SOCIETY OF ACTUARIES

Article from:

Risk Management

August 2014 – Issue 30

Tails, You Lose: Making Sense of Tail-Hedging Indexes

By Edward K. Tom, Stanislas Bourgois and Grace J. Koo

THE THREE-YEAR PERIOD DEFINED BY THE START OF THE CREDIT CRISIS IN 2008, the intervening Flash Crash, and the subsidence of the Sovereign Debt Crisis in 2011 marked one of the most volatile regimes in market history. Of particular note were the successive waves of "tail events," market dislocations deemed a priori, to be statistically improbable. Although differing in both intensity and duration, these events, collectively known as "fat tail" or "black swans," precipitated abrupt and immense drawdowns as stock prices unraveled from company and macroeconomic fundamentals.

Edward K. Tom is a managing director and global head of equity and equity derivatives trading strategy at Credit Suisse in New York, N.Y. He can be reached at *ed.tom@credit-suisse.com.*

Stanislas Bourgois is a director and head of equity derivatives strategy EMEA at Credit Suisse in New York, N.Y. He can be reached at *stanislas.bourgois@credit-suisse.com*

WHY HEDGE TAILS?

As an example of the potential impact of tail events upon a market portfolio, consider the magnitude of the drawdowns experienced during the heart of the Credit Crisis in 2008. As seen in Figure 1, under the assumptions of normality embedded into modern portfolio theory, it is anticipated that over the course of a trading career, one would observe at most one one-day drawdown in excess of four standard deviations (i.e, 5+ percent). Yet as shown in Figure 1 and 2, during the four-month period from Aug 2008–Dec 2008, the market experienced ten such declines—negating in the span of four months, six years of equity growth. On the surface, therefore, the most obvious and oft-cited reason to hedge against tail events is to mitigate the severity of the market drawdown.

Figure 1: A Priori Probability of One-Day Market Declines

Sigma	Expected Frequency Outside Range	Frequency for Daily Trading Event
1	1 in 3	Twice as week
2	1 in 22	Once a month
3	1 in 370	Once very 1.5 years
4	1 in 15,787	Once very 63 years
5	1 in 1,744,278	Once very 7,000 years
6	1 in 506,797,346	Once very 2 million years
7	1 in 390,682,215,445	Once very 1.5 billion years

Source: Credit Suisse Derivatives Strategy

Figure 2: A Priori Probability of One-Day Market Declines

Date	SPX Decline	Pre-Crisis Expected % Occurrence
9/29/08	-8.79%	763,083,992 years
10/7/08	-5.74%	584 years
10/9/08	-7.62%	2,011,100 years
10/15/08	-9.03%	3,180,535,165 years
10/22/08	-6.10%	2,501 years
11/5/08	-5.27%	96 years
11/12/08	-5.19%	96 years
11/19/08	-6.12%	2,501 years
11/20/08	-6.71%	34,267 years
12/1/08	-8.93%	1,550,262,586 years

Source: Credit Suisse Derivatives Strategy

A more subtle and arguably more important benefit of a tail hedge however, is that it addresses the most disruptive feature of a tail shock—specifically, the impact associated market distortions that often accompany tail events. These following market distortion undermine 1) the underlying principles of financial valuation—causing a departure of asset prices from their "fair" values and 2) the stabilizing assumptions of portfolio construction including:

- breakdown in portfolio diversification (via correlation)
- negative feedback loops (via volatility clustering)
- beta instability (via cross-asset contagion)
- discontinuous trading

Volatility Buffer

Often during these events in which in-house volatility based risk limits are suddenly breached, portfolio managers ("PMs") and traders are forced to sell out of tactically unattractive but strategically desirable positions. Tail hedges can provide a volatility buffer to slow the exit of positions (or lessen the impact).

"The primary challenge during the current low volatility environment, however, is that the cost of static, "always on," tail insurance is often expensive to hold."

