FIRST QUADRANT

July 2014

How "Tail Risk" Changes Over the Market Cycle

FQ Perspective

ED PETERS
Partner

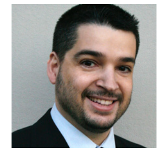

BRUNO MIRANDA PhD
Associate Director

Investors have been more conscious of "tail risk" since the financial crisis of 2008. The term refers to the frequency distribution of stock market returns having fatter "tails" than the normal distribution (the traditional "bell-shaped curve") which has been used for years to model market returns and evaluate risk. The fatter tails mean that the probability of a significant positive or negative event is larger than the normal distribution would lead us to believe. Since many risk management and asset allocation tools use the normal distribution as a starting assumption, the risk of a large or catastrophic event is often understated. The on-line financial dictionary Investopedia defines "tail risk" as

> "A FORM OF PORTFOLIO RISK THAT ARISES WHEN THE POSSIBILITY THAT AN INVESTMENT WILL MOVE MORE THAN THREE STANDARD DEVIATIONS FROM THE MEAN IS GREATER THAN WHAT IS SHOWN BY A NORMAL DISTRIBUTION."

Investors have also become more accepting of the existence of "volatility regimes." That is, the market can be defined by periods of persistently higher or lower than average volatility. Volatility is usually defined as annualized standard deviation. But when investors discuss a "three-standard-deviation event" they are likely speaking of the average standard deviation over time rather than recent standard deviation. Investors typically use 15% annualized standard deviation for the stock market as a rule of thumb since that fits the long-term standard deviation of many large cap indices such as the S&P 500 or the MSCI equity indices. So a "three-standard-deviation event" in monthly terms would be a decline of approximately -13%.

We can postulate that whether annualized volatility of the above indices is running at a 10% or 25% rate in the recent past, investors will still be referring to a -13% or greater monthly decline when they are describing "tail risk." For many, a two-standard-deviation decline of -8.66% would also be considered a "fat-tail event."

But is anchoring onto what would statistically be referred to as the "population standard deviation" (which assumes tail risk is constant) reasonable if we have significant volatility regimes? Previous work (Peters (2009)) has found that the market can be defined by periods of persistently higher or lower than median volatility with market characteristics quite different in the two regimes. These periods are partially explained by looking at the VIX, an index of the implied volatility of stock index options. Since the VIX is the implied volatility of options that falls out of the Black-Scholes option pricing formula (Black and Sholes (1973)), many confuse this number with a forecast of realized volatility. Unfortunately the VIX is a poor predictor of actual volatility. Instead, it is a better indicator of a level of market, or even macroeconomic, uncertainty. That is, it is a signal that the markets are entering a period of exceptional high or low uncertainty with changing market characteristics. What is not widely appreciated is that the VIX can also be a signal of increasing tail risk as well as a sign that the market is prone to "negative skew," meaning that there will be more large down-market than large up-market moves.

In this paper, we will document that for many sectors of the stock and bond markets, tail risk and negative skew are conditional upon volatility regime. Tail risk is not constant. We will first see this through a graphical representation of the distribution of returns. Then we will look at this phenomenon through new statistics we call conditional kurtosis and conditional skewness. That is, given we are in a high or low uncertainty environment, what is the probability of a "fat-tail event"? We will find that the risk of a large move and a negative skew in market returns is much higher when the VIX is above its median than below. So traditional statistics both understate and overstate tail risk. The actual methodology for calculating conditional kurtosis and skewness is simple and intuitive though it has not appeared in the literature as far as we know. It is also different than the conditional statistics discussed in the literature about GARCH models.

This research has significant implications for asset and security allocation and is also one more indicator of weakness in standard capital market theory such as the Sharpe-Litner-Mossin Capital Asset Pricing Model (CAPM) which remains the primary model for long-term strategic asset allocation (Sharpe (1964), Litner (1965), Mossin (1966)).

The Distribution of Market Returns

Quantitative finance has long assumed that market returns are normally distributed, or follow the familiar bell-shaped curve. This was necessary in order for many mathematical techniques to have validity (Mandelbrot (1964) and Peters (1994) among others). However, it has been well known since the 1950s that market returns are not normally distributed and have a high peak at the mean and fatter tails than the normal distribution (Osborne (1959)).

