

Stable vs. Unstable Markets: A Tale of Two States

Executive Summary



ED PETERS

Partner, Investments

NIGHT TIME. TWO MEN ARE UNDER A LAMPPOST. ONE MAN IS SEARCHING THE GROUND ON HIS HANDS AND KNEES. THE OTHER STANDS OVER HIM.

“WHAT ARE YOU LOOKING FOR?” THE STANDING MAN ASKS.

“MY KEYS. I DROPPED THEM.”

“WHERE WERE YOU WHEN YOU DROPPED THEM?”

THE MAN ON THE GROUND POINTS OUT INTO THE DARKNESS.

“THEN WHY ARE YOU LOOKING HERE?” THE STANDING MAN ASKS.

“THE LIGHT’S BETTER HERE,” THE MAN ON THE GROUND REPLIES.

One of the bedrocks of modern capital market theory is that market risk and related statistics are stable over the long run. Nobel prizes have been won for this insight, and it is taught in the best business schools. Regulations have also been written based upon this assumption. Yet, experience does not support this idea. We know that often markets have periods of relative stability, but they can also be followed by years where it seems that all is chaos (and not in the physics sense of the word). In academia, there are theories that compete with the Capital Asset Pricing Model (CAPM), the main proponent of stable markets, but these competing theories are generally considered impractical since they don't lend themselves to easy solutions. The answers we often receive from the usable

models, however, go horribly wrong when markets go south. Could it be that we're using those models simply because “the light is better here?” But how can we use the models that are “impractical?”

In this paper, we show convincing evidence that there are actually two separate market states, each corresponding to these competing models. In essence, the CAPM and its critics are both right, but only part of the time.

The implications from this for asset allocation and plan management are profound. There will be periods where using standard techniques for asset allocation or investment management will work well. When the environment changes, however, those processes may no longer work with reliability. This is particularly true of diversification because assets that diversify one another in one state fail in the other state when they are truly needed. The rules change, and if investors hope to adapt successfully, they will need to know the new rules when the change occurs. This paper offers an explanation for that problem - two different market states causing assets to behave very differently. A subsequent paper will provide a roadmap to help asset owners anticipate and navigate deftly through the changes in market states.

How do we distinguish these two market states? In one state, markets are statistically well behaved. They can be modeled using standard statistical analysis. Volatility is stable and low. Correlations are stable. Tail events (3 standard deviation or larger returns in either direction) are rare. The periods correspond to low volatility

Past performance is no guarantee of future results. Potential for profit is accompanied by possibility of loss.



periods as defined by factors like the Implied Volatility Index (VIX) produced by the CBOE. We can call these “low uncertainty” states and they correspond broadly to equity bull markets and periods of economic expansion. “Buy and hold” works as does “buying on the dips.”

The second state is quite different. In “high uncertainty” states, markets are not statistically well behaved. Volatility and correlations change continuously. Four standard deviation events happen with regularity. For instance, according to a standard random walk, a four-standard deviation or higher event should happen once every 126 years (assuming a 260-day trading

year). Yet, in the last 25 years, the MSCI World Equity Index has experienced 38 of them, 22 on the downside, and all during periods of high uncertainty. These periods correspond to high volatility (above median) using the VIX and also correspond to large bear and bull markets, as well as periods of overall negative economic growth. In these environments, standard asset allocation techniques may not provide stable returns, and portfolio diversification and downside mitigation become increasingly important even as the ability to truly diversify are reduced.



Stable vs. Unstable Markets: A Tale of Two States

FQ Perspective

Uncertainty and risk are synonymous with investing. We know that positive returns in excess of cash yields require taking some risk of loss. Models such as the Sharpe-Lintner-Mossin Capital Asset Pricing Model (CAPM) postulate that there is a direct relationship between risk and return. The more risk the more potential return. Unfortunately, many studies have shown that this is not true, perhaps because there are many ways to define “risk.” Since Markowitz (1952), risk has been defined as the standard deviation of returns with the standard deviation calculated along the lines of the well-behaved normal distribution. While there have been variations on this basic theme, most go back to risk being “volatility” as defined by the standard deviation and diversification measured by the covariance of returns.

