1] whereas *VaR* is not a coherent risk measure.

that aggregated firm-wide trading positions and measured daily *VaR* based on historical variances and covariances between asset returns — a paramteric *VaR* approach.[4] In an attempt to capture gradual changes in variances and covariances, exponential declining weights were used in the estimation of the second moments. To account for fat tails, the 99.6% percentile of a mixed normal distribution was used to compute the loss threshold such that each year the daily loss would be expected to exceed the threshold not more than once.[5] The simple intuitive lesson learned from the stock market crash of 1987 and the bond market crash of 1994 was that firm-wide directional exposure was particularly risky. Relative value trades, on the other hand, with their mean reverting properties and very low (or zero) historical correlations[6], had very attractive expected return / risk characteristics under most backward looking *VaR* paradigms. The result was a global proliferation of a wide range of relative value trades — with *VaR* based risk limits implemented at the trader, desk, divisional and firm-wide levels.

Relative value trading in fixed income and equity markets was quite profitable over the ensuing few years and proprietary trading positions of investment banks and hedge funds grew rapidly. Long-Term Capital Management (LTCM), for instance, generated 28.15% gross returns for ten months in 1994 after their launch; 58.77% in 1995, and 57.47% in 1996. LTCM's spectacular performance and their notoriety spurred many imitators and capital flocked to these types of trades. Although different models were used to identify attractive trades, very similar positions were established across many different hedge funds and proprietary trading desks. Price aberrations in the cash markets, that created relative value trading opportunities, were quickly arbitraged away by the large pool of capital dedicated to relative value trades. This resulted in a substantial decline in the realized volatilities and the profitabilities associated with both individual relative value trades and portfolios of relative value trades.

LTCM's gross returns, for example, declined to 25.28% in 1997; and the monthly standard deviation of their returns declined to 1.64% in 1997 compared to 3.46% in 1994.[7] With the benefit of 20/20 hindsight, the reduction in volatilities that occurred in mean reverting trades was a sign of crowded trades, which indicated an *increase* rather than a decrease in catastrophic risk. At the end of 1997, LTCM returned \$2.7 billion of capital to its investors — leaving it with an equity capital base of \$4.67 billion.[8] At the time, LTCM's Risk Aggre-

[4]While the aggregation of trading exposures across a wide range of trading instruments was a major accomplishment, the shortcomings of this system were exposed by the relative value hedge fund crisis in 1998.

[5]For the purposes of its annual report, Goldman Sachs reports the *VaR* at the 95% level for a one-day horizon based on a normal distribution. J.P. Morgan (JPM) began reporting *VaR* in its annual report beginning in 1994. In contrast to Goldman, JPM reports its *VaR* based on the non-parametric approach using the current positions and an historical simulation of changes in market value over the past twelve months (J.P. Morgan Chase & Company 2005 annual report [9], pp. 76.)

[6]To each other and to directional positions.

[7]The monthly standard deviation of returns was 1.80% in 1995 and 2.68% in 1996.

[8]Thirty-seven percent of its capital was returned in December 1997. The reduction in the volatility of the relative value trades, in part a reflection of the crowding of the trades,

gator (a parametric form of a *VaR* model using input parameters that were in part based on historical data and in part based on expectations about future risk) forecasted its daily standard deviation of P&L at roughly \$60 million.[9] Given the capital base of \$4.67 billion, this was equivalent to an annual standard deviation of approximately 20%.[10] As of January 30, 1998, the realized daily volatility of percentage changes in the S&P 500 index during the previous 252 trading days was approximately 19%. Thus LTCM's forecasted and realized daily volatilities were quite similar to that of the S&P 500. However, due to the complicated and multi-legged structure of many of its trades, the distribution of future potential volatilities was much wider than the S&P 500.

In 1998, a series of events including the voluntary liquidation of the Salomon Brothers arbitrage unit at Citibank, the Russian devaluation of the ruble and default on its internal debt, and a large widening in credit spreads led to losses and increased volatilities of relative value trades.[11] The proliferation of risk management paradigms based on *VaR* measures, implemented with recent historical return data, had an impact on losses on relative value trades that was analogous to the impact of portfolio insurance on equity price declines in 1987. Initial losses on relative value trades and increases in volatility caused *VaR* risk limits to be violated and caused further liquidations of relative value trades. Seemingly unrelated trades, that were uncorrelated historically, became highly correlated simply because they were being liquidated simultaneously.[12] The markets for the various components of relative value trades became less liquid and some actually seized up for a period of time. Since firms: (i) did not view themselves as price-takers, (ii) did not know the positions of other firms or their liquidation plans, and hence (iii) did not want to *bang the market*, positions were slowly liquidated which resulted in serial correlations and losses extending over multiple months. Trades that Goldman Sachs held in common with Long-Term Capital Management and likely with many other firms experienced two and threefold increases in volatility, correlations between returns on these trades increased substantially, and these positions experienced large cumulative losses over a three month period.

undoubtedly made LTCM more comfortable returning the capital than if volatility had increased.

[9] LTCM's Risk Aggregator assumed that positions at shorter horizons, such as a month, were more highly correlated than trades at longer horizons, such as over the course of the year. In December 1997, the Aggregator predicted a daily standard of \$81 million assuming the higher correlations and a daily standard deviation of \$60 million at the lower correlations. Until the crisis in the Fall, LTCM's Risk Aggregator consistently predicted higher risk than was actually realized. For instance, in January 1998, the realized daily standard deviation of P&L was \$41 million. This information is taken from Modest [7].

[10] And a realized annual standard deviation of approximately 14% based on the daily realized volatility of \$41 million in January 1998.

[11] LTCM, for instance, for the first time loss money for consecutive months in May and June of 1998.

[12] The Brunnermeir and Pedersen [5] model captures many aspects of this type of phenomenon. Chan, Getmansky, Haas and Lo [6] provide insights on how the serial correlation of hedge fund returns can be used to provide insights on the illiquidity of the underlying positons.

In the summer of 2007, losses in the sub-prime credit market, extreme movements in credit spreads, and the steep declines in prices of many buy-out related equities led to de-levering by several large hedge funds and proprietary trading desks and ultimately to forced liquidations of positions held by many market-neutral quantitative equity strategies. This resulted in price movements far in excess of what could have been reasonably predicted based on past historical data. The extreme magnitudes of the relative price moves in early August are reflected, for example, in the net asset values for the publicly traded: Highbridge Statistical Market Neutral (HSKAX) mutual fund. This fund, between December 1, 2005 and the December 31, 2006, experienced a daily standard deviation of returns of 0.161%. During the seven days between August 2, 2007 and August 10, 2007, the fund had returns equal to: -0.25%, -0.75%, -0.81%, -1.26%, -2.30%, -2.09% and +2.14%. These correspond to the following moves in units of historical standard deviation: -1.54, -4.62, -5.04, -7.82, -14.25, -12.97 and +13.25. Goldman Sachs CFO, David Viniar, in commenting on the August 2007 experienced by the firm's flagship quantitative hedge fund, *Global Alpha*, stated that "We were seeing things that were 25-standard deviation moves, several days in a row."[13] The historical probability distributions in non-crisis periods, that serve as the basis of Viniar's comments and were estimated with many decades of daily historical data, have little or no relevance to this crisis period. The observed daily returns in August were generated by the simultaneous liquidations of similar trading positions by many levered quantitative equity funds and proprietary trading desks. A wide range of seeming unrelated quantitative equity strategies became correlated because they were being liquidated at the same time. Liquidations increased in intensity from August 2nd through August 9th as some highly levered hedge funds were desperate to de-lever and some proprietary trading desks were ordered to liquidate positions to reduce the firms' franchise risks. The post liquidation recovery only partially mitigated the losses of funds and desks that dramatically reduced their positions at market lows.

Hence, an important lesson from the 1998 and 2007 relative value crises is that variances, covariances, and serial correlations estimated from recent historical return data are misleading indicators of potential losses in *trade-driven* financial crises. Losses can be endogenously caused by crowded trades and the reaction of traders to initial trading losses and increases in market volatility. In retrospect, this should also have been a lesson learned from the 1987 stock market crash (where portfolio insurance was a crowded trade) and the 1994 bond market crash (where the long bond carry trade was a crowded trade). The response of Goldman Sachs and other firms to their experience in 1998 was to place greater reliance on stress tests and scenario analysis over longer time horizons in managing trading risks. For example, a credit spread widening scenario over a three month horizon was used to set risk limits for Goldman Sachs credit sensitive fixed income positions. The process of establishing trading limits based on stress testing credit spreads established a risk culture at

[13] "Goldman pays the price of being big", *Financial Times*, August 13, 2007.