Exhibit 01 illustrates the distribution of monthly market returns (the MSCI World Index in local currency in excess of 3-month LIBOR, "MSCI") for 24 years ending 2013. The normal distribution is also shown using the mean and standard deviation of the MSCI.

We can see that over this period the MSCI does indeed have a slightly fatter negative tail than the normal distribution and a higher peak at the mean though its positive tail is slightly smaller than the normal distribution. The chance of a drop of -13% or higher is 0.70% rather than the 0.03% predicted by the normal distribution. While still not large, it is significantly higher than that predicted by the normal distribution though

EXHIBIT 01: MSCI DISTRIBUTION OF MARKET RETURNS
(1990 – 2013)

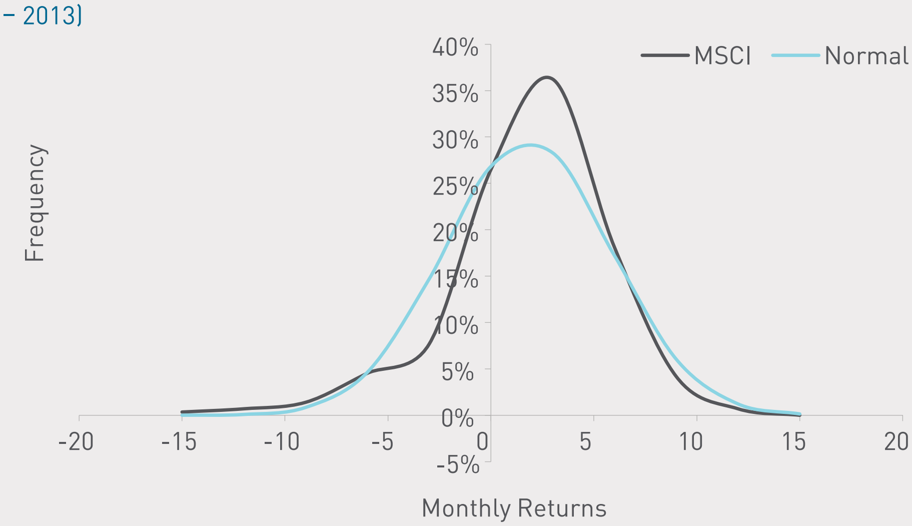

Sources: Global Financial Data, CBOE, FQ Analytics this slight increase may make many wonder what all the excitement is about.

The answer lies with market behavior during volatility regimes explored in Peters (2009) using volatility indices. As we will see below, the average of the tail risk in high and low volatility periods largely cancel each other out as it is higher than normal in high volatility regimes and the opposite in low volatility regimes. Tail risk changes over the market cycle in a fairly dramatic way.

Return Distributions with Volatility Regimes

High and low implied and realized volatility regimes are themselves the symptoms rather than the causes of the disease. Volatility represents the level of uncertainty about the markets and/or the economic backdrop of which the markets are a part. To update Peters (2009), we show in Exhibit 02 the history of a composite VIX over time. This composite VIX is a combination of a continuous maturity three-month S&P 500 VIX future, the three-month moving average of the EuroStoxx spot VIX (V2X), and the three-month moving average of the WTI oil VIX (OVX) weighted 60/30/10. The weightings were chosen to represent the cap weightings of the MSCI constituents with oil as a proxy for emerging markets. For the rest of this paper we will refer to this construct as the CVIX.

Looking at the graph we can see that there are long periods where the CVIX is above and below its median. These periods roughly correspond to periods of strong stock market returns and global economic expansions during lower-than-median CVIX, and periods of volatile market returns and global economic contractions during higher-than-median CVIX. Thus the CVIX can be postulated to be a sign of high and low macro and market uncertainty. Table 01 shows annualized MSCI returns as partitioned in Exhibit 01 for these periods.