This paradigm assumes that risk is fairly stable over the long term. Markowitz (1952) assumes exactly stable statistics. So long-term asset allocation studies typically use 20- to 50-year measures of risk, expected return, and covariance even though we know that these numbers can deviate significantly from the averages for long periods of time. But conventional wisdom assumes that the average is indicative of the long term and investors can ride out the deviations which are considered rare short-term events.

But suppose that the market is rarely at the average but instead spends long periods of time in very different environments, and it is the average which is rare. The average then would be like the old joke “If my head is in the freezer and my feet in the oven, on average I feel fine.” If the market does indeed spend most of its time in the freezer or the oven, using the average would clearly over- and under-state risk, though on paper the market would look well behaved and “feel fine.”

This paper postulates such a two-state model. That is, that market returns for most risky assets are generally in one of two states, a high uncertainty state or a low uncertainty state. The behavior and risks of the markets, as described by the first four moments of a probability distribution (return, standard deviation, skew and kurtosis), as well as the correlations across assets are so different in the two states that using the average is a head-in-the-refrigerator-feet-in-the-oven problem. Complicating things even further, the transition from one state to the other is not a stable point though we can estimate the conditions for such a transition to occur. Contrary to popular belief, being at the average is the exception rather than the rule. We will also see that there is a difference between risk and uncertainty and there may not be a positive relationship between risk and return.

The results are similar to those found by Chow, et al (1999) but are more based upon a fundamental theory to explain the two states rather than empirical analysis alone. They are

Past performance is no guarantee of future results. Potential for profit is accompanied by possibility of loss.



also similar to the results reported by Ang and Timmermann (2011) who also found a two-state regime model with similar characteristics. Their model was strictly mathematical and volatility based, however. The analysis in this paper is, again, based upon fundamental theory and is much simpler.

These findings have significant implications for all facets of investing, but particularly for asset allocation. This model unites the CAPM approach of a stable variance assumed by Sharpe (1964), Lintner (1965) and Mossin (1966) with the Stable Paretian Hypothesis of Mandelbrot (1964) which postulated an unstable variance. Both are true at different times and have a significant effect on something as basic as the shape of the distribution of returns as well as the resulting statistics.

Instability of Market Returns

Investors continue to make decisions based upon long-term measures of risk despite the fact that we have known for decades that stock returns do not have a stable statistical structure. A simple test was pioneered by Fama (1964) before he formed the Efficient Market Hypothesis. His test showed that stock returns cannot follow a conventional random walk even if you measure risk for long periods. The test is called a “sequential standard deviation” or SSD. SSD calculates the standard deviation as you expand the window of observations. Once you have enough observations, then the SSD should converge to the population, or true standard deviation. Exhibit 01 shows such a calculation for 3,500 standard normally distributed numbers with a mean of zero and standard deviation of one.

We can see that once we have about 1,000 observations the numbers converge very closely to the pre-defined population standard deviation of 1.0. Fama originally did this calculation using daily Dow Jones Industrials returns. Peters (1994) did the same from 1888 to 1990 using five-day returns. The resulting sequential standard deviation of returns never converges to a stable value. Instead, the SSD

EXHIBIT 01: CONVERGENCE OF THE SEQUENTIAL STANDARD DEVIATION FOR NORMAL RANDOM NUMBERS



Source: FQ Analytics

suddenly jumps when a fat-tail event occurs. Exhibit 02 shows the SSD for S&P 500 returns from 1960-2014 using 54 years of daily data or 13,505 observations.

We can see that not only is there no convergence, but stable periods are interrupted by jumps in volatility. The Crash of 1987 is clearly visible in the center of the chart, for instance. Each time the

EXHIBIT 02: SEQUENTIAL STANDARD DEVIATION - S&P 500 DAILY OBSERVATIONS (JANUARY 1961 - AUGUST 2014)



Sources: Data Stream, Global Financial Data

graph tries to converge (shown by the flat or declining portion of the graph), there is inevitably a fat-tail event which pushes the graph away from convergence. So even after 54 years there is no convergence to a population standard deviation. Similar results led Fama (1964) and Peters (1991, 1994) to agree with Mandelbrot (1962) that the unconditional standard deviation of the market is undefined. That is, the second moment of the probability density function (which is the standard deviation) does not lend itself to a closed form solution unlike the normal distribution.