Goldman Sachs that controlled its exposure to the subprime mortgage crisis in the summer of 2007. Unfortunately, such elementary risk controls were apparently not in place at Merrill Lynch, which wrote down subprime mortgages by $7.9 billion, and Citgroup, which stated in November 2007 that it might suffer a write-down for subprime losses of $12 billion. The Chairman and CEO of Merrill Lynch, Stanley ONeal resigned and was replaced by John Thain, who as the CFO at Goldman Sachs, encouraged the use of stress test limits for fixed income securities in the third quarter of 1998. The chairman of CEO of Citigroup, Charles Prince, also resigned. Nevertheless, firms continue to use *VaR* implemented with historical variances and covariances because of the analytical tractability of this model in aggregating risk across different types of trades; and its mechanistic appeal to regulators.

The present paper builds on these lessons to develop an analytically tractable risk management metric that more accurately measures potential exposures to financial crises and also captures volatility during non-crisis times. We develop a multiple-regime stress-loss risk framework that assumes markets are characterized by quiescent (non-crisis) periods most of the time; interspersed with infrequent crisis periods where 4-5 sigma events can occur with non-negligible probabilities. The framework is flexible and can incorporate an arbitrary number of crises. One of the primary lessons of 1998 and 2007 is that returns can be correlated due to the capital underlying a collection of trades (or strategies), regardless of any underlying economic rationale. This is an important feature of our model. We include crises that are directional in nature and capture severe directional moves such as those which occurred in 1994 and 1987, and we incorporate crises that capture strategy-based (or trade-based) crises such as occurred in 1998 and 2007.[14]

A crisis is associated with a negative strategy return and a dramatic increase in volatility where the magnitude of the negative shock depends on many factors including the liquidity of the instruments underlying the strategy, the aggregate size of the positions pursuing the strategy, the crowdedness of the trades in the strategy, and the trade complexity. When a crisis occurs, all investments with exposure to an affected strategy are impacted causing correlations of returns to tend towards one during crisis periods — an empirical feature that any realistic risk model must capture. Our model also has the desired features: (i) it can accurately predict portfolio volatility during normal, non-crisis times, (ii) it captures in a realistic fashion stress moves during crisis periods, (iii) it is consistent with the empirical observation that returns in many financial markets are characterized by distributions with fat left tails, and (iv) it provides a unified framework for comparing trades whose risk may be composed of different proportions of typical day-to-day non-crisis volatility and crisis risk.

The remainder of the paper is organized as follows. Section 2 presents our

[14]It could be argued that the stock market crash of 1987 was also trade-based, due to the the crowding of the portfolio insurance *trade*, and not based on any change in underlying fundamentals. The bond market crash of 1994 came on the heels of six straight increases in short-term interest rates by the Federal Reserve, but was exacerbated by the proliferation of the long bond carry trade.

model. In section 2.1, we characterize financial markets as consisting of quiescent or non-crisis volatility periods and periods of severe market stress, and define the set of feasible states. In section 2.2, we discuss the computation of expected tail loss and show that it can be expressed as the probability weighted average of state contingent put options. Section 3 shows how the model can be used to decompose the risk of a portfolio between crisis and non-crisis risk, and how to decompose the strategy (or individual asset) contributions to the two types of risk. The model is also used, in a Black Litterman spirit, to examine the expected returns that are consistent with a given portfolio allocation and how expected returns would have to change to justify a portfolio tilt away from an initial allocation. In section 4, we discuss the practical implementation of the model in the context of a fund of hedge funds manager. Section 5 concludes.

2 The Model

2.1 Crisis Risk and the State Space

In this paper, we assume that financial markets are characterized by quiescent, non-crisis periods, infrequently interspersed by crisis periods of severe market stress. Let us assume there are C types of financial crises. For C>2, the number of states, S, is equal to:

$$S = C + 2 + \sum_{k=2}^{C-1} \frac{C!}{k!(C-k)!} \tag{1}$$

For example if there are three types of crises, there are eight possible states: one no-crisis state (the quiescent period that occurs most of the time), three single-crisis states, one three-crisis state, and three two-crisis states.[15] In each state, returns are assumed to be generated by a state-dependent normal distribution where the state-dependent moments can vary over time. The state probabilities for a given period, π_s, can be chosen to be consistent with either independently-distributed crises or crises that are correlated. These probabilities can also vary over time, although in our implementation we assume they are constant.

Assets impacted by a given type of crisis have their state-dependent standard deviations increase by a crisis multiplier, and suffer a crisis-dependent downward mean shift equal to a Z-value times their state-dependent standard deviation. This downward mean shift is equivalent to a perfectly correlated component of returns. This will cause the correlation of all *assets* impacted by a particular crisis to tend towards one during the crisis, although they will not be perfectly correlated. In implementing the model, it is natural to think of an *asset* at the primitive level, such as an individual equity or specific bond. However, one of the important lessons of 1998 and 2007 is that forced liquidations can cause *trades* to be correlated and, in fact, send individual *assets* in directions counter

[15] For $C = 1$, $S = 2$: one no-crisis state and one crisis state. For $C = 2$, $S = 4$; one no-crisis state, one two-crisis state, and two single crisis states.

to that which would be predicted by the economics driving the security's cash flows.

Take the case of Royal Dutch and Shell Transport in 1998. The Royal Dutch - Shell Transport group of companies was created in February 1907 when the Royal Dutch Petroleum Company of the Netherlands and the Shell Transport and Trading Company Ltd of the United Kingdom merged their operations. The terms of the merger gave 60% of the new company to the Dutch shareholders and 40% to the British holders. Until July 20, 2005, the group was a dual listed company. Royal Dutch was a member of the S&P500 and traded primarily in New York and the Netherlands. Shell Transport traded primarily in London although ADRs did trade in New York. Royal Dutch tended to trade at a premium to Shell Transport (relative to the 60/40 split of earnings) in part because of its inclusion in the S&P500. LTCM had a $2.1 billion dollar arbitrage position: long Shell Transport and short Royal Dutch. In 1997, the average discount between the two sets of shares was less than 8.83%.[16] During LTCM's difficulties in the August-October 1998 period, the discount expanded to a peak of 18.53% – as the market became concerned LTCM would be forced to liquidate its position rapidly.[17] Thus the 1998 crisis resulted in two nearly identical assets moving in opposite directions. This phenomenon displayed itself again with great force in August 2007 when equities typically held long by market-neutral equity hedge funds plummeted and equities typically held short by the same funds rose dramatically — with overall market levels relatively unaffected. This suggests the need for a risk system to not only be focused on individual assets, but also on *trades*.

For the purposes of illustrating our model, in section four below, we assume that the probability of a given state occurring, π_s, is constant over time. A period in our model must thus be no shorter than the typical length of a crisis event — to insure that the probability of being in a crisis at date $t+1$ is independent of whether a crisis occurred at date t. We choose a period to equal a quarter. A quarter horizon seems sufficient to allow for the prolonged impact of a crisis on illiquid trades and also reflects the importance of the investment banks' quarterly reporting cycles. Figure 1 graphically illustrates the conditional state-dependent and unconditional returns distributions in a simple two-state stress loss framework.

[Figure 1 to go here]

The lighter colored dashed distribution is the standardized distribution *conditional* on no-crisis; the darker dotted distribution is the standardized distribution *conditional* on a crisis. The *conditional* crisis distribution reflects: (i) a downward mean shift to reflect the downward *perfectly correlated component* of

[16] The percentage spread is defined as the market value of Royal Dutch in USD less 1.5 times the market value of Shell Transport in USD divided by 1.5 times the market value of Shell Transport in USD. The calculation is done using closing ADR prices from New York.

[17] The peak spread was reached on October 8, 1998.

a crisis, and (ii) an increase in volatility relative to the non-crisis state. The distribution drawn in solid black is the unconditional distribution that is the probability weighted distribution over the no-crisis (quiescent) and crisis states.

2.2 Crisis Risk and Expected Tail Loss (ETL)

A natural measure of crisis risk in our model, as discussed above, is the expected tail loss, ETL at a pre-specified percentile level, $(1-\alpha)$, and over a designated horizon, τ.[18] The expected tail loss ETL at the 95% percentile level is depicted in Figure 2.

[Figure 2 to go here]

Our framework allows for the aggregation of ETL across assets and the measurement of the contribution of individual assets (or strategies) to portfolio $ETLs$. Under the assumption of normality within each state, the solution for the state-contingent ETL is similar to the well known formulation for the value of a put option when asset returns are normally distributed. The overall expected tail loss, ETL, is expressed as the probability weighted average of state contingent put options. The assumptions of state-dependent normal distributions combined with a finite number of crisis classifications result in a well-defined fat tail distribution as the number of asset increases. This contrasts with the use of mixed normal distributions with independent mixing probabilities across assets, where the limiting distribution is normal as the number of assets increases.