Credit Reserve

It is somewhat ironic that downside tail events also provide the best opportunity to outperform. In fact, a historical analysis of returns shows that there are almost as many upside tail shocks and of similar as there are downside tails. Take for example the Crash of 87 in which the market collapsed 23 percent over the course of one day but recoups the bulk of the losses over the course of the next two days. A good way to recover returns lost due to a tail shock is therefore to invest during times of market duress. However, in many cases, trader positions are often drastically pared down as the aforementioned risk limits are breached. An important function of tail hedges is therefore to provide a source of funding which accrues as the market is in decline and which can then be used to lever into a long position to allow the portfolio to more quickly recover.

Algorithmic (Signals-Based) Tail Hedging

The primary challenge during the current low volatility environment, however, is that the cost of static, "always on," tail insurance is often expensive to hold. Accordingly, if a tail event fails to materialize, the buyer of a systematic tail strategy risks significantly underperforming his unhedged peers. To moderate the cost of carry, hedgers often shift towards dynamic tail risk strategies during times of market stability.

Over the last few years, a vast number of dynamic strategies in the form of algorithmic indices[1] have been designed to profit from the realization of tail events and offered as a hedging product to end investors. Algorithmic indices (algos) are liquid, transparent and easily investable through delta-one wrappers such as swaps, notes or more advanced products involving the use of derivatives and/or leverage in order to produce a highly asymmetrical payoff.

Algorithmic Tail Risk Construction

As of the time of writing, the marketplace currently has over 200 active tail risk algorithm (algo) products spanning five asset classes. However, due to the leverage to downside shocks and the greater liquidity offered by equity volatility products in times of market distress the majority of algo products invest in equity volatility. Figure 3 provides a cross-section of Credit Suisse's more popular tail hedging algos (by notional invested), its asset class exposure, and a short description of the trading rules.

Algorithmic Tail Risk Construction

Algorithmic tail risk construction generally follows a five-step process.

1. Tail Definition
2. Benchmark Selection
3. Trigger Design
4. Simulation
5. Test of Efficacy

In the following pages, we will use the development of our Equity Dynamic Tail Hedge Index (Ticker: DYTL) as a case study to illustrate the process of constructing a tail risk algo.

Step 1: Tail Definition

The obvious first step to developing a tail risk algo is to first define what is meant by "tail." Given the breadth of investment styles, the definition of the term "tail-risk" itself (and therefore the solution) may vary greatly among investment professionals. Take for example, the "Flash Crash" in which the market plummeted 10 percent over the course of one hour and then recovered 8 percent over the next hour to finish down 2 percent for the day. For an investor such as a high frequency trader or an active delta-hedger who was actively trading during that period and so realized profit and loss ("P&L") during those volatile two hours of the day, such an event may in fact qualify as a tail event. However, if one were a "low-frequency," long-term investor such as a pension fund that did not trade during that day, then a tail event may refer to a protracted deterioration in one's portfolio caused by a breakdown of the core investment strategy. For the purposes of this case study, we will define a tail risk as **a sizeable abrupt market decline which triggers a persistent volatility regime shift from a low to high volatility environment**.

CONTINUED ON PAGE 16

Grace Koo is a managing director of the investment banking division at Credit Suisse in New York N.Y. She can be reached at *grace.koo@credit-suisse.com.*