But perhaps things have changed with the advent of modern electronic trading and arbitrage? Exhibit 03 shows the post crash period starting in 1988.

Unfortunately, things have not changed. While this sample has 6,469 observations, that is still more than enough for convergence; but the high volatility tech bubble period starting with the emerging markets crisis of 1997 is still clearly visible. While there was a stable period starting in 2003, this ended with the “Quant Meltdown” of 2007 which lead into the Credit Crisis of 2008. Not until 2012 does stability return.

In addition, the basic distribution of returns remains non-normal with a high peak at the mean and fat tails. Exhibit 04 shows the distribution of daily

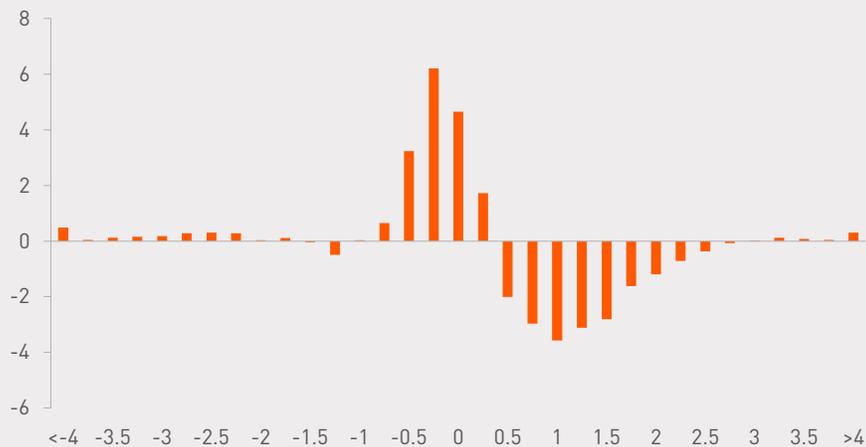
EXHIBIT 03: SEQUENTIAL STANDARD DEVIATION - S&P 500 DAILY OBSERVATIONS (JANUARY 1988 - AUGUST 2014)



returns from January 1988 – August 2014 in z-scores less the observations from the standard normal distribution (mean=0 and standard deviation=1). We can see the higher number of observations at the mean and the fatter tails as well as less observations around 2 standard deviations.

Graphs shown in Peters (1991, 1994) looked similar. The original proposal by Mandelbrot (1964) stated that the market distribution was Stable Paretian rather than normal. Such a

EXHIBIT 04 - FREQUENCY DISTRIBUTION OF S&P 500 DAILY RETURNS (Z-SCORED) (JANUARY 1988 - AUGUST 2014) LESS THE NORMAL DISTRIBUTION





distribution has a high peak at the mean and fat tails. While the mean is stable for a Stable Paretian distribution, the variance is undefined. This means in an unconditional sense, a sample variance, or standard deviation, will not converge to a population standard deviation since the latter does not exist. Of course, this is exactly what we see with the sequential standard deviation shown in Exhibits 02 and 03. This leads us to the conclusion that using standard deviation as a measure of risk can be misleading since the sample standard deviation typically used as a measure of risk will not be stable and is sample dependent. The implications of this on portfolio theory are so far-reaching that it has been mostly ignored by the finance community and the dangerous practice continues of using standard deviation as a stable measure of risk.

However, the sequential standard deviation graphs hold a potential solution to this issue. We can see that there are long periods where the sequential standard deviation steadily declines implying that those periods are more stable than the disruptions caused by fat-tailed events. We can postulate that there are, perhaps, two periods: one with a stable variance and another that is unstable. This would be compatible with the Fractal Market Hypothesis of Peters (1994).

The Fractal Market Hypothesis

The Fractal Market Hypothesis of Peters (1994) proposes that:

- (1) The market consists of many investors with different investment horizons.
- (2) The information set that is important to each investment horizon is different. The longer-term horizons are based more upon fundamental information, and shorter-term investors base their views on more technical information. As long as the market maintains this fractal structure, with no characteristic time scale, the market remains stable because each investment horizon provides liquidity to the others.
- (3) When long-term investors begin to question the validity of their information,

their investment horizon shrinks making the overall investment horizon of the market more uniform.