More formally, the Expected Tail Loss, ETL, on portfolio (or asset) p may be expressed as:

$$ETL \equiv \mathbf{E}[\tilde{R}_p \mid \tilde{R}_p \leq A] = \sum_{s}^{S} \pi_s \mathbf{E}[\tilde{R}_p \mid \tilde{R}_p \leq A, s] \tag{2}$$

where:

- S is the number of states,
- π_s is the probability of state s occurring,
- R_p is the portfolio rate of return in excess of the risk-free rate,
- $\mu_{p,s}$ is the portfolio mean rate of return in state s,
- $\sigma_{p,s}$ is standard deviation of the portfolio's return in state s, and
- A is a return threshold that depends on the percentile choice, α.[19]

[18] For instance, as noted in their annual reports, Goldman Sachs focuses on VaR at a $1-\alpha$ percentile level of 95% and a daily horizon; J.P. Morgan Chase reports its VaR at a $1-\alpha$ percentile level of 99% and the same daily horizon.

[19] Although A should more formally be written $A(\alpha)$, for simplicity of notation we write it as A. A is also an implicit function of the state dependent means and standard deviation as shown in relation (4) below. A also corresponds to what is called the $(1-\alpha)$ percentile VaR.

Substituting for the expected value operator, ETL can be re-written:

$$ETL = \sum_{s}^{S} \pi_s \int_{-\infty}^{A} x f(x, \mu_{p,s}, \sigma_{p,s}) dx/\alpha \tag{3}$$

where $f(\cdot)$ denotes the standard *state-dependent* normal probability density function of the normalized portfolio returns.

Analytically integrating the right hand side of equation (3) gives the closed-form solution for the expected tail loss — *conditional* on A, the α percentile return threshold for the mixed normal distribution:

$$ETL = \sum_{s}^{S} \pi_s \left[\mu_{p,s} F\left(\frac{A - \mu_{p,s}}{\sigma_{p,s}}\right) - \sigma_{p,s} f\left(\frac{A - \mu_{p,s}}{\sigma_{p,s}}\right)\right] /\alpha \tag{4}$$

where $F(\cdot)$ denotes the standard *state-dependent* normal cumulative distribution functions of the normalized portfolio returns.[20]

The threshold A can be estimated numerically using:

$$\sum_{s}^{S} \pi_s F\left(A, \mu_{p,s}, \sigma_{p,s}\right) = \alpha \tag{5}$$

The 1980s and 1990s were very turbulent times in global financial markets. Especially for those participants active in the markets, it seemed that hundred-year floods were occurring at least every five years. At Goldman Sachs, there was considerable discussion about the frequency of recurrence of such events for the purpose of determining the adequacy of the non-crisis profitability levels of different trading businesses. One quarter in five years (20 quarters) was subjectively determined as the basis for determining which trading activities were sufficiently profitable to continue. In this paper, we make the same assumption and assume a 5% probability of a given crisis in any quarter, $\alpha = 0.05$, in implementing the model.

3 Contributions of Individual Assets to Portfolio Risk

3.1 Marginal Contribution to Expected Tail Loss

Consider portfolio $p^{\dagger}$ that has a portfolio weight of w_i in the i^{th} asset and $(1-w_i)$ in portfolio p. The state dependent mean, $\mu_{p^{\dagger},s}$,and standard deviation, $\sigma_{p^{\dagger},s}$ of this new portfolio are:

$$\mu_{p^{\dagger},s} = w_i \mu_{i,s} + (1 - w_i)\mu_{p,s} \tag{6}$$

[20]This derivation is similar to Brennan's [4] derivation of the value of a call option under normally distributed returns and exponential utility. Under risk neutrality, the product of alpha and the ETL formula is the forward value of a put option that exercises when the price is below A, but does not require the payment of A.

and

$$\sigma_{p^\dagger,s} = \left[w_i\sigma_{i,s}^2 + (1-w_i)\sigma_{p,s}^2 + 2w_i(1-w_i)\sigma_{ip,s}\right]^{\frac{1}{2}} \tag{7}$$

where $\sigma_{ip,s}$ denotes the state-dependent covariance between the returns on asset i and the returns on portfolio p.

To evaluate the marginal impact of the i^{th} asset on the ETL of the portfolio, equation (4) is differentiated with respect to w_i subject to the constraint given by (5) and evaluated at $w_i = 0$. The following expression defines the sensitivity of the portfolio's expected tail losses to changes in the portfolio weights:

$$\begin{aligned} \left.\frac{dETL}{dw_i}\right|_{w_{i=0}} &= \frac{1}{\alpha}\sum_s^S \pi_s \Bigg\{ \left(\frac{d\mu_{p^\dagger,s}}{dw_i}\right) F(\cdot) + \mu_{p^\dagger,s}\left(\frac{dF(\cdot)}{dw_i}\right) \\ &- \left(\frac{d\sigma_{p^\dagger,s}}{dw_i}\right) f(\cdot) - \sigma_{p^\dagger,s}\left(\frac{df(\cdot)}{dw_i}\right) \Bigg\}\Bigg|_{w_{i=0}} \end{aligned} \tag{8}$$

where the derivatives with respect to the portfolio weights are evaluated at $w_i = 0$. The differentiation of equation (4) with respect to w_i is a closed form solution holding A constant. While A must be solved for numerically based on constraint (5), the implicit differentiation of A based on (5)can be done analytically as shown in the Appendix.

The fractional contribution of the i^{th} asset to the ETL of portfolio $p^\dagger$ is equal to the product of its portfolio weight, w_{ip}, and its *expected tail loss* beta with respect to portfolio $p^\dagger$ which is given by:

$$\beta_{i,p^\dagger}^{ETL} = 1 + \frac{\frac{dETL}{dw_i}|_{w_{i=0}}}{ETL} \tag{9}$$

The calculation of fractional contributions to portfolio ETL is analogous to the calculation of the marginal contribution of the i^{th} asset to the variance of a portfolio as the product of its portfolio weight, w_{ip}, and its volatility beta, $\beta_{i,p^\dagger}(\sigma^2)$, where:

$$\beta_{i,p^\dagger}^{\sigma} = 1 + \frac{\frac{d\sigma_{p^\dagger}}{dw_i}|_{w_{i=0}}}{\sigma_{p^\dagger}} \tag{10}$$

However, unlike the more familiar volatility beta, the ETL beta can not be estimated as a slope coefficient of a least squares regression.

3.2 The Optimal Portfolio and Implied Risk Premia

Consider the optimization problem of an investor whose utility function trades off expected excess return and expected tail loss, and poses her investment decision as a series of one period problems.[21] Since a portfolio's risk premium

[21] In practice, some asset owners may care about both expected tail loss and portfolio volatility during normal times. Investment banks, for instance, typically care about expected tail loss — as these can threaten the viability of the firm's future. They also care about expected quarter to quarter volatility of trading profits and losses, since equity analysts use this as a gauge of management's risk management prowess. In this paper, for simplicity, we assume that ETL is the sole relevant measure of portfolio risk.

(R_p) and ETL per dollar of equity are proportional to leverage, the *unlevered* portfolio with the highest ratio of risk premium to ETL when combined with borrowing or lending is optimal. This is depicted in figure 3, where the tangency portfolio is denoted by p^*.

[Figure 3 to go here]

The contributions of individual assets to the portfolio's ETL are used to derive the necessary and sufficient conditions for maximizing this reward to risk ratio. The first order conditions for an optimal portfolio are:

$$\frac{E_i(p^*)}{\beta_i^{ETL}} = \frac{E_j(p^*)}{\beta_j^{ETL}} = E_{p^*} \qquad \forall i, j \tag{11}$$

where p^* is the optimal tangency portfolio and the *crisis* betas reflect marginal contributions of the individual assets contribution to the portfolio's ETL, and $E_i(p^*)$ is the expected return on asset i that is consistent with portfolio p^* being the optimal portfolio.

It is well known that an unconstrained mean-variance optimizer is highly sensitive to small differences in expected returns in that relatively these differences can produce portfolios with extreme weights.[22] Hence, solving for optimal portfolio weights is not a well-posed problem in that the results are highly unstable. Black and Litterman [3] proposed an alternative approach that determines the necessary return forecast to justify a portfolio tilt. The approach of the current paper is in the spirit of their paper. Let p represent a base or reference portfolio. Given the choice of portfolio p, the first order condition for the optimal portfolio weights in equation (11) imply the risk premia for each asset that are necessary for the portfolio weights to be optimal.