Figure 3: Credit Suisse Tail Hedging Algos

Index	Short Name	Bloomberg Code	Underlying	Dynamic/ Static	Source of Tail Exposure	Short Description
Credit Suisse Advance Defensive Volatility	ADVOL	CSEAADVL	Equity	Static	Long VIX Futures	The index offers efficient long volatility exposure by systematically going long short- or medium-term VIX futures based on current levels of VIX futures
Credit Suisse Dynamic Tail	DYTL	CSEADYTL	Equity	Dynamic	Long SX5E Volatility Skew	The index dynamically allocates to the Credit Suisse Equity Tail Hedge based on the level of the SX5E Skew or the iTRAXX credit index
Credit Suisse Equity Tail Hedge	TAIL	CSEATAIL	Equity	Static		The index offers efficient long SX5E skew exposure by going short delta-hedged put ratios
Credit Suisse Dynamic Tail S&P	DTSP	CSEADTSP	Equity	Dynamic	Long S&P Volatility Skew	The index dynamically allocates to the Credit Suisse Tail Hedge S&P based on the level of the S&P Skew or the CDX credit index
Credit Suisse Equity Tail Hedge S&P	TLSP	CSEATLSP	Equity	Static		The index offers efficient long S&P skew exposure by going short delta-hedged put ratios
Credit Suisse Cheapest Slide	CHPS	CSEACHPS	Equity	Static	Long SX5E Forward Variance Swaps	The index offers efficient long volatility exposure by systematically going long the cheapest-to-carry SX5E forward starting variance swap
Credit Suisse Advanced Volatility Index - Foreign Exchange Opportunistic Vol	AVI FX	CSVI	FX	Static	FX Volatility	The index opportunistically goes long/ short volatility across 12 major currency pairs based on a Jump model, with a systematic net long volatility bias
Credit Suisse Tail Risk Overlay Protection Strategy	TOPS	CSEATSERUS	Fixed Income	Static	Long Treasury, German Bonds; Long Euro rate futures	The index offers exposure to tail events by opportunistically going long CBOT Note Futures, Eurex German Bond Futures, or Euronext and CME Euro rates futures
Benchmark	Benchmark	NA	Equity	Static	Long SX5E and S&P Volatility	The index goes systematically long 2-month 90% put options on SX5E and S&P and carries them to maturity

Source: Credit Suisse Derivatives Strategy

Step 2: Benchmark Selection

The second step is to create a "naïve" or systematic hedging benchmark index (the Benchmark) using a plain-vanilla options strategy in order to gauge the relative performance of the tail hedging strategy. In our example, our Benchmark is designed as follows:

- Strategy: On every listed expiry, we purchase new S&P-500 90 percent-struck put options with a two-month maturity. At any time we would therefore have two options in the portfolio with maturities equal to front month and back month expiries. All options are either 1) let to run until they expire or 2) unwound in the event that the delta reaches 100.

"The objective in employing a trigger mechanism is to decrease the weighting (and therefore the cost) of the downside hedge in times of quiet markets and ratchet up exposure in anticipation of a tail event."

- The notional of the purchased options is equal to one-fourth of the mark-to-market value of the Benchmark on that same day in order to match exposures.

- Performance calculation: the Benchmark is calculated in USD. Payoffs or premiums are paid in and out of a synthetic USD cash account earning Fed Funds.

The simulated history of the benchmark is shown on Figure 4. We also show the cumulative P&L of the S&P-500 index, the cumulative P&L of the S&P with a one-to-one overlay of the Benchmark as a hedge.

Figure 4 demonstrates the conundrum faced by many systematic plain-vanilla hedging strategies:

When a tail event does materialize, such a strategy can successfully cushion the initial blow of a tail event. In our example, for $100 invested in the portfolio in April 2008, the hedging strategy would have saved the investor up to $20 by November 2008. However, if a tail event does not materialize, it also shows how the long-term running cost (the Carry) of the strategy may gradually eat-up the accrued hedging benefits. During our five-year backtesting period, implementing the Benchmark hedge would have left the investor worse-off by $12 per $100 investment at the end of April 2013.

This then illustrates the disadvantage of a static tail hedge strategy: by systematically investing in the same notional, it tends to be under-invested in the period leading up to the shock, causing the investor to be under-hedged, and it tends to be over-invested immediately after the tail event when the price of options is high and the risks have dissipated, resulting in higher performance drag.

