Calibrating Event Risk with the Market Sentiment Meter

Blu Putnam, Chief Economist, CME Group

18 July 2019

Event risk comes in a number of different forms, with perhaps some of the more difficult risk management challenges being posed when market participants split into two divergent camps associated with strikingly different views of the world. That is, there are two conflicting scenarios for how the future may develop and both have meaningful probabilities. This type of two-scenario event risk has developed in association with the Brexit referendum in the UK in June 2016, the US elections in November 2016, in the aftermath of the corn belt drought of 2012 as worries about whether 2013 would bring the same or a return to rain, the controversy over the OPEC production decision in November 2014, the deal or no-deal possibilities related to the US-China trade war in the spring of 2019, and so on. Some of these examples had specific dates on which the outcome would become known (e.g., elections, referendums, OPEC meetings), while others operated in a more nebulous time frame during which a resolution was expected (e.g., trade war, drought). In every case, though, the key market characteristic was the presence of a pre-event bi-modal risk probability distribution reflecting highly conflicted, distinctly different scenarios, both with meaningful probabilities. Once the outcome was known, markets quickly reflected the new reality and resolved to a single-mode, bell-shaped risk probability distribution centered on the outcome.

The risk management challenges of this conflicted two-scenario event risk are complex. First, there is a distinct possibility of a quick and abrupt price movement when the outcome of the conflicted two-scenario situation resolves to one side or the other. Price gap risk is not easily hedged and may cause special problems for options-related delta hedging strategies that have embedded the assumption that market prices operate in a smoothly continuous fashion and sharp price breaks are never allowed to occur. Second, the meaning of implied volatility derived from options prices may be distorted. The problem here is that common options pricing models based on Black-Scholes[1] or Merton[2] theories assume away the possibility of price breaks (i.e., prices are continuous), which means that when two-scenario event risk is present, many market participants may choose to incorporate price breaks into their expectations. When price break expectations are present and implied volatility is calculated from an options pricing model that assumes no price breaks, then the implied volatility is over-stated by the degree of the price break expectations. That is, one number, implied volatility, contains both expectations of future volatility and expectations of a one-off major shift to a new center of gravity for price expectations.

To put the risk management challenges of conflicted two-scenario event risk in perspective, let's look at two different probability distributions custom-tailored to have the same expected mean and same expected standard deviation. We doubt anyone would expect the use similar risk management strategies when faced with probability distributions that are so strikingly different.

Normal Distribution (μ=0%,σ=12%).

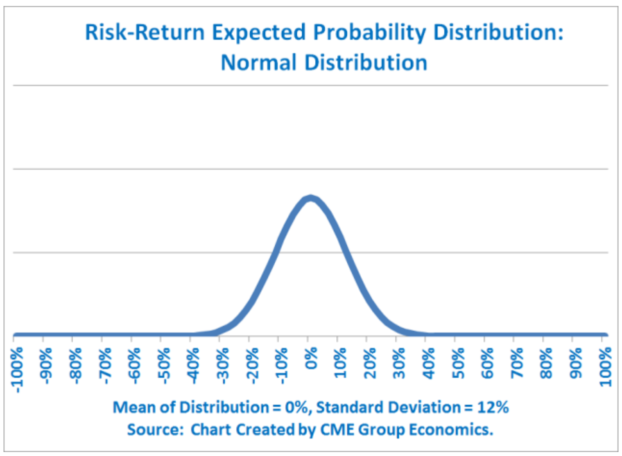

Bi-Modal Distribution (μ=0%,σ=12%).