Possible choices for the base portfolio p include the market portfolio of all assets (including hedge funds and other alternative assets), a hedge fund market portfolio such as the Morgan Stanley Capital International (MSCI) asset-weighted hedge fund index, or a fund of hedge fund's existing portfolio. The first choice is appropriate for analyzing the portfolio of a high net worth investor and the first order conditions may be interpreted as market equilibrium conditions under the assumption of identical beliefs across investors. This gives a CAPM in Expected Return - ETL space. The second choice is appropriate for an enhanced index fund of hedge funds that views its mandate as determining hedge fund strategy tilts from the MSCI strategy weights that would increase the portfolio's risk premium per unit of ETL, E_p/ETL, analogous to the Sharpe ratio in Mean-Standard deviation space.[23] The final choice is appropriate for a fund of hedge funds manager desiring to increase the E_p/ETL ratio through reallocating either hedge fund strategy weights or hedge fund manager weights.

[22] See, for instance, Best and Grauer [2].

[23] The risk premium refers to the expected return over the relevant risk free rate.

Consider a candidate portfolio p^* that a manager is considering as an alternative to the base portfolio p. For p^* to be a portfolio optimum, $E_i(p^*)$ has to equal the Bayesian posterior estimator of the risk premium for asset i. Thus, the portfolio equilibrium conditions on risk premia at the new equilibrium are the implicit posterior Bayesian estimator. This estimator has to also be the minimum variance convex combination of the Bayesian prior, $E_i(p)$ (which is equal to the original risk premium implied by the original base portfolio being an equilibrium portfolio), and the forecasted risk premium, $E_i(f)$ (Sample estimator). Assuming that the forecast error of the prior and the sample estimators are independent the Bayesian posterior estimators may be expressed as:

$$E_i(p^*) = \frac{E_i(p) + k^2 E_i(f)}{(1+k^2)} \qquad \forall \quad i \tag{12}$$

where k is the confidence in the forecasted risk premium (sample estimator) relative to the base equilibrium risk premium (prior estimator), which is measured as their relative standard error of forecasts.[24] When $k = 0$ and there is zero confidence in the forecasted risk premium, the posterior estimate is simply the base equilibrium / prior risk premium. For $k = 1$ equal confidence is placed in the forecasted risk premium and the prior equilibrium risk premium and they are equally weighted. Similarly, as $k \to \infty$ and the confidence in the prior equilibrium risk premium, $1/k$ approaches zero, the posterior estimate converges to the forecasted (sample) risk premium.

Since the implied equilibrium risk premia for both the base portfolio and the reallocated portfolio are known from the first order conditions and the base portfolios risk premium/ ETL ratio, the equation may be solved for the implied forecasted risk premia necessary for the forecast to be consistent with the new portfolio being an equilibrium portfolio.

$$E_i(f) = E_i(p^*) + \frac{1}{k^2}\left[E_i(p^*) - E_i(p)\right] \qquad \forall \quad i \tag{13}$$

4 Fund of Hedge Funds Implementation

This multiple-state stress-loss model was developed in 2003 during our time together at Azimuth Trust, a hedge fund of funds. The need for a model like this is apparent from a cursory examination of the fat left tails associated with hedge fund strategy returns. Consider, for instance, the historical volatility of returns for the fixed income arbitrage strategy. During the period from 1990-2006, this strategy had annual volatility of returns equal to 4.15%.[25] Coupled with the assumption of normally distributed returns, this implies a maximum predicted three-month loss of -4.30% during any given five-year period. This

[24]Let $\sigma_i(f)$ denote the standard error of the sample / manager forecast and $\sigma_i(p)$ denote the standard error of prior forecast. k is then given by $(1/\sigma_i(f))/(1/\sigma_i(p) = \sigma_i(p)/\sigma_i(f)$.

[25]The data are taken from *Hedge Fund Research* for the January 1990 - May 2006 period. The annual volatility of the returns presumes the monthly returns are serially correlated over time.

compares with the -13.96% that was experienced in the Fall of 1998. The returns for credit and distressed security hedge fund strategies show a similar pattern: annual volatility of -5.94%, a predicted maximum three month loss of -6.14% based on normally distributed returns, and a realized worst three-month loss of -12.82%.

While the model presented above is sufficiently general to handle an arbitrary number of types of crises, in practice the implementation of the model is more straightforward with a small number of crises. For example, Azimuth Trust implemented the model with three types of crises: (1) a relative value hedge fund crisis, (2) an equity market collapse, and (3) a macro-economic hedge fund crisis. Our approach at Azimuth Trust was to use this framework to determine strategy allocations to the different hedge fund strategies. Once strategy allocations were determined, we attempted to find the best collection of funds that would give us the desired strategy exposures.

Each hedge fund that is affected by a given type of crisis has it standard deviation go up by a fund specific Crisis Multiplier.[26] Non-affected funds' standard deviations are unchanged.[27] The portfolio of hedge funds state dependent standard deviation is calculated using the hedge funds state dependent standard deviations and an unchanged correlation matrix. The *perfectly correlated* component of a crisis is modeled as a downward mean shift for each affected fund that is equal to: 1.645 times its state dependent crisis standard deviation. The portfolio mean is calculated by weighting the non-crisis means of unaffected funds and the crisis means for funds affected by the crisis. The crisis multiplier for an individual hedge fund is subjectively determined based on many factors including: (1) the liquidity of the underlying instruments, (2) the size of its positions , (3) the crowdedness of the trades, (4) the complexity of its trades, (5) the amount of leverage used by the funds, and (6) the funds' risk management policy.

One of the important lessons from 1998 and 2007 is that hedges, which from an underlying fundamental perspective should be expected to reduce risk, can in crises serve to exacerbate trading losses. Consider the case of convertible bond arbitrage in 1998. One sensible way to hedge a long convertible bond position is to use: swaps to hedge the interest rate exposure, the underlying stock to hedge the net delta, index options to hedge exposure to aggregate changes in the level of volatility and a credit instrument to hedge the credit risk. In the fall of 1998, one would have been better off not hedging the aggregate vega exposure of convertible positions as the implied volatility of convertibles dramatically contracted while at the same time the implied volatilities of index

[26] Hedge funds may be impacted by more than one crisis. Consider a long/short equity hedge fund that operates with a net long exposure of 35%. This fund can be expected to be impacted by both a relative value hedge fund crisis, which will affect the hedged part of its portfolio, and an equity market collapse which will severely affect is net long exposure.

[27] It is sensible to think of an individual hedge fund's standard deviation as consisting of a systematic part that is a result of exposure to one or more strategies, and a fund-specific idiosyncratic part. The idiosyncratic part should reflect both idiosyncratic market risk factors (e.g. risk concentration), and operational and credit risk exposures.

options dramatically increased.[28] This is an example of why more complex trades can be expected to behave worse in a crisis and why they require a higher crisis multiplier. In a similar spirit, more crowded trades, trades that are larger in absolute size, and trades in more illiquid securities can also be expected to be subject to higher crisis losses and hence warrant a higher crisis multiplier — as forced liquidation for liquidity or risk reasons are more likely to lead to distressed selling.

One of the nice feature of our multiple-regime stress-loss framework is that it allows a decomposition of the risk of the portfolio by both contribution to ETL and contribution to non-crisis portfolio volatility in a consistent unified framework. This decomposition analysis is presented in Table 1. Overall, the portfolio has an expected volatility of returns of 4.5% per annum and an expected tail loss of -5.6%.[29] The table contains information on the hedge fund strategy weights for the portfolio[30], the strategy expected volatilities which are a subjective input to the model[31], the expected maximum three-month stress loss[32], the contribution of each strategy to the portfolio's ETL and volatility, the strategy ETL and volatility betas, and the set of expected excess returns (over Libor) that are consistent with this portfolio being optimally constructed.

[Table 1 to go here]

It is interesting to contrast the risk contributions of the global macro and fixed income arbitrage strategies to the portfolio's overall risk. In this example, 21.00% of the fund's assets are invested in global macro strategies. Although global macro strategies tend to have relatively high volatility, they tend to perform well in stress periods when asset prices tend to move in a sustained manner. Thus global macro accounts for 16.12% of the portfolio's expected tail loss compared to 23.99% of the portfolio's volatility during non-crisis periods. Fixed in-

[28] This reflected LTCM's significant long position in convertibles and its short position in index volatility, especially at the longer end of the volatility term structure.