Step 3: Trigger Mechanism
To enhance the performance of the basic "benchmark" tail hedge, we thus introduce the use of a timing indicator or trigger mechanism. The objective in employing a trigger mechanism is to decrease the weighting (and therefore the cost) of the downside hedge in times of

Figure 4: Benchmark Tail Strategy vs. S&P and SX5E ('08-'13)

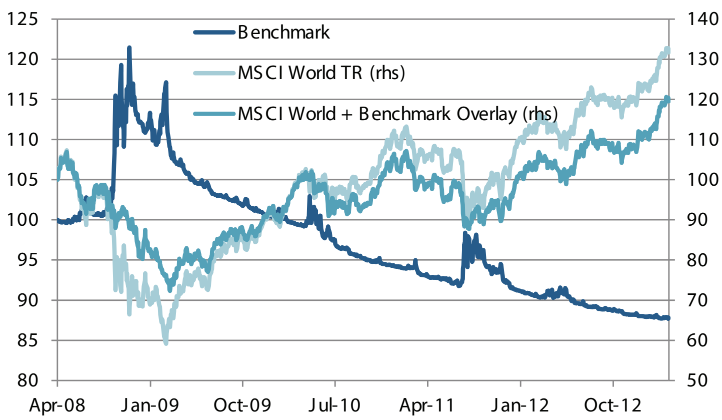

Source: Credit Suisse Equity Derivatives Strategy

Figure 5: S&P 3M normalized skew vs. S&P index level

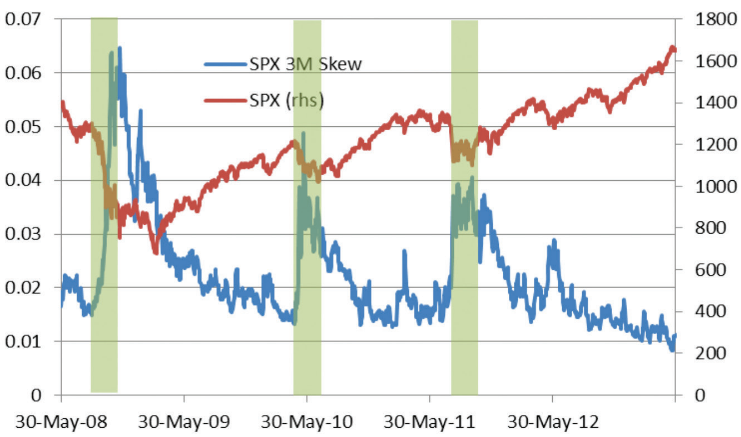

Source: Credit Suisse Equity Derivatives Strategy

CONTINUED ON PAGE 18

quiet markets and ratchet up exposure in anticipation of a tail event. In our example, we discuss the use of two triggers taken from two asset classes: 1) equity volatility skew and 2) CDS spreads from the fixed income markets in the construction of the Credit Suisse Equity Dynamic Tail Hedge index.

Signal 1 – Skew
Implied equity market skew is defined as the difference between implied volatility for lower strike options (typically put options purchased for protection) and implied volatility for higher strike options (typically call options purchased for leveraged upside exposure). Historically, during severe market downturns implied equity market skew has increased significantly. (Figure 5) This may be explained by an increase in demand for downside protection, pushing up implied volatility levels for lower strike levels.

The indicator analyzes the historical distribution of the three-month 80-100 skew on the underlying equity index, over the last three months. If the skew level is above 1.5 standard deviations from the mean, the signal for a distressed market is activated. This indicator has been historically reactive to market events signaling the beginning of a tail episode.

Signal 2 – CDS Spreads
The indicator is linked to the five-year CDS spread of companies for the relevant underlying equity market. If the CDS index is above 125 percent of its three-months moving average, the signal for a distressed market is activated (Figure 5, green shaded bars). If the CDS index is below 100 percent of its three-months moving average, the signal for a distressed market is deactivated. Otherwise, the signal remains unchanged. The indicator captures medium-term risk and is reactive to changes in the macro environment.