TABLE 01: RETURN AND RISK UNDER VOLATILITY REGIMES MSCI CVIX REGIMES (1990 - 2013)

	High Vol	Low Vol	Full Period
Excess Return	-3.33	10.10	2.98
Standard Deviation	18.60	9.46	14.99
Serial Correlation	18.53	-1.12	15.91
Sample Kurtosis	0.18	-0.26	0.80
Sample Skewness	-0.62	-0.38	-0.87
Observations	147	141	288

Excess Return is over 30 day T- Bills as Cash
Sources: Datastream, Global Financial Data, CBOE, FQ Analytics

Periods of below-median CVIX (which will be referred to as "low uncertainty") correspond to higher excess returns and lower risk than periods of above-median CVIX (which will be referred to as "high uncertainty") where excess returns are negative and risk as described by annualized standard deviation is nearly twice as high. Serial correlation is also positive in high uncertainty but is largely absent in low uncertainty. Positive serial correlation shows that high uncertainty is not the time to "buy on the dips" expecting a reversion to the mean. Low uncertainty, on the other hand, is more likely to have reversals. Serial correlation is a sign of the fragility of markets in high uncertainty and their resilience in low uncertainty. So we can see that market behavior is quite different in the two regimes.

Kurtosis and skewness are statistics used to determine how far a distribution deviates from the normal distribution. Kurtosis is typically normalized so a reading of zero means the tails were the same size as the normal distribution. Positive kurtosis means that the tails are fatter than the normal distribution while negative kurtosis means the tails are thinner. Likewise, if skewness equals zero the distribution is symmetric. Negative skewness means that the tails on the left are larger than the tails on the right. The converse is true if skewness is a positive value. The normal distribution is symmetric so its skewness value is also zero.

In Table 1, the MSCI kurtosis is larger than the normal distribution in high uncertainty but in low uncertainty becomes smaller than the normal distribution. Skewness is a larger negative in

EXHIBIT 02: HISTORY OF THE COMPOSITE VIX
(1990 – 2013)

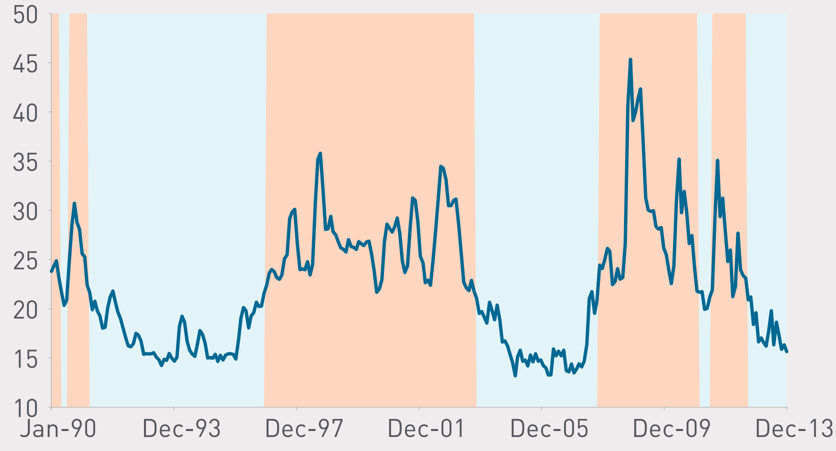

Sources: Global Financial Data, CBOE, FQ Analytics

high uncertainty than in low uncertainty but still negative for both periods. Based upon this statistical information we can see that periods of high uncertainty are quite different than periods of low uncertainty and neither is close to being "normal." Since the mean, variance, kurtosis and skewness are the first four moments of a distribution, we can see that they are quite different in the two sub-periods lending credence to the idea that they are indeed regimes.

Exhibit 03 shows the distribution of returns for the high and low uncertainty periods compared again with the normal distribution we used in Exhibit 01.

Graphically, the high uncertainty period is characterized by a much larger negative tail than the normal distribution but the positive tail is virtually the same. The low uncertainty environment, by contrast, has much smaller negative and positive tails than the normal distribution and a much higher peak at the mean. The regimes also show why the tail risk in Exhibit 1 looks slight. It is an average of the fat tails of the high uncertainty regime and the thin tails of the low uncertainty regime which almost cancel each other out.