[29] Given a volatility of 4.5% and the assumption of normally distributed returns, the expected tail loss would be -4.6%. As noted above, the ETL measures the actual magnitude of the expected tail loss whereas the traditional VaR measure only provides a loss threshold. Based on the multiple-state stress-loss model, the traditional 95% one-quarter VaR is -4.4%. Given an assumption of normality, the 95% one-quarter VaR would be -3.7%.

[30] These weights roughly corresponded to the asset-weighted weights in the MSCI hedge fund index as of January 2007. We have reallocated funds from their multi-process category into other hedge fund strategies. We assume that the global macro strategy is affected by its own macroeconomic crisis; long/short equity and event-driven strategies are affected by an equity market crisis; and the remaining strategies are impacted by a relative value hedge fund crisis.

[31] As discussed by Chan, Getmansky, Haas and Lo [6], the monthly returns of many hedge fund strategies exhibit significant serial correlation. Our subjective expected volatilities are a weighted combination of the standard volatility estimator and a synchronicity-adjusted volatility estimator.

[32] This is also an input to the model and depends on our subjective estimate of the strategy crisis multiplier. The factors that are important in determining the appropriate multiplier were discussed above.

come arbitrage behaves in a very different manner. As Table 1 illustrates, 5.75% of the fund's assets are invested in fixed income arbitrage strategies. This investment accounts for 6.57% of the portfolio's expected tail loss, but only 1.35% of the portfolio's volatility during non-crisis periods. The largest category of hedge funds in the *MSCI* index is long/short equity hedge funds which includes funds with a significant long bias, market-neutral funds, and funds with a persistent short bias.[33] Those type of funds account for 41.02% of the dollar portfolio allocation. Their contribution to *ETL* and portfolio volatility are 39.80% and 48.05% respectively.

The strategy *ETL* and volatility betas, as defined in equations (9) and (10) and also presented in Table 1, are a different way of representing strategy's contributions to the portfolio's *ETL* and standard deviation. The strategies with the two lowest *ETL* betas are statistical arbitrage (0.51) and global macro (0.77). Given its low $ allocation, a marginal increase in the allocation to this strategy has very little impact on the portfolio's overall *ETL*. A marginal increase in the allocation to global macro has relatively little impact on the *ETL* because it is impacted by its own macroeconomic crisis and provides diversification against relative value hedge fund and equity crises. The two strategies with the highest *ETL* betas are credit and distressed securities (1.38) and fixed income arbitrage (1.14). This reflects the size of their losses during relative value crises and the importance of relative value hedge fund crises in determining the portfolio's overall *ETL*. While fixed income arbitrage has one of the largest *ETL* betas, no strategy has a lower volatility beta (0.24) — reflecting its modest portfolio weight and its volatility during non-crisis periods. The big contributors to the portfolio's volatility are event-driven and merger arbitrage with a volatility beta of 1.21 and long/short equity with a volatility beta of 1.17. Global macro's volatility beta is also above average at 1.14 — reflecting a relatively high $ allocation and high volatility during non-crisis periods.

Given a set of assumptions about the non-crisis standard deviations, the correlations of strategy returns, the crisis multipliers, and a choice of base portfolio p; our framework can be used, as discussed above, to generate the set of implied expected returns that are consistent with the given portfolio weights being optimal. The results of this exercise are also presented in Table 1. The last column of numbers in Table 1 provides the implied expected excess returns that are consistent with the given strategy portfolio weights being optimal — given that the overall expected return on the portfolio is 6.4%. Fixed income arbitrage accounts for 5.75% of the dollar investment in this portfolio and 6.57% of the portfolio's expected tail loss. The table indicates that the implied expected excess return for fixed income arbitrage is 7.4% over the riskless rate (Libor). This expected return makes the 5.75% portfolio weight optimal, in the context of the other holdings, and assuming the manager makes her investment decision solely on the basis of expected excess return and *ETL*. If instead, the investor made her investment decision solely on the basis of expected excess return and volatility, the implied expected excess return for fixed income arbi-

[33] On average, funds in this category have a beta that is roughly 1/3 of the S&P500.

trage would be less than 100 basis point over Libor. Global macro is a strategy which provides good diversification during relative value hedge fund crises although it contributes significantly to the normal month to month volatility of portfolio returns (23.99%). For this portfolio construction to be optimal, this strategy must yield an expected excess return of 4.9% — which reflects its low contribution to the portfolio's ETL.

As discussed above in Section 3.2, this framework can be used to guide a fund of hedge funds manager who takes the MSCI hedge fund index as a benchmark and attempts to add value through strategy tilts and asset allocation. Table 1 presented, in the context of our risk model, the implied risk premia that are consistent with the MSCI hedge fund strategy weights being optimal. For portfolio p^* to be an optimal improvement over the benchmark, the implied risk premia for each strategy in the new portfolio ($E_i(p^*)$) has to equal the Bayesian posterior estimator of the risk premium for asset i given by equation (12). Tables 2 and 3 present the implied risk premia required to tilt away from the MSCI strategy weights under two different assumptions: (i) equal confidence in the forecasted risk premia (sample estimator) relative to the base equilibrium risk premia (prior estimator), i.e. $k = 1$, (Table 2); and where the confidence in the forecast is one-half the prior, i.e. $k = \frac{1}{2}$ (Table 3).[34]

The first row of Tables 2 and 3 present the implied risk premia required for tilting away from the MSCI strategy allocation of 21% to global macro. As expected, an investor should expected a higher expected return to justify a greater allocation to global macro and a lower expected return to justify a smaller allocation. As Table 2 indicates, increasing the strategy weight for global macro from 21% to 26.25% (a 25% increase in portfolio weight) is justified if an investor expects the excess return on the strategy to be 6.3% rather than 4.9% — assuming equal confidence in the forecast and in the prior. If the investor has less confidence in her forecast, the required expected risk premia must be higher. This is illustrated in comparing the corresponding risk premia in Tables 2 and 3. In table 3, where $k = \frac{1}{2}$, the risk premia required to justify a 25% increase in strategy weight is 8.5% — 220 basis points larger than the corresponding risk premium in Table 2.

The risk premia in Table 2 show that there is an important and intuitive relation between the changes in risk premia required to justify various portfolio tilts and: (i) the size of the base strategy allocation and (ii) the nature of its crisis exposure. In general, strategies with small portfolio allocations (and hence small impacts on the overall portfolio ETL) require relatively small changes in risk premia to justify tilts compared to strategies with larger allocations. For example, the strategy with the smallest base allocation is statistical arbitrage (3.2%). It only takes a 5 basis point risk premium change to justify either a 25% increase or decrease in its strategy allocation. Contrast this with the strategy with the largest allocation that is also affected by relative value hedge fund crises: credit and distressed securities arbitrage. Table 2 shows that it would

[34]When the portfolio weight is shifted away from a strategy, we assume it is proportionally reallocated to the other strategies.

require a 100 basis point increase in risk premia (from 8.9% to 9.9%) to justify a 25% increase in the allocation to the credit strategy and a 130 basis point decrease in risk premia (from 8.9% to 7.6%) to justify a 25% decrease. The largest MSCI strategy allocation is to the long/short equity category, 41.02%. A 25% increase in the allocation to this strategy requires a relatively large 150 basis point increase in risk premia (from 6.2% to 7.7%) and a corresponding decrease in allocation would require a 110 basis decrease (from 6.2% to 5.1%) in risk premia to justify the tilt.

The increases and decreases in risk premia required to justify the portfolio tilts in Table 3, when there is less confidence in the forecast relative to the prior, are significantly more pronounced than the corresonding changes in premia reported in Table 2. In fact, negative risk premia are required to justify 50% decreases in strategy allocations to global macro and long/short equity. Similar intuition drives both of these results. Global macro is the only strategy affected by a macroeconomic crisis in this implementation of the model. Since the probabilities of a given crisis occurring are assumed to be independent, global macro provides a very important diversification function for hedge fund investors. Hence to diminish its portfolio allocation by 50% (from 21% to 10.5%), one would have to expect the strategy to earn 320 basis points *less* than the riskless rate — otherwise it would be optimal to have greater exposure. Long/short equity has the largest allocation in the MSCI asset-weighted index, with an allocation of 41%. A decrease of 50% in this allocation and the concomitant reallocation to the other strategies, would dramatically increase the portfolio's exposure to both relative value hedge fund crises and macroeconomic crises. Hence an investor would have to expect a risk premium 470 basis point less than Libor to justify a portfolio tilt reducing exposure to long/short equity by 50%. The other numbers in Table 3 also show more dramatic swings than those in Table 2, but are less noteworthy.