(4) When the market's investment horizon becomes uniform, the market becomes unstable because trading becomes based upon the same information set which is interpreted in a more uniform way. So good news causes increased buying while bad news results in increased selling.

(5) Liquidity dries up causing high volatility in the markets because most of the trading is on one side of the market.

(6) Eventually the long term becomes more certain and stability returns to the market as investment horizons broaden and become more diverse.

The Fractal Market Hypothesis (FMH) gives a fundamental underpinning to the volatility regimes we see in the SSD graphs. It implies that at those times when investment horizons are diverse, the markets should trade in an efficient and orderly fashion; while at times of crisis, the market will be prone to high volatility and a more disorderly trading environment.

Recently, Anderson and Noss (2013) produced a mathematical model showing the properties of the FMH. Li, Nishimura and Men (2014) have simulated the FMH using an agent-based model. Kristoufek (2013) showed empirically that the FMH described market activity during the bursting of the Tech Bubble in 2000 and the Credit Crisis of 2008 very well.

The Two-State Market

From the FMH and the SSD graphs we can postulate that the market has two states rather than one as assumed by the CAPM and other market models. There is a low uncertainty state which is fairly stable and exhibits what might be considered stable statistical behavior, and a high uncertainty state which is unstable and does not lend itself to standard statistical analysis.

The question is can we identify one state from another in an objective fashion? Clearly volatility



or standard deviation is a sign of change, but is not the cause. The cause is a rise in uncertainty which then gives rise to increased volatility in market trading. Increasing standard deviation is similar to an increase in the temperature of a person infected with a disease. The fever is not the cause of the illness, but is instead a symptom, usually one of many; but it may be the best sign of infection. In a similar way, increased volatility can also be a symptom of increased uncertainty. But standard deviation is problematic in that it measures what has happened in the recent past, and a number of past observations are required for its calculation. Realized volatility is forward looking only if we think past volatility is a good indicator of future volatility. Unfortunately, this is often not the case. What is needed is a contemporaneous measure of expected volatility.

Luckily, we have such a measure. Volatility (or the standard deviation of returns) is one input to the Black and Scholes (1973) option pricing formula. By using current option prices, the expected volatility implied by the current option price can be calculated. Since 1986, the Chicago Board of Options Exchange has published an implied volatility index (VIX) of S&P index options which essentially gives us a measure of the rise and fall of market uncertainty to an extent that it

has been nicknamed the “Fear Index.” This oversimplifies the nature of the index, but the VIX is a significant indicator of market uncertainty since it reflects the cost of hedging a broad index of stocks. The VIX has become so popular that there are now numerous volatility indices for other markets available, as well as futures contracts on the level of the VIX itself.

Below is a composite VIX which consists of a continuous maturity three-month VIX future, a three-month moving average of the EuroStoxx VIX (V2X), and a three-month moving average on an oil ETF (OVX). Exhibit 05 shows this composite VIX (CVIX) from January 1990 to June 2014.

We have shaded the areas where the CVIX is above its long-term median of 21.7.

With this graph, we can see that implied volatility spends long periods above and below its median and that these periods roughly correspond to the periods of instability that we saw in Exhibit 03 which covers the same time period. These periods also roughly correspond to periods of economic uncertainty. The first comes from Gulf War I, the second from the Emerging Market/Tech Bubble crises and the third from the Credit Crisis. As an outside measure of market instability, we will use these periods of above-and below-median

EXHIBIT 05 - COMPOSITE VIX
(JANUARY 1990 - JUNE 2014)



Sources: Data Stream, Global Financial Data

levels for the CVIX to approximate the high and low uncertainty periods described in the FMH. It is important to emphasize that using the median of the CVIX to define the states is an approximation. As we will discuss below, the actual transition is not known and probably changes with each business cycle.

Two-State Frequency Distributions

First we look at the basic characteristics in the two sub-periods. From a straight return and risk standpoint the two environments are very different from each other. As Table 01 shows, returns are significantly higher in the low uncertainty environment even as volatility declines.