In many ways, this confirms the information in Table 01. That is, periods of high uncertainty

EXHIBIT 03: MSCI DISTRIBUTION OF MARKET RETURNS BY REGIME
(1990 – 2013)

Sources: Global Financial Data, CBOE, FQ Analytics have less desirable characteristics than periods of low uncertainty. But we can return to our idea that most investors use the total period standard deviation to measure tail risk. The monthly standard deviation for the MSCI over the full 24-year period is 4.33%. According to Exhibit 03, the risk of a -13% decline or higher (an overall three-standard-deviation event) is 1.36% in high uncertainty but zero in low uncertainty vs. 0.03% for the normal distribution. For a "two-sigma event," or -8.66%, the risk in high uncertainty rises to 8.20% while it remains zero in low uncertainty vs. 4.50% in the normal distribution. Graphically it appears that all of the risk of a "two-standard-deviation drop" or greater occurs during periods where the CVIX is above its median. In addition, the probability of a two-standard-deviation advance or greater also occurs during high uncertainty. Low uncertainty is characterized by smaller returns with skewness closer to normal though still slightly negative. So based upon the graphical evidence, we can say that almost all of the probability of a drop of -8.66% or higher is in the high uncertainty period. We can also say, "Given that we're in a high uncertainty environment, the risk of a -8.66% decline or greater is 8.20% which is a significant risk." Yet the statistics in Table 1 do not show the significant difference between the skewness and kurtosis values in the two regimes as we see in Exhibit 03. In particular, the statistics in Table 1 have the unintuitive, if not incorrect, quantitative value showing kurtosis and skewness are better in the high uncertainty regime than in the full period. What are we missing?

Conditional Statistics

The problem appears to lie in the nature of statistical calculations. Given that we are taking two large subsamples from an already large dataset covering a long time period, we would expect that any two subsamples would have statistically insignificant differences with one another. We have seen that the mean and the standard deviation are significantly different from one another. And there lies the problem with the standard kurtosis and skewness calculation.

Kurtosis is the fourth moment of the distribution. It is equal to the average of the series deviations from the mean to the fourth power divided by the standard deviation to the fourth power. Skewness is the third moment and is calculated in a similar way to kurtosis. But skewness is equal to the average of the series deviations from the mean cubed, divided by the standard deviation cubed. In most standard statistical packages the mean and standard deviation of the subsample is used to calculate kurtosis and skewness assuming that the sample is large enough to give a good estimate of the population (or true) mean and standard deviation. The assumption is that the mean and standard deviation should be the same in both subsamples given that each is about half of a large sample covering 25 years. Yet we have seen that the mean and standard deviation of these two subsamples are quite different. So it would make sense that we should calculate the kurtosis and skewness of each subsample using the returns of each subsample but the mean and standard deviation of the entire time period encompassing both subsamples.

We call this conditional kurtosis and skewness. "Conditional" here is used in the Bayesian sense of the term. Bayes' theorem answered the question" What is the probability of event A given that event B has occurred?" Here we are asking, "What is the kurtosis (or tail risk) given that we are in a high (or low) uncertainty environment?" We can ask the same conditional question about skewness. By calculating the kurtosis and skewness of the subsamples using the mean and standard deviation of the whole series, we can answer that question. In addition, the conditional skewness and kurtosis become a form of attribution. The weighted average of the two sub-periods equals the whole time series for both skewness and kurtosis. In this way, we can attribute the number of observations that cause kurtosis or skewness to deviate from normal to a particular period. If the series were a random one characterized by the normal distribution, then the skewness and kurtosis should be statistically the same for both subsamples.

Table 02 shows the values of the conditional skewness and kurtosis calculations along with their statistical significance as measured by their standard error. A reading of +/- 1.96 is the critical number for significance at the 5% level. The conditional statistics are compared to their conventional sample equivalents.

TABLE 02: MSCI CONDITIONAL SKEWNESS AND KURTOSIS

MSCI CVIX REGIMES (1990 – 2013)

	Conditional	Sample
High Kurtosis	5.17	0.18
Significance	12.92	0.46
Low Kurtosis	-2.60	-0.26
Significance	-6.34	-0.87
Total Kurtosis	0.80	0.80
Significance	2.76	2.76
High Skewness	-1.73	-0.62
Significance	-8.61	-3.06
Low Skewness	0.06	-0.38
Significance	0.31	-1.87
Total Skewness	-0.87	-0.87
Significance	-6.00	-6.00

Sources: Global Financial Data, CBOE, FQ Analytics

We can see that the conditional statistics are significantly different in the two regimes from each other as well as their sample calculation equivalents.