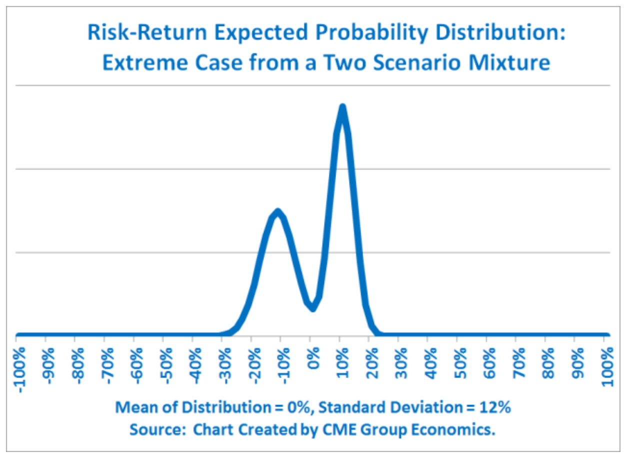

The challenge for risk managers is whether their quantitative systems can even identify whether the conflicted, two-scenario, bi-modal probability distribution exists, since we are dealing with the unobservable. And then, after we construct our hypothetical distributions, we need to assess the degree and intensity of the event risk. Our system uses a two scenario mixture probability distribution or double gamma distribution to allow for the possibility of a bi-modal situation. We measure the degree and intensity of the bi-modal probability distribution by examining two metrics: the distance between the two modes, and the area of the probability distribution between the two modes.

[1] Black, Fischer; Scholes, Myron (1973). "The Pricing of Options and Corporate Liabilities". *Journal of Political Economy*. 81 (3): 637–654. doi:10.1086/260062.

[2] Merton, Robert (1973). "Theory of Rational Option Pricing". *Bell Journal of Economics and Management Science*. 4 (1): 141–183. doi:10.2307/3003143.

We also should note that event risk comes in other types than the special case of conflicting scenarios which we are studying. For example, considerable research has been done into extreme risk cases, where there is a very, very small probability of an event with massive consequences. The earthquake and tsunami that hit Japan in 2011 would be an example of extreme value risk. Extreme value analysis is most applicable for these exceptionally low probability cases, and risk management approaches often focus on an insurance model using deep out-of-the-money options.

There are also other types of risks that may seem like event risk, but would not qualify under our definition. That is, one might know the time and date of an important data release, such as the monthly employment situation report in the US. Just knowing a date, however, is not a sufficient criterion to earn the label of event risk. Most data releases are best described by bell-shaped probability distributions, as there is usually a strong consensus around an expected mean with an acknowledged appreciation of the volatility present in the specific data. Three or four standard deviation events may certainly occur, and, indeed, they seem to occur much more often than suggested by normal distributions. Again, an extreme value analysis in which the tails of the distribution are augmented is probably appropriate. Our system incorporates some important features of extreme value analysis and has the ability to identify and quantify unusually skewed or fat-tailed distributions.

I. Philosophy of our Joint Research with 1QBit and CME Group

Our risk-return probability research has been conducted jointly with 1QBit, a Vancouver-based machine learning and quantum software company. The curated data from this research, powered by 1QBit, known as the "Market Sentiment Meter", will formally be available for subscription purchases in September 2019 on CME's DataMine platform for a number of futures products, including the E-Mini S&P500® equity index, US Treasuries, Euro FX, gold, oil, natural gas, corn, and soybeans. A free trial package is available now through CME Group Economics (contact: Bluford.Putnam@CMEGroup.com).

In effect, we have reimagined probability risk distributions. While we use information derived from options market activity, our perspective is quite different from models based on implied volatility. Our research approach contains four key tenets, which we will discuss, and then provide a few case studies to make the practical applications clear. Our four key beliefs are as follows:

- We do not accept that standard deviations are adequate measures of risk.
- We believe that starting points matter. Starting one's risk analysis with implied volatility introduces some hidden biases that may be surprisingly hard to overcome.
- We believe that any practical risk system must be capable of capturing two-scenario event risk with very special characteristics.
- Probability risk distributions are inherently unobservable, yet we believe that a variety of price and volume metrics can be examined to make some useful inferences about the risk distribution and how it shifts through time.