Methodology
The underlying fundamental strategy can be broken down into five steps:

1. The algorithm completes a monthly sale of vanilla ratio-put-spreads on the underlying equity index consisting of:
    - short a number of three-month 95 percent puts and
    - long a number of three-month 80 percent puts
2. The quantity of puts is chosen such that each leg generates one volatility point (i.e., 1 percent vega exposure) per one point decline in the underlying index. The position thus naturally adapts to the prevailing level of equity volatility. Specifically, during times of low volatility when options' vega is low, the quantity of options needed to generate one volatility point increases, resulting in higher exposure to a tail event before it has happened. Likewise, when a tail event has realized and equity volatility and options' vega are high, the quantity of options needed to generate one volatility point is lower, and the strategy naturally deleverages itself at each reset. The ratio of 95 percent puts to 80 perfect puts has a historical average of 1-by-3.15.
3. The position is delta-hedged. (Once the directional component of the position is removed via delta-hedging, what remains is pure exposure to volatility.)
4. The puts are unwound a day before expiration to avoid expiration day effects and rolled on a monthly basis.
5. Any cash balance accrues at relevant rate.

To drive the allocation between cash and Index the two signals are run daily:

- If one of the signals is switched ON, 50 percent of the exposure is allocated to the CS Equity Tail Hedge SPX Index.
- If both signals are ON, 100 percent is allocated to the hedge index. If neither of the signals is ON, 100 percent is invested in cash (US Federal Funds Rate or EONIA).

Historically, at least one of the signals has been ON for 31 percent of the time period. Typically, a distressed macro environment would first activate the CDS signal, indicating that the likelihood of a tail event has

"One must therefore evaluate the benefit of reducing carry costs in times of stable markets versus the risk of potentially missing the event because the signals have been "switched off"."

increased. The skew signal would activate when the market crisis takes momentum and equity skew breaks out of range.

Step 4: Simulation

To determine the effectiveness of the tail hedging algo, one would typically divide the dependent data set (in this case daily S&P and Eurostoxx returns) into an in-sample data set which is used to construct the algo and an out-of-sample data set which is used to test the stability and effectiveness of the algo going forward. In our case, given the limited scope of CDS spreads data sets, we created a proxy data set for CDS spreads extending back to 1996 using the Merton model. The algo was then constructed using the in-sample data set from 1996 to 2006 and run out-of-sample using the actual CDS spreads from 2006 to 2013 as shown in Figure 6.

In general, our out-of-sample simulation for the CS Equity Dynamic Tail Hedge Index embodied the two traits we felt was desirable in a tail-hedging algo, delivering outsized returns during periods of market crisis, and efficiently reducing the effect of negative carry over stable market periods via the dynamic signals.

An important consideration is that tail risk strategies which incorporate some element of market timing regardless of whether it is actively determined by a PM or signal based face the very real risk that a hedge may not be in place when it is needed. One must therefore evaluate the benefit of reducing carry costs in times of stable markets versus the risk of potentially missing the event because the signals have been "switched off." The final step to the process of algo construction is therefore to conduct an additional test of efficacy above and beyond the basic simulation in order to determine 1) whether the inclusion of the proposed signals provide adequate cost reduction to compensate for the risk of the hedge being "deactivated" during the days leading up to a tail event and 2) how the chosen algo stacks up against alternative tail risk algos.

Step 5: Additional Tests of Efficacy

The primary criteria we use to evaluate the efficacy of tail risk algos is to compare the tail-to-carry ratio of each strategy with one another. The tail-to-carry ratio is computed by dividing the average performance during tail events by the negative annualized carry. The metric essentially conveys how many years of negative carry can be paid for by one single tail event. The higher the ratio, the more efficient the hedge.

Figure 6: CS Dynamic Tail (DYTL) and CS Dynamic Tail S&P (DTSP): Simulated Performance 2008 - 2013