In high uncertainty, kurtosis is a high positive number with a high level of significance showing that the tail is much fatter than the normal distribution. By contrast, the low uncertainty kurtosis is also significant but kurtosis is negative, meaning that tails are thinner than the normal distribution.

Likewise, skewness is a strong negative number and very significant in high uncertainty showing that in this regime there is a much higher likelihood of a large negative return rather than a positive one. In low uncertainty, skewness is close to zero and insignificant reflecting a more symmetric distribution.

Both statistics are much closer to describing the graphical evidence in Exhibit 02. You will also notice that the sample kurtosis and skewness seem to have little relationship to the whole time series values. Kurtosis over the entire period is 0.80 with a significance of 2.76. By contrast, the sample statistics show insignificant values which are smaller than the whole period. While the sample skewness values are more significant, their values are also smaller than the full time series. By contrast, the weighted average of the conditional skewness and kurtosis values are precisely the full period since each regime has about half the observations (the exact numbers are 53% high uncertainty and 47% low uncertainty).

Table 03 (next page) shows conditional statistics for a broader range of indices. We can see that all the indices show significant differences in the two regimes when we look at their conditional values.

Some, such as the FTSE NAREIT All REITs Index and the iBoxx USD Liquid HY Index show tail risk in high uncertainty that dwarfs that of equities and commodities. World government bonds, as represented by the Citi World Government Bond Index (WGBI), also show variations, but are much tamer. Interestingly, the skew values for government bonds are the reverse of the other assets showing that bonds do have a natural hedging ability in high uncertainty but lose that characteristic in low uncertainty. Emerging market equities and commodities have negative skewness in high uncertainty but significant positive skewness in low uncertainty.

Taken together, the graphical evidence and the conditional statistics show that tail risk for the so-called "risk-on" assets is largely concentrated in periods of high uncertainty. In periods of low uncertainty, not only is volatility lower, but the chance of a three-standard-deviation event as measured by the total period is very low. Empirically, it is actually zero. Does that mean the tail risk is zero in low uncertainty? Of course not. There is always a chance that an exogenous shock such as a terrorist attack or a large ecological disaster can strike. But based upon this evidence, the risk of a large tail event drops dramatically when the CVIX is in its lower state.

TABLE 03: INDEX RETURNS FOR CVIX REGIMES

CVIX REGIMES (1990 – 2013)

	60/40	R2000	MSCI EM	NAREIT	WGBI	High Yield	GSCI
High Kurtosis	5.26	4.32	4.95	15.97	1.65	16.49	3.62
Significance	13.15	10.79	12.36	39.90	4.13	41.20	9.04
Low Kurtosis	-2.26	-2.55	-1.02	-1.83	-0.79	-2.88	-0.43
Significance	-5.50	-6.23	-2.50	-4.47	-1.92	-7.03	-1.06
Total Kurtosis	0.92	0.56	1.16	4.17	0.26	4.03	0.94
Significance	3.18	1.94	4.03	14.44	0.91	13.97	3.24
High Skewness	-1.67	-0.90	-1.08	-1.33	0.29	-1.50	-0.75
Significance	-8.30	-4.49	-5.36	-6.58	1.44	-7.46	-3.74
Low Skewness	0.08	-0.01	0.40	-0.22	-0.43	-0.03	0.41
Significance	0.37	-0.06	1.94	-1.06	-2.07	-0.13	2.00
Total Skewness	-0.83	-0.48	-0.51	-0.79	-0.06	-0.79	-0.19
Significance	-5.73	-3.28	-3.52	-5.48	-0.40	-5.47	-1.31

Sources: Global Financial Data, CBOE, FQ Analytics

DEFINITIONS: 60/40 is a hypothetical capitalization weighted portfolio used for illustrative purposes and is comprised of 60% MSCI World Index (local currency), 40% Citi World Government Bond Index (local currency). R2000 is the Russell 2000 Index. MSCI EM is the MSCI Emerging Markets Index. NAREIT is the FTSE NAREIT All REITs Index. WGBI is the Citigroup World Government Bond Index. High Yield is the Markit iBoxx USD Liquid High Yield Index. GSCI is the S&P GSCI.