TABLE 01: ANNUALIZED RETURN AND RISK BY UNCERTAINTY STATE
DAILY S&P 500 (JANUARY 1988 - AUGUST 2014)

	High	Low
Return	2.88	17.02
Risk	23.62	11.53
Return/Risk	0.12	1.48
Observations	3065	3403

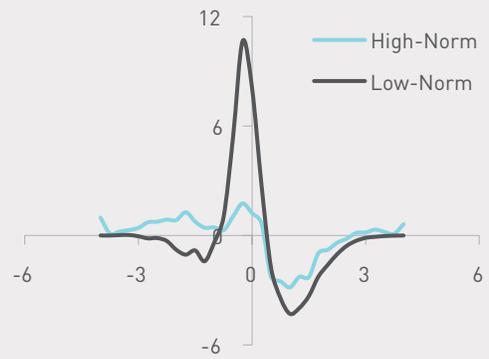
Sources: Data Stream, Global Financial Data

Given the difference in these very basic statistics, we would expect that the return distributions would also look quite different. So next we examine the frequency distributions of returns in the two sub-periods. We will use the same z-scored data points used in Exhibit 04 but divide them into the two states. The high uncertainty state has 3065 observations, while the low uncertainty state contains 3403 observations.

Exhibit 06 is the difference between these sub-periods and the normal distribution.

This chart illustrates that all of the “fat-tail” observations in both the positive and negative side of the distribution are observed in high uncertainty, while the vast majority of the increased number of observations about the mean occur in low uncertainty. In other

EXHIBIT 06: DIFFERENCE BETWEEN HIGH AND LOW UNCERTAINTY STATES AND THE NORMAL DISTRIBUTION
DAILY S&P 500 (JANUARY 1988 - AUGUST 2014)



Sources: Data Stream, Global Financial Data

words, we can see that the probability of a fat-tail event increases dramatically in high volatility. If we put this in conditional or Bayesian terms we can say that given we are in a high uncertainty state, the probability of a large event is much higher than in a low uncertainty state. The probability of observations larger than 3 standard deviations on the downside is approximately 0.14% in the normal distribution. For the S&P 500 in this time period, the probability of a 3-sigma downside event in low uncertainty is 0.12%, or slightly lower than normal, but in high uncertainty this increases to 2.12%. This is not a low probability. In a 260-day trading year, it means that a 3-sigma event or higher downside event would happen approximately once every 780 days (3 years) in low uncertainty, but once every 50 days in high uncertainty. The chance of a 2-sigma downside event is 2.28% according to the normal distribution. For the S&P 500, this probability is 0.85% in low uncertainty and 7.44% in high uncertainty. The upside probabilities, while not as dramatic as the downside, are also much higher in the high uncertainty state. Given that the upside probabilities are shown to be lower than the downside probabilities, we can see why the skew of the market is negative. This is also a

EXHIBIT 07: SEQUENTIAL STANDARD DEVIATION FOR THE LOW UNCERTAINTY STATE DAILY S&P 500 (JANUARY 1988 - AUGUST 2014)



Sources: Data Stream, Global Financial Data

EXHIBIT 08: SEQUENTIAL STANDARD DEVIATION FOR THE HIGH UNCERTAINTY STATE DAILY S&P 500 (JANUARY 1988 - AUGUST 2014)



Sources: Data Stream, Global Financial Data

function of the high uncertainty rather than the low uncertainty environment.

Sequential Standard Deviation by Uncertainty State

We can also ask about stability. Exhibit 06 and Table 01 both show that the low uncertainty state has less volatility than the high uncertainty state. But is low uncertainty more stable than high uncertainty as predicted by the FMH? In Exhibit 07, we show the SSD for the low uncertainty state. It does nicely converge to the daily equivalent of the volatility shown in Table 01 within 1000 days.

On the other hand, Exhibit 08 shows that the High Uncertainty state remains unstable and never converges to a stable value.

As predicted by the FMH, the high uncertainty state is where market instability lies. So we can now think of the market in conditional terms. That is, we can look at probabilities and risk given the current uncertainty environment. It also means that when it comes to long-term risk assumptions, we can say that for the S&P 500 about 11.5% is a reliable number in low uncertainty; but in high uncertainty, there is no stable assumption for the level of risk.