Although it is beyond the scope of this paper, the same framework could be used to analyze the contribution of individual hedge funds to the portfolio's ETL and portfolio's volatility, and to the strategy's ETL and the strategy's volatility. In applying the framework to individual hedge funds, it is natural to think of some portion of an individual hedge fund's returns coming from a systematic component, the strategy returns, and some portion coming from an idiosyncratic component. This suggests, as mentioned above, that the funds' expected volatilities will depend on both the forecasted strategy and idiosyncratic volatilities. These idiosyncratic volatilities will presumably reflect fund-specific market, operational and credit risk factors.[35]

[35] In practice, idiosyncratic risk can be quite important given many of the notable hedge fund failures of the two decades such as: the Granite Fund (1994), Long-term Capital Management (1998), Manhattan Fund (2000), Maricopa Funds (2000), Lipper & Company (2002), Beacon Hill Asset Management (2002), Eifuku Master Fund (2002), Lancer Offshore Fund (2003), Millenium Partners (2003), Bailey Coates Cromwell Fund (2005), Aman Capital (2005), MotherRock (2006) and Amaranth (2006).

5 Concluding Remarks

In this paper, we have developed a multiple-regime stress-loss risk framework that incorporates the idea that financial markets are characterized by quiescent periods most of the time interspersed by occasional periods of crisis. The framework allows an arbitrary number of crises and permits crises that are both directional in nature and crises that may be associated with trades — building on one of the central lessons of 1998 and 2007: that positions may be correlated not only because of the underlying economic fundamentals but because the capital across trades is correlated. The crisis periods are characterized by a sharp decline in returns and periods of increased volatility. The returns distributions, conditional on a given crisis, are assumed to be normally distributed. The unconditional distribution of returns, which is a mixture of normal distributions, has fat left tails that characterize unconditional empirical returns distributions. The risk framework is able to capture, in a *unified* setting, non-crisis period volatility during typical periods and stress losses during crisis periods. This stands in sharp contrast to most practitioner risk management frameworks which use *VaR* to measure potential losses during normal periods and then supplement it with ad-hoc stress loss scenarios.

Our risk management framework requires calibration based on subjective *ex-ante* assessments of crisis probabilities and stress losses during crises. The rareness of crisis events makes it virtually impossible to estimate these probabilities exclusively based on past data. In addition, the changing nature of financial crises – in large part due to the rapid pace of innovation in financial markets — makes it extremely difficult to model losses during crisis periods solely on the basis of past data and without regard to the structure of current *trades*.[36] For an investor in hedge funds, the potential losses are likely to depend on many factors including the liquidity of the financial assets underlying the strategies, the aggregate position sizes pursuing the strategy, the crowdedness of the trades, and the trades' complexity.

One natural measure of risk in our framework is expected tail loss ETL. We show that for an investor who trades off expected excess returns and ETL, the first-order portfolio optimality conditions require the ratio of expected excess return to ETL beta to be equal for each asset in the portfolio. In a world of identical investors, an ETL-CAPM would hold rather than the usual Sharpe-Lintner-Mossin CAPM. We apply the framework to the portfolio problem of a fund of hedge funds investor, and show how the framework can be used to decompose the risk of the portfolio and derive the implied expected excess returns for each hedge fund strategy that is consistent with a set of portfolio weights being optimal. We also analyze how an investor can combine subjective return forecasts with prior beliefs implied by a benchmark portfolio to tilt strategy allocations in an optimal and internally consistent manner.

In this implementation, we assume that a period is a quarter, the occurrence

[36]The philosophy underlying our risk management framework thus contrasts sharply with the extreme value theory literature which seeks to improve estimation of extreme tail probabilities solely based on historical asset returns.

of different types of crises are independent, the probability of a crisis is time and state invariant, the state dependent moments of the distribution are constant over time, and there are three types of crises. However, the framework is flexible, allows for multiple crises, and the probabilities of the states of the world and the moments of the returns distributions can be made to be time varying. We also presumed that the manager of the fund, in constructing her optimal portfolio, poses her investment problem as a series of one-period decisions in which she trades off expected excess return and expected tail loss. In practice, the manager might also care about the volatility of her portfolio — since the month to month volatility is likely to impact her ability to attract capital.[37] In this case, the framework could be extended to allow the investor to trade off expected excess return and a weighted average of ETL and non-crisis volatility. Another interesting extension would be to consider the problem faced by a pension fund manager who seeks to construct an optimal allocation to hedge funds and other alternative assets to complement her exposure to traditional stocks and bonds. We leave this for further work.

[37]This would be similar to the situation faced by the top management of an investment bank. Although losses during crises periods could threaten the firm's future viability, quarter-to-quarter volatility is tracked by stock analysts as an indicator of management's risk control prowess. Hence, top management is likely to want to monitor and control both non-crisis period volatility and stress losses during crisis events.

Appendix

In this appendix we provide the analytical derivatives needed to compute the *expected tail loss* beta of asset i with respect to portfolio $p^\dagger$. As in the text of the paper:

- S is the number of states,
- π_s is the probability of state s occurring, and
- A is a return threshold that depends on the percentile choice, α.

In addition, define the normalized return threshold in state s, a_s, as:

- $a_s = \frac{A - \mu_{p^\dagger,s}}{\sigma_{p^\dagger,s}}$

Consider portfolio $p^\dagger$ that has a portfolio weight of w_i in the i^{th} asset and $(1 - w_i)$ in portfolio p. The Expected Tail Loss, ETL, may be expressed as:

$$ETL = \sum_s^S \pi_s \left[\mu_{p^\dagger,s} F\left(a_s\right) - \sigma_{p^\dagger,s} f\left(a_s\right)\right] / \alpha \tag{A.1}$$

where $F(\cdot)$ denotes the standard *state-dependent* normal cumulative distribution function of the normalized portfolio returns, $f(\cdot)$ denotes the corresponding standard *state-dependent* normal probability density function and equation (A.1) is subject to the constraint:

$$\sum_s^S \pi_s F(a_s) = \alpha \tag{A.2}$$

The ETL in each state depends on the state-dependent mean, $\mu_{p^\dagger,s}$, and standard deviation, $\sigma_{p^\dagger,s}$ of the portfolio which are given by:

$$\mu_{p^\dagger,s} = w_i \mu_{i,s} + (1 - w_i)\mu_{p,s} \tag{A.3}$$

and

$$\sigma_{p^\dagger,s} = \left[w_i^2 \sigma_{i,s}^2 + (1 - w_i)^2 \sigma_{p,s}^2 + 2w_i(1 - w_i)\sigma_{ip,s}\right]^{\frac{1}{2}} \tag{A.4}$$

where $\sigma_{ip,s}$ denotes the state-dependent covariance between the returns on asset i and the returns on portfolio p.

The derivative of the ETL with respect to w_i is:

$$\begin{aligned} \left.\frac{dETL}{dw_i}\right|_{w_i=0} &= \frac{1}{\alpha} \sum_s^S \pi_s \left\{ \left(\frac{d\mu_{p^\dagger,s}}{dw_i}\right) F(a_s) + \mu_{p^\dagger,s} \left(\frac{dF(a_s)}{dw_i}\right) \right. \\ &- \left. \left(\frac{d\sigma_{p^\dagger,s}}{dw_i}\right) f(a_s) - \sigma_{p^\dagger,s} \left(\frac{df(a_s)}{dw_i}\right) \right\}\Bigg|_{w_i=0} \end{aligned} \tag{A.5}$$

where the derivatives with respect to the portfolio weights are evaluated at $w_i = 0$ and where equation (A.5) is subject to the constraint:

$$\sum_{s}^{S} \pi_s \frac{dF(a_s)}{dw_i}\bigg|_{w_i=0} = 0 \tag{A.6}$$

The analytical derivative given by equation (A.5)is the key component needed for the *expected tail loss* beta of asset i with respect to portfolio $p^\dagger$.

We now provide the analytic derivatives for the components of equation (A.5). The derivatives of the state-dependent mean and standard deviation with respect to w_i are:

$$\frac{d\mu_{p^\dagger,s}}{dw_i}\bigg|_{w_i=0} = \mu_{i,s} - \mu_{p,s} \tag{A.7}$$

$$\frac{d\sigma_{p^\dagger,s}}{dw_i}\bigg|_{w_i=0} = \rho_{i,p}\sigma_{i,s} - \sigma_{p,s} \tag{A.8}$$

where $\rho_{i,p}$ is the correlation between the returns of the i^{th} asset and portfolio p.