II. Volatility is Not Risk

To begin with, volatility is a poor measure of risk.[3] Many analysts like volatility because the historical standard deviation is easy to calculate and fits nicely into basic risk systems and mean-variance portfolio optimization models[4].

One problem is that an investor, or a financial institution for that matter, may have asymmetrical risk preferences, preferring to avoid substantive losses rather than to make equivalently large gains. That is, if avoiding large losses is the primary risk, then a symmetrical standard-deviation based metric that only looks at the average noise level and not the extremes is certainly not appropriate.

Another challenge, already highlighted, is that implied volatilities are typically calculated from straightforward options pricing models that embed the heroic assumption that prices move up or down with continuous trading – that is, price breaks or price gaps are assumed never to occur. If market participants fear the possibility of price breaks or gaps, options prices will reflect this risk with a higher calculated implied volatility. But it will not be easily apparent that the implied volatility is reflecting price gap risk instead of an upward shift in the volatility regime. And, price gap risk is not the same risk as volatility regime shift risk. Depending on one's financial exposures, one of these risks could be much more important than the other. For those managing options portfolios, for example, the risk of an abrupt price break can do considerable damage to delta hedging strategies, while a volatility regime shift represents a different risk, commonly known as "vega" risk. What one needs to create is a comprehensive view of the whole risk probability distribution providing a robust perception of risks, allowing for decidedly different risk scenarios, and not being biased toward bell-shaped curves.

III. Starting Points Matter

To build a risk probability distribution that is not necessarily bell-shaped or even of a single mode and can capture the extremes in a robust manner, we prefer to start from a very different point of view. We start with the Bayesian prior (i.e., our initial views before we even examine the data) of a very unusual distribution – in our case, a bi-modal distribution that might reflect a type of binary or two-scenario risk often associated with event risk. Then, we examine market data to see if the risks are more bell-shaped. While the implied volatility is one of the market metrics we examine, it does not necessarily have the primary influence it does when it is the starting point for the risk analysis.

Put another way, if we start from a prior of an extreme and unusual distribution, we know that it can exist, and we have not assumed it away. Starting from a standard deviation approach, such as implied volatility, may inadvertently make it very hard to estimate when extreme and highly dangerous risk distributions are present. The math behind this observation is quite old and goes back to the Russian mathematician, Pafnuty Lvovich Chebyshev (1821 – 1894). What most people take away from Chebyshev's Inequality Theorem is that if you know only the standard deviation you have a very good

[3] Please see Chapter 9, "Volatility & Uncertainty" in the recently published book, ***Economics Gone Astray*** (World Scientific Publishing Co.), by Blu Putnam, Erik Norland, and KT Arasu, 2019.

[4] Please see Chapter 11, "Portfolio Optimization", ***Economics Gone Astray*** (World Scientific Publishing Co.), by Blu Putnam, Erik Norland, and KT Arasu, 2019.

idea of the typical ranges in which values will fall most of the time. What we take away from the Inequality Theorem is that if you only know the standard deviation, you know absolutely nothing about the extremes of the distribution where the most dangerous risks reside.

Pafnuty Lvovich Chebyshev, 1821-1894 (from Wikipedia).