Conditional Kurtosis and Skewness

Skewness and kurtosis are standard measures for a distribution. Skewness tells us whether the returns are more likely to be positive or negative. The normal distribution is a random process so the likelihood of a positive or negative event is equal for a skewness value of zero. Kurtosis tells us whether the tails of the distribution are fatter or thinner than the normal distribution if kurtosis is positive or negative, respectively.

In Peters and Miranda (2014), we discussed the calculation of conditional skewness and kurtosis. The test was a modification of the traditional skewness and kurtosis calculation by using the mean and standard deviation of the total period to calculate the skewness and kurtosis statistics for a subsample. The traditional approach would use the subsample mean and standard deviation with the assumption that for a large enough subsample, they should be good estimates for the population, or true mean and standard deviation. But as we saw in Table 01, the sample standard deviation and mean for the high and low uncertainty subsamples are quite different from one another despite the large sample size;



so using the sample statistics makes no sense. By using the population mean and standard deviation (whole sample vs. subsample), we can calculate skewness and kurtosis in a way that attributes them to the subsamples.

TABLE 02: CONDITIONAL KURTOSIS AND SKEWNESS FOR THE S&P 500 IN HIGH AND LOW UNCERTAINTY STATES (JANUARY 1988 - AUGUST 2014)

	Skew	Kurtosis
High	-0.1689	21.4935
Significance	-3.8190	242.9729
Low	-0.0085	-2.2523
Significance	-0.2020	-26.8387
Total	-0.0845	9.0118
Significance	-2.7760	147.9988

Sources: Data Stream, Global Financial Data

Table 02 shows the conditional kurtosis and skewness calculations for the high and low uncertainty states in addition to the total period. As we found in Peters and Miranda (2014), for the MSCI World Equity index using monthly returns, the daily returns of the S&P 500 show negative skewness and high kurtosis in the high uncertainty state. In the low uncertainty state, the distribution is symmetric and kurtosis is negative, meaning that the tails are thinner than the normal distribution. This confirms the visual representation in Exhibit 06. Basically, all of the negative skew and fat-tail risk attributed to the market occurs in the high uncertainty state in this 6,469 day sample.

TABLE 03: CONDITIONAL ANNUALIZED RETURN AND RISK (1990 - 2013)

	Global Stocks	Small Cap Stocks	EM Stocks	REITS	Global Bonds	Credit	Commodities
High Return	-3.55	2.30	-4.00	8.19	2.82	3.53	-0.76
High Risk	18.65	24.57	29.11	22.74	3.50	12.50	25.31
Low Return	10.27	11.85	16.50	16.08	3.03	7.12	6.71
Low Risk	9.44	12.71	16.92	13.05	3.02	4.29	16.07

Sources: Data Stream, Global Financial Data

DEFINITION: Global Stocks is the MSCI World (local currency), Small Cap Stocks is the Russell 2000 Index, EM Stocks is the MSCI Emerging Markets Index, REITS is the FTSE NAREIT All REIT Index, Global Bonds is the Citigroup World Government Bond Index, Credit is the BofA ML High Yield Master Index and Commodities is the Bloomberg Commodity Index (Total Return).

Of course, this does not mean that there is no downside risk in the low uncertainty state. There are still negative returns. But these tests show that crashes and stampedes from January 1988 to August 2014 all occurred when the CVIX was above its median. This suggests increases and decreases in tail risk may be anticipated because the conditions for large events can be measured. In this case, we used one measure, the CVIX. Knowing when the CVIX is above or below its median gives us one condition which historically indicates the probability of a fat-tail event is high vs. low.

CAPM vs. Stable Paretian Models

The CAPM states the expected excess returns increase with risk as investors expect to be compensated for taking additional risk. Risk is typically defined as the standard deviation of returns. The Stable Paretian Hypothesis states that standard deviations are undefined due to the fat-tail high peaked structure of the market return distribution. The two-state approach reconciles these two models. In an unconditional sense, standard deviation is undefined. However, given we are in a low uncertainty environment, standard deviation is finite and expected returns increase with risk. In a high uncertainty environment, this no longer holds. In Table 03 are annualized excess returns and risk for some standard asset classes based upon monthly returns from 1990 - 2013. We can see that in low uncertainty, return and risk behave as expected. Except for commodities, higher risk is compensated with higher returns in periods of low uncertainty.