Using the chain rule, we can write the derivatives of the cumulative distribution and density functions of the normalized portfolio returns with respect to the portfolio weights as:

$$\frac{dF(a_s)}{dw_i}\bigg|_{w_i=0} = \left(\frac{da_s}{dw_i}\right)\left(\frac{dF(a_s)}{da_s}\right)\bigg|_{w_i=0} \tag{A.9}$$

$$\frac{df(a_s)}{dw_i}\bigg|_{w_i=0} = \left(\frac{da_s}{dw_i}\right)\left(\frac{df(a_s)}{da_s}\right)\bigg|_{w_i=0} \tag{A.10}$$

The component derivatives are:

$$\frac{da_s}{dw_i}\bigg|_{w_i=0} = \frac{1}{\sigma_{p^\dagger,s}}\left\{\frac{dA}{dw_i}\bigg|_{w_i=0} - (\mu_{i,s} - \mu_{p,s}) - a_s(\rho_{i,p}\sigma_{i,s} - \sigma_{p,s})\right\} \tag{A.11}$$

$$\frac{dF(a_s)}{da_s} = f(a_s) \tag{A.12}$$

$$\frac{df(a_s)}{da_s} = -a_s f(a_s) \tag{A.13}$$

Substituting the right-hand-side of equation (A.11) for $\frac{da_s}{dw_i}\Big|_{w_i=0}$ in relations (A.9) and (A.10) and using (A.12) and (A.13) gives:

$$\frac{dF(a_s)}{dw_i}\bigg|_{w_i=0} = \frac{f(a_s)}{\sigma_{p^\dagger,s}}\left\{\frac{dA}{dw_i}\bigg|_{w_i=0} - (\mu_{i,s} - \mu_{p,s}) - a_s(\rho_{i,p}\sigma_{i,s} - \sigma_{p,s})\right\} \tag{A.14}$$

$$\frac{df(a_s)}{dw_i}\bigg|_{w_i=0} = -a_s\frac{dF(a_s)}{dw_i}\bigg|_{w_i=0} \tag{A.15}$$

Substituting the right-hand-side of equation (A.14) into (A.6) and solving for the derivative of the $\alpha\%$ tail upper bound with respect to the portfolio weights (i.e. $\left.\frac{dA}{dw_i}\right|_{w_{i=0}}$) gives:

$$\left.\frac{dA}{dw_i}\right|_{w_{i=0}} = \sum_{s}^{S} \pi_s \frac{f(a_s)}{\sigma_{p,s}} \left[(\mu_{i,s} - \mu_{p,s}) + a_s(\rho_{i,p}\sigma_{i,s} - \sigma_{p,s})\right] \div \left(\sum_{s}^{S} \pi_s \frac{f(a_s)}{\sigma_{p,s}}\right) \tag{A.16}$$

Note that A is the $(1-\alpha)$ percentile VaR for the state dependent normal distribution. Hence, this last derivative may be interpreted as the sensitivity of the VaR to changes in the weight of asset i in the portfolio.

This completes the analytical expression for $\left.\frac{dETL}{dw_i}\right|_{w_{i=0}}$.

References

[1] Phillippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath, *Coherent Measures of Risk*, *Mathematical Finance*, **9**, 1999, pp. 203-228.

[2] Michael J. Best and Robert R. Grauer, *On the Sensitivity of Mean-Variance-Efficient Porfolios to Changes in Asset Means: Some Analytical and Computational Results*, *Review of Financial Studies*, **4**, January 1991, pp. 315-342.

[3] Fischer Black and Robert Litterman, *Asset Allocation: Combining Investor Views with Market Equilibrium*, Goldman Sachs & Company Fixed Income Research, September 1990.

[4] Michael J. Brennan, *The pricing of contingent claims in discrete time models*, *Journal of Finance*, **34**, 1979, pp. 53-68.

[5] Markus K. Brunnermeier and Lasse Heje Pedersen, *Market Liquidity and Funding Liquidity*, working paper, February 2007.

[6] Nicholas Chan, Mila Getmansky, Shane M. Haas, and Andrew W. Lo, *Systemic Risk and Hedge Funds* in *Risks of Financial Institutions*, Mark Carey and Rene Stulz, eds., University of Chicago, Chicago, 2007.

[7] David M. Modest, *Long-Term Capital Management: An Internal Perspectivie*, presentation at Yale School of Management, April 10, 2001.

[8] Philippe Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk,* Third Edition, McGraw-Hill, New York, 2007.

[9] *J.P. Morgan Chase & Company Annual Report*, 2005.

[10] Charles P. Kindleberger and Robert Aliber, *Manias, Panics and Crashes,* Fifth Edition, John Wiley & Sons, Hoboken, New Jersey, 2005.

[11] Joshua V. Rosenberg and Til Schuermann, *A General Approach to Integrated Risk Management with Skewed, Fat-tailed Risks*, forthcoming *Journal of Financial Economics.*

Figures and Tables for:
Crisis and Non-Crisis Risk in Financial Markets

Robert H. Litzenberger
and
David M. Modest

Figure 1:
Graphical Illustration of Simple Two-Regime Stress-Loss Framework

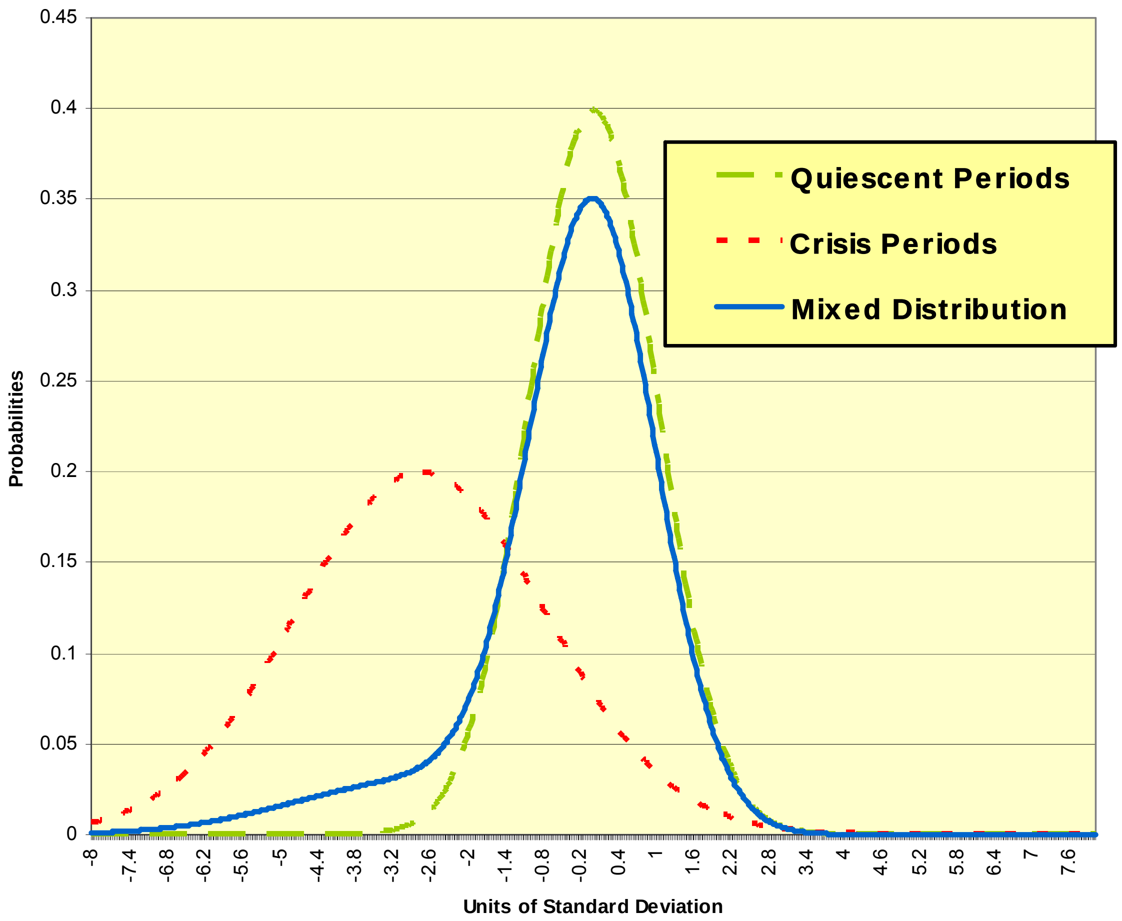

Figure 2:
Expected Tail Loss

(Expected losses are on the y-axis and units of standard deviation on the X-axis)