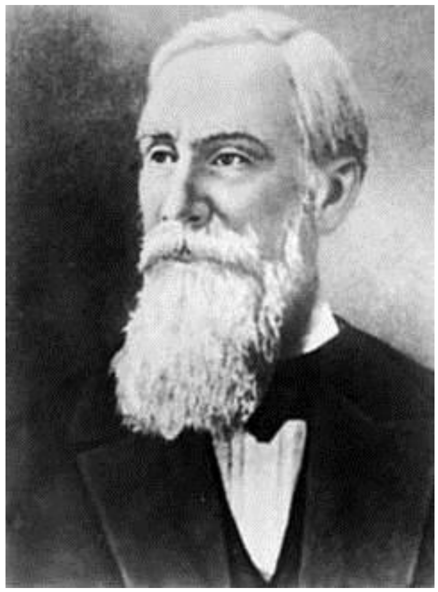

IV. Conflicted Two-Scenario Event Risk Has Special Characteristics Not Captured by Standard Deviation Metrics

The motivation for our research was the observation that in financial markets, especially since 2016, we have been seeing important episodes of event risk associated with elections – UK Brexit Referendum of June 2016, US President election of November 2016, French and UK elections in 2017, Brazilian elections of October 2018, US Congressional elections of November 2018, etc. This led us to a study of how markets cope with two strikingly different scenarios – a type of event risk. When there are two possible scenarios, then pre-event, the market is going to price the probability-weighted outcome, or the middle ground. So, post-event, when the outcome becomes known, the market immediately moves away from the middle ground to the "winning" scenario – a price break. For example, with Brexit, the "Leave" vote generated a sharp downward move in the British pound (vs USD), while a "Remain" vote would have presumably generated a sharp almost instantaneous rally in the pound – either way, the pound was no longer going to trade in the middle. Even if they are extremely rare, if one's risk system cannot create the possibility of a bi-modal probability distribution, then price break risk and tail risk may be greatly underestimated.