In high uncertainty, there appears to be little relationship between return and risk.

Peters and Miranda (2014) also found that kurtosis and skewness for these same asset classes followed the same pattern as Table 02. That is, except for bonds, the other asset classes had high kurtosis and negative skew in high uncertainty and negative kurtosis and no skew in the low uncertainty state. Table 03 confirms that partitioning the data according to the market state produces those characteristics as well.

Diversification

Perhaps the largest impact on asset allocation of the two-state model is that what diversifies in low uncertainty loses much of that power in high uncertainty when it is really needed. Table 04 is a correlation matrix of the assets in Table 03 over the same period, but partitioned into periods of high and low uncertainty.

Across the equity assets, we can see that correlations increase in high uncertainty.

Credit, or high yield bonds, acts like an equity. Commodities become more highly correlated with the smaller equity components. Bonds become negatively correlated to all the other assets in high uncertainty offering the only true diversification in this group; though in low volatility the correlations become slightly positive. In 2008, there was surprise that diversification “failed” as many assets which were considered diversifying all declined together. We can see in Table 04 that this is a normal part of the cycle, and one for which investors should be prepared.

Summary of the Empirical Results

Based upon these results, we can postulate that the S&P 500 does, indeed, have 2 distinct states as predicted by the FMH. The high uncertainty state is characterized by high, unstable volatility, fat tails, negative skew, and low returns. The low uncertainty state, by contrast has low and stable volatility, thin tails, no skew and high returns. The data was partitioned using

TABLE 04: CONDITIONAL CORRELATION
(1990 - 2013)

	High Uncertainty						
	Global Stocks	Small Cap Stocks	EM Stocks	REITS	Global Bonds	Credit	Commodities
Global Stocks	1.0000	0.7912	0.7981	0.5557	-0.2249	0.6658	0.1370
Small Cap Stocks		1.0000	0.7568	0.6979	-0.1992	0.6880	0.2458
EM Stocks			1.0000	0.5381	-0.2453	0.6682	0.3003
REITS				1.0000	-0.0485	0.6615	0.2267
Global Bonds					1.0000	-0.0930	-0.2492
Credit						1.0000	0.2281
Commodities							1.0000

	Low Uncertainty						
	Global Stocks	Small Cap Stocks	EM Stocks	REITS	Global Bonds	Credit	Commodities
Global Stocks	1.0000	0.6504	0.5549	0.3153	0.1569	0.3579	0.1606
Small Cap Stocks		1.0000	0.4865	0.5115	0.0657	0.3664	0.0722
EM Stocks			1.0000	0.2258	0.1223	0.3271	0.1493
REITS				1.0000	0.2264	0.4522	-0.0521
Global Bonds					1.0000	0.4218	-0.0440
Credit						1.0000	0.0492
Commodities							1.0000

Sources: Data Stream, Global Financial Data

DEFINITION: Global Stocks is the MSCI World (local currency), Small Cap Stocks is the Russell 2000 Index, EM Stocks is the MSCI Emerging Markets Index, REITS is the FTSE NAREIT All REIT Index, Global Bonds is the Citigroup World Government Bond Index, Credit is the BofA ML High Yield Master Index and Commodities is the Bloomberg Commodity Index (Total Return).



an index of implied volatility which is not directly related to the S&P 500 return data though it is, of course, correlated to it. The CVIX is derived from option pricing rather than being directly computed from S&P 500 index returns, and thus reflects the cost of hedging the S&P 500.

Since the low uncertainty period is also similar to periods of economic expansion and bull markets, these findings explain why many investment models work well in the low volatility period but fail during the bear markets. The entire statistical structure of the market changes to such a degree that an entirely new model would likely be called for.

For markets in general, the risk/return trade-off postulated by the CAPM is largely contained in the low uncertainty state. In high uncertainty, those conditions do not apply and there appears to be little, if any, relationship between the risk of an asset and its realized return. In addition, the ability to diversify changes in the two states, and the assets that diversify in the low uncertainty state, offer smaller benefits in high uncertainty.

What Does This Mean?