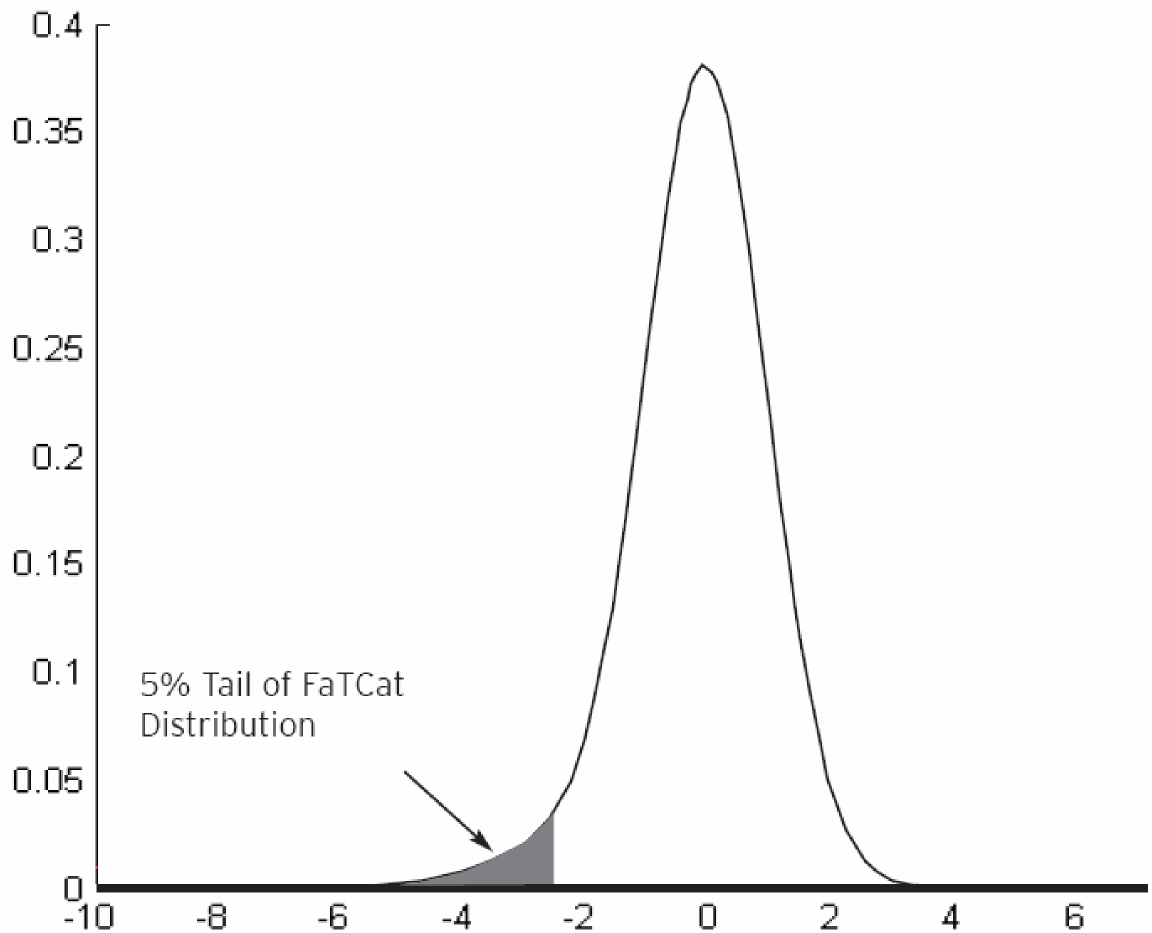

Figure 3:
Expected Return (y-axis) /
Expected Tail Loss Frontier (x-axis)

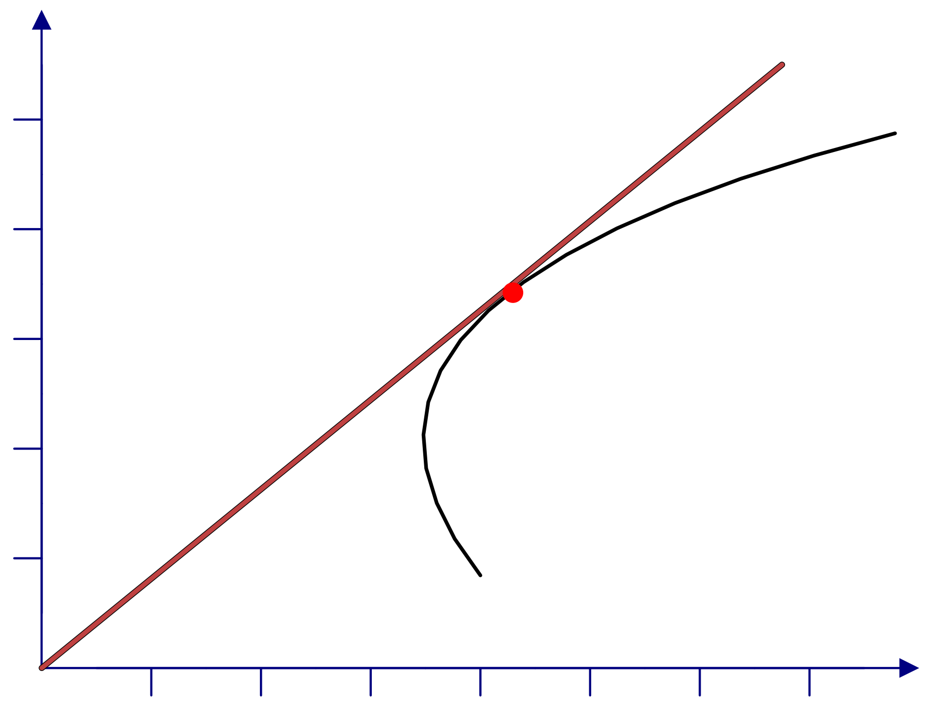

Table 1:
Hedge Fund Strategies Asset Allocation Example

Strategy	$ Allocation	Strategy Expected Volatility	Strategy ETL	Contribution to Portfolio ETL	β^{ETL}	Contribution to Portfolio Volatility	β^{σ}	Implied Expected Excess Return
Global Macro	21.00%	6.00%	-8.00%	16.12%	0.77	23.99%	1.14	4.90%
Long / Short Equity	41.02%	5.50%	-7.00%	39.80%	0.97	48.05%	1.17	6.20%
Credit and Distressed Securities	17.68%	4.50%	-9.00%	24.40%	1.38	15.96%	0.90	8.90%
Fixed Income Arbitrage	5.75%	2.50%	-8.00%	6.57%	1.14	1.35%	0.24	7.40%
Convertible and Volatility Arbitrage	5.27%	4.50%	-7.00%	4.83%	0.92	2.88%	0.55	5.90%
Statistical Arbitrage	3.59%	2.50%	-4.00%	1.83%	0.51	0.87%	0.24	3.30%
Event-Driven and Merger Arbitrage	5.70%	6.00%	-8.00%	6.45%	1.13	6.89%	1.21	7.30%

Table 2:
Portfolio Tilts and Required Excess Returns:
Equal Confidence in Forecast and Prior

Strategy	-50% Portfolio Weight	-25% Portfolio Weight	Base Implied Excess Return	+25% Portfolio Weight	+50% Portfolio Weight
Global Macro	0.8%	3.4%	4.9%	6.3%	7.5%
Long / Short Equity	3.3%	5.1%	6.2%	7.7%	8.6%
Credit and Distressed Securities	6.0%	7.6%	8.9%	9.9%	10.7%
Fixed Income Arbitrage	6.4%	6.9%	7.4%	7.8%	8.3%
Convertible and Volatility Arbitrage	5.5%	5.7%	5.9%	6.1%	6.2%
Statistical Arbitrage	3.2%	3.25%	3.3%	3.35%	3.4%
Event-Driven and Merger Arbitrage	6.9%	7.1%	7.3%	7.5%	7.6%

Table 3:
Portfolio Tilts and Required Excess Returns:
Confidence in Forecast is One-Half of Prior

Strategy	-50% Portfolio Weight	-25% Portfolio Weight	Base Implied Excess Return	+25% Portfolio Weight	+50% Portfolio Weight
Global Macro	-3.2%	1.0%	4.9%	8.5%	11.4%
Long / Short Equity	-4.7%	1.4%	6.2%	9.8%	12.2%
Credit and Distressed Securities	1.8%	5.7%	8.9%	11.4%	13.4%
Fixed Income Arbitrage	4.9%	6.2%	7.4%	8.5%	9.6%
Convertible and Volatility Arbitrage	5.0%	5.5%	5.9%	6.3%	6.7%
Statistical Arbitrage	3.1%	3.2%	3.3%	3.4%	3.5%
Event-Driven and Merger Arbitrage	6.3%	6.8%	7.3%	7.7%	8.2%