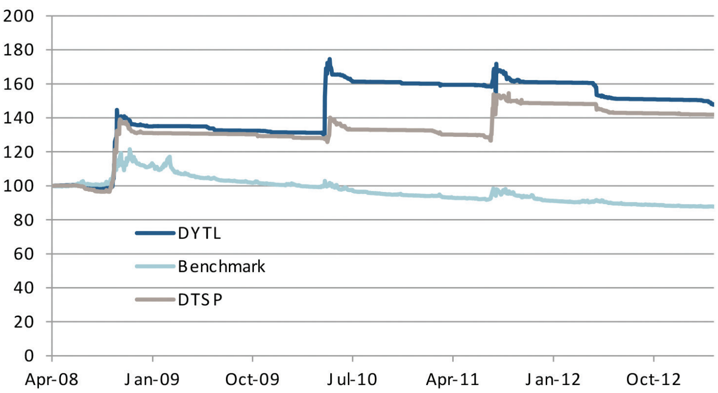

Efficacy of Signal Overlay

In our first example, we test the efficacy of our signal overlay, by comparing our signal based Dynamic Tail strategy index (DTSP), to its unconstrained parent strategy, the Tail Hedge S&P index (TLSP), which is 100 percent invested at all times. DTSP is invested 100 percent in cash when no risk indicator is on, 50 percent in cash, 50 percent in TLSP when only one risk indicator is activated, and 100 percent in TLSP when both risk indicators are on.

CONTINUED ON PAGE 20

Figure 7: DTSP vs. TLSP performance during Quiet markets (Carry) and Tail events (Tail)

From	To	Tail/Carry	DTSP	TLSP
22-Apr-08	1-Sep-08	Carry	-9.3%	-28.7%
1-Sep-08	1-Dec-08	Tail	36.4%	39.6%
1-Dec-08	1-Apr-10	Carry	-2.1%	-25.2%
1-Apr-10	31-May-10	Tail	7.9%	5.5%
31-May-10	1-Jul-11	Carry	-5.5%	-19.2%
1-Jul-11	30-Sep-11	Tail	15.5%	33.4%
30-Sep-11	3-Mar-14	Carry	-4.1%	-17.0%
		Average Tail	-5.2%	-22.5%
		Average Carry	19.9%	26.2%
		Tail To Carry Ratio	3.8	1.2

Source: Credit Suisse Equity Derivatives Strategy

Figure 8: Signal Activation during 2011 US Debt Downgrade

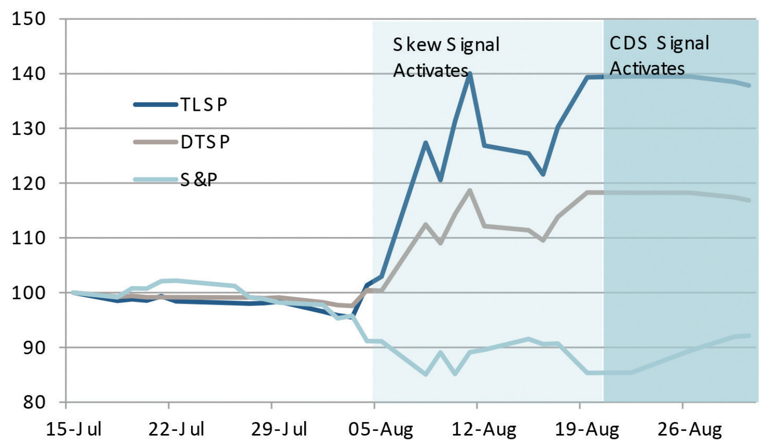

Source: Credit Suisse Equity Derivatives Strategy

Figure 7 compares the performance of DTSP vs TLSP from 2008 to 2011. At first glance, one might conclude the unconstrained "always-on" strategy is superior given that DTSP provided comparable returns to TLSP during the Lehman Collapse and the emergence of Greek Sovereign crisis in 2008 and 2010 but as shown in Figure 8, because the CDS signal activated late into the Tail strategy in summer 2011, DTSP underperformed. Note, however, that during periods of market stability, DTSP reduced the cost of carry, on average by a factor of five, producing a higher and therefore efficient tail to carry ratio.

FUNDING A TAIL RISK STRATEGY

The bane of every tail risk hedger is managing the punitive decay profile of their held options. Or put another way, tail risk strategies can be very costly over time, especially if there is no payout. Therefore, investors continue to explore ways to help defray the cost of maintaining a tail risk hedge—one of which is to overlay a mean-reversion strategy to the portfolio.