Earlier we used the old head-in-the-refrigerator-feet-in-the-oven joke to describe the markets. Of course, even a two-state market would not have two discrete states. The move from high to low uncertainty is not a phase shift similar to when water boils at 100 degrees centigrade and transforms from a liquid into a gas. There is likely a period of transition from one state to the other. Contrary to popular imagination, the stock market does not suddenly collapse for economic reasons (political events, natural disasters and other high impact non-economic events are a different story). There are always warnings. In fact, if we need an analogy for the market cycle, we should use the weather instead. Most of the Earth has two real seasons: summer and winter. Autumn and spring are transitions from winter to summer and back again. However, the “start” of summer or winter is not obvious by looking at

the weather itself. The calendar developed in the western world uses lunar movements and the tilt of the Earth to mark the first day of summer and winter, but we have all seen significantly colder than average first days of summer and similarly warm first days of winter. The market transitions from low to high uncertainty in a similar way. Unfortunately, we do not have a calendar to give us even a broad idea of when that transition occurs so we have to use other means.

In this paper, we used the median of the CVIX as a cut-off point creating two discrete states. While fine for the purposes of this study, it would be dangerous to use the CVIX as a lone indicator in practice. Examine the graphical evidence in Exhibit 08, the SSD of the high uncertainty state. There are some periods of stability there just as the SSD of the low uncertainty state in Exhibit 07 has some significant “bull market corrections” in it. But for the purposes of this paper, we can clearly see that there are two main states for the market. Those two states are very different and so investment management during those two periods should be different. Unfortunately, investment management strategies base their approach on long-term averages and so dress for spring even if it is winter or summer outside.

This leads us to the need for a new market model which assumes two distinct states.

Amending the FMH

A two-state model is already described in the FMH, but now we can amend the FMH to add:

(7) During periods of low uncertainty, markets will exhibit well-behaved, finite variance statistics while in high uncertainty markets will exhibit fat-tailed risks and unstable variance more associated with the stable Paretian distribution as described by Mandelbrot (1962).

This reconciles the market models associated with CAPM and those associated with Fractal analysis. Essentially, in low uncertainty, the traditional methods of quantitative analysis using standard deviation as a measure of risk



and covariance to describe diversification will work quite well. However, in high uncertainty, the stationary conditions needed for that type of analysis no longer hold. Assuming a stationary process in high uncertainty can be misleading and even dangerous to investment strategies.

Summary

In this paper, we found that the markets have two very different states with very different conditions. Assuming the long-term average as “normal” can lead to very misleading conceptions of risk as well as estimates of return. Unfortunately, the transition from one state to another is not well defined and likely changes with each cycle. In this paper, we approximated the transition point using implied volatility indices, which would be appropriate if the stock market were always the trigger. However, there are likely other triggers such as financial or macroeconomic conditions. This would lead to the necessity for a methodology to better estimate the transition in real time if this knowledge is to be practical. That will be the subject of another forthcoming paper.



References

- Anderson, N. and Noss, J. (2013) “The Fractal Market Hypothesis and its implications for the stability of financial markets,” Bank of England Financial Stability Paper
- Ang, A. and Timmerman, A. (2011) “Regime Changes and Financial Markets,” NBER Working Paper Series
- Chow, E. Jacquier, E., Kritzman, M., and Lowry, L. “Optimal Portfolios in Good Times and Bad,” Financial Analysts Journal, May/June 1999
- Fama, E. (1964), “Mandelbrot and the Stable Paretian Hypothesis,” in P. Cootner ed., The Random Character of Stock Prices, Cambridge: MIT Press
- Kristoufek, L. (2013), “Fractal Market Hypothesis and Global Financial Crisis: Wavelet Power Analysis,” Scientific Reports 3
- Li D., Nishimura Y., and Men M. (2014), “Fractal markets: Liquidity and investors on different time horizons,” Physica A 407
- Lintner, 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
- Mossin, J. (1966), “Equilibrium in a Capital Asset Market,” Econometrica 34
- Peters, E. (1991), Chaos and Order in the Capital Markets, New York: Wiley, 1991
- Peters, E. (1994), Fractal Market Analysis, New York: Wiley, 1994
- Peters, E. and Miranda, B. (2014) “How tail risk changes over the market cycle,” FQ Perspectives
- Sharpe, W. (1964) “Capital asset prices: A theory of Market equilibrium under conditions of risk” Journal of Finance 3