We can also see that using the sample method of skewness and kurtosis can significantly understate or overstate these risks for specific subsamples of the data. When an investor is preparing to hedge against these risks, it would be advantageous to have a clean measure for them. Conditional skewness and kurtosis supply those measures.

Implications

As we stated in the opening, investors typically use the average statistics over the long term to define what they mean by "tail risk." But we have seen that the first four moments of the distribution of market returns (mean, variance, kurtosis, and skewness) are very different in periods, or regimes, we labeled as high uncertainty and low uncertainty. So "tail risk" can change over the market cycle and "anchoring" onto the long-term standard deviation can overstate and understate actual risk. The regimes in this study were defined by partitioning a time series of MSCI excess returns into periods when the CVIX was above or below its long-term median. The CVIX is a practical and measurable indicator so these results are more than esoteric. What is more, the CVIX regimes appear to be an indicator of a broader state since Table 3 showed that many other market indices from different asset classes appear to be affected by the CVIX regimes as well. The true regime, which is likely unknowable, appears to be related to macro level uncertainties though that is not proven in this paper. These results, however, are conditioned upon knowing, with certainty, whether the market is in a high or low uncertainty environment. Will the definitions we used here for those two regimes hold in the future? That is, can we depend upon the median of the CVIX to be the same value going forward? The answer of course is, no. But it does appear to be more reasonable than assuming the averages hold over long periods. In fact, the two regimes are so different from one another that it is unlikely that we are ever at the "average" which assumes a level of statistical normality unlikely to exist.

While "tail risk" has always existed in the markets, it is mostly since the financial crisis of 2008 and the resulting bear market that investors have become accepting of its persistence. This paper shows that the probability of a large drop in the market varies over the market cycle. While we cannot be sure that the results of this study are definitive, they are nonetheless impressive. A sign that the markets are susceptible to the kind of shocks that can cause a two- or three-sigma or greater drop as measured by the overall standard deviation of returns can be considered conditional upon the current uncertainty regime. That is, the market can be considered to be "fragile" in high uncertainty where shocks can cause tail events while it is more "resilient" in the low uncertainty periods.

It should also be stated that these results apply to the broad markets which are driven by macro risks and not individual securities which are significantly affected by their own idiosyncratic risks.

The conditional statistics themselves have implications for asset allocation optimization. Those investors who are capable of adjusting their asset allocation through regimes have an opportunity to implement conditional strategic asset allocation if they choose. Making different distributional assumptions in different risk regimes would provide a valuable advantage to those who manage their asset allocation conditional upon risk regime. Not only are the volatility of the assets different, but tail risk and skewness vary significantly in the two regimes and over the market cycle and Peters (2009) showed correlations change as well. This alone should influence investment policy and asset allocation. At least it is evidence that investors should consider especially if they are concerned about the "quality of ride" of their results.

References

Black, F. and Scholes, M. (1973), "The Pricing of Options and Corporate Liabilities," Journal of Political Economy, May/June

Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," Journal of Econometrics 31

Litner, J. (1965), "The Valuation of Risk Assets and the Selection of Risk Investments in Stock Portfolios and Capital Budgets," Review of Economic Statistics 47

Mandelbrot, B. (1964), "The Variation of Certain Speculative Prices," in P. Cootner ed., The Random Character of Stock Prices, Cambridge: MIT Press

Markowitz, H. (1952), "Portfolio Selection" Journal of Finance 7

Merton, R.C. (1976), "Option Pricing When Underlying Stock Returns are Discontinuous", Journal of Financial Economics, 3

Mossin, J. (1966), "Equilibrium in a Capital Asset Market," Econometrica 34

Peters, E. (1994), Fractal Market Analysis, New York: Wiley, 1994

Peters, E. (2009) "Balancing Betas: Essential Risk Diversification," FQ Perspectives

FIRST QUADRANT, L.P. | 800 E. COLORADO BLVD. SUITE 900, PASADENA, CALIFORNIA 91101
MARKETING SERVICES INFO@FIRSTQUADRANT.COM | OFFICE 626.683.4223 | WEB FIRSTQUADRANT.COM