For example, in market regimes where we observe side-ways, or "saw-toothed" trading patterns, an investor could benefit from mean-reversion harvesting, and can convert a negative decay profile into one of positive accrual.

The CS Fixed Mean Reversion on S&P 500 is an algorithmic strategy that monitors the most recent five-day performance of the S&P 500, and if that performance is negative, a long position is established. Otherwise, it takes a short position. Moreover, the size of position is subject to caps and floors to avoid over-leverage.

"Dynamic tail-hedge strategies in the form of algorithmic indexes can provide a liquid, transparent, and easily investable solution to mitigate the impact of a "fat tail" or black swan market event."

In Figure 9, we present a 100 percent overlay of the CS Fixed Mean Reversion on S&P 500 on an underlying position of the CS Equity Dynamic Tail Hedge of the S&P 500.

We note that combining the two strategies augments the performance of the tail hedge significantly. Furthermore, we also observe that the correlation between them has historically been low, except during the second half of 2011 when the Eurozone crisis intensified. But, even in that instance, as the dynamic tail hedge performance spikes, the mean reversion strategy corrects, after declining sharply for a short period of time (see Figure 10).

CONCLUDING REMARKS

An ironic aspect of tail events is that it is not the expected or foreseeable events (the known unknowns) that causes the greatest market upheavals but rather the events from left field (the unknown unknowns). More often than not, true tail events often 1) have little or no historical precedent and 2) are difficult to anticipate a priori. Backtesting, by contrast, is by definition a backward-looking process that is optimized "to fight the last war." As a result, hedging strategies that are designed for a specific event or asset class that have been responsible for tails in the past may be optically attractive from a backtesting perspective but may not necessarily outperform if a future tail event is greatly dissimilar to prior shocks.

Nonetheless, dynamic tail-hedge strategies in the form of algorithmic indexes can provide a liquid, transparent and easily investable solution to mitigate the impact of a "fat tail" or black swan market event. In conclusion, the volatility buffer provided by a tail hedge not only serves to reduce the downswing in overall portfolio performance, but also could allow a credit reserve to put money to work after a market shock. A systematic tail hedge that also avoids a heavy cost-of-carry can keep PMs off the sidelines during the very time they should be the most active in navigating periods of market duress.

 The original article was published in May/June 2014 *Journal of Indexes*. The article was revised and additional strategy was added.

ENDNOTES

[1] Algorithmic indexes are rules-based, systematic investment strategies that are created to be transparent, liquid and investable. These indexes can, in turn, be packaged into structured notes, OTC swaps and options, and even funds. Algorithmic indexes differ from "trading algorithms" which typically focus on the execution of stocks and baskets of stocks.

Figure 9: 100% CS Fixed Mean Reversion Overlay (CSEAFMRS) and CS Dynamic Tail S&P (DTSP): Historical Performance December 2009 – March 2014

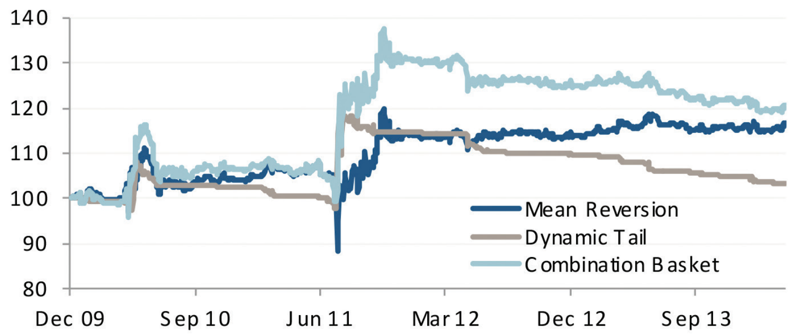

Figure 10: 12 month Rolling Correlation between CS Fixed Mean Reversion Overlay (CSEAFMRS) and CS Dynamic Tail S&P (DTSP): December 2009 – March 2014