UK Brexit Referendum

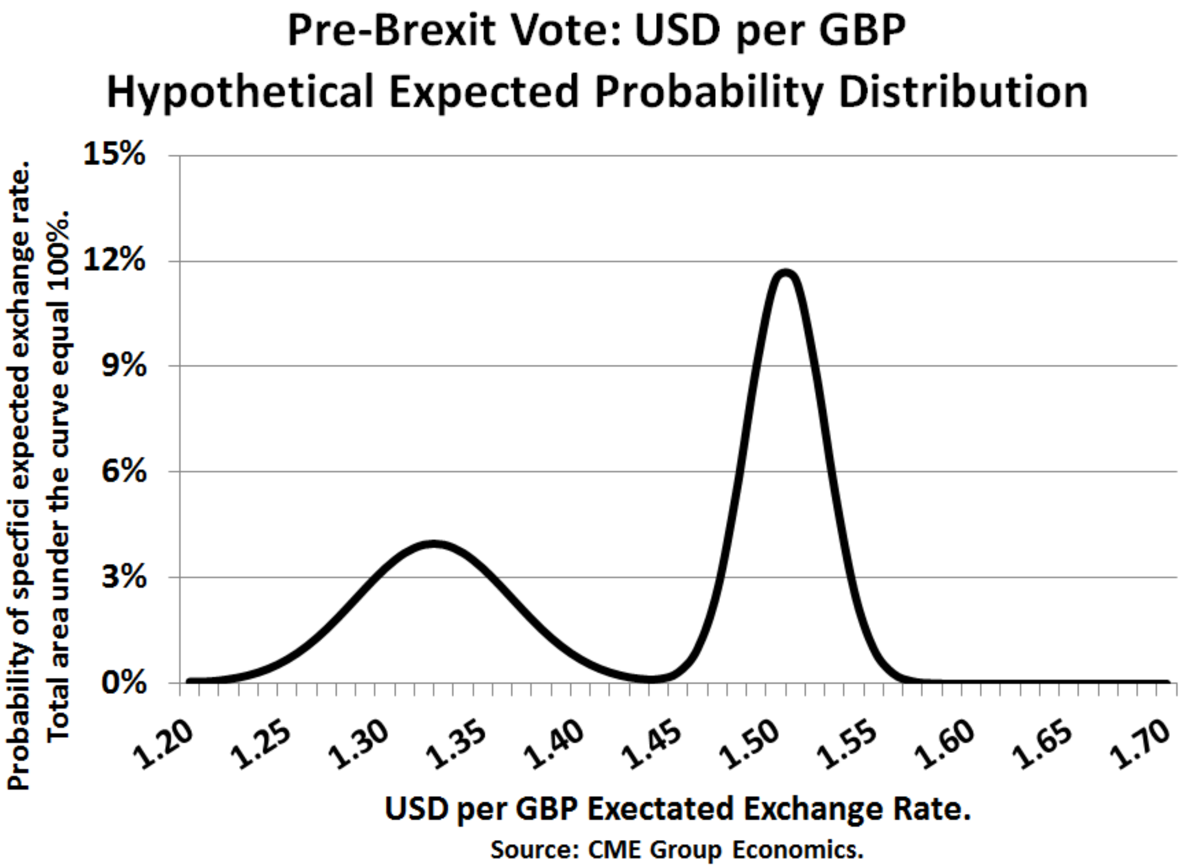

IV. Expected Probability Distributions are Unobservable

From a practical perspective, starting with the prior (e.g., our view of the world before examining any data) of an abnormal, bi-modal risk probability distribution requires some creativity that might put off some risk managers. The challenge is that expected risk-return probability distributions cannot be directly observed. What we can do is to estimate some of their characteristics from looking at market behavior – prices, volumes, futures versus options, intra-day activity, etc.

While our research is still at the early stages, we have found a few metrics that are especially enlightening relative to the shape of the probability distribution. Our three primary metrics are: (1) the evolving pattern of put option trading volume relative to call option volume, (2) intra-day market activity, especially high/low spreads, and (3) implied volatility from options prices relative to historical volatility.

Studying put/call volume patterns helps us understand if one side of the market is more at the center of the current debate than the other side. For example, immediately after former Federal Reserve (Fed) Chair Ben Bernanke threw his famous "Taper Tantrum" in May 2013, he set off a debate about when the Fed would withdraw quantitative easing (QE) and raise interest rates. Put volume on Treasury note and bond prices soared relative to call volume as an indicator that a two-scenario situation had developed. While there is a buyer and a seller for every trade, one side thought prices would fall (yields rise) and volatility might rise very soon (buyer of puts), while the other side thought the process of exiting QE would take a long time (seller of puts).

Intra-day market dynamics help us appreciate risk in a different way. The observed high price to low price intra-day trading spread is informative in helping us assess the degree to which fat-tails might be present. Mathematically, work by Mark B. Garmin and others back in the 1970s and 1980s has shown that if one assumes a normal distribution then there is a straightforward way to estimate the standard deviation of daily returns from the intra-day high-to-low spread. Put another way, if the relationship between intra-day dynamics and the day-to-day standard deviation diverge in a significant manner, then this is strong evidence that the risk probability distribution is not normally distributed.

To ascertain the risk of price breaks we track the evolving pattern of implied volatility relative to historical volatility. While it is usual for implied volatility to exceed recent historical standard deviations, a shift in the pattern toward a much higher implied volatility may indicate that expectations for the potential of a sharp price break are building in the market. And, if a price break occurs, scenarios resolve one way or the other, so post-outcome we often see a quick decline in the implied volatility representing a shift back to a single-mode bell-shaped distribution.

To gather all our risk information and create a probability distribution, we use a probability mixture technique that is distribution independent – that is, it is not constrained to take on a given specified shape. Most of the time, bell-shaped curves are appropriate descriptions of the probability distributions – balanced risk distributions. Our method does, however, occasionally generate some especially tall distributions (i.e., relatively lower volatility), which we classify as "complacent" and worthy of special study to see if the market may be underestimating risks. We also see on occasion some very flat distributions, not unlike the Wall Street maxim about the equity markets "climbing a wall of worry" which we call "anxious" risk distributions. And, finally, on rare occasions our metrics support the idea of a two-scenario, event risk, bi-modal distribution. That is, we classify expected risk distributions into four types: "Complacent" which are very tall and thin, "Balanced" or neutral risks with a typical bell-shape, "Anxious" reflecting a relatively flat bell-shape with very fat tails and possibly skewed one way or the other, and finally our bi-modal (aka, "Conflicted") or event risk distribution which are trying to anticipate what happens if one of two very divergent scenarios is the outcome.

Figure 7: Market Sentiment Meter categories of perceptions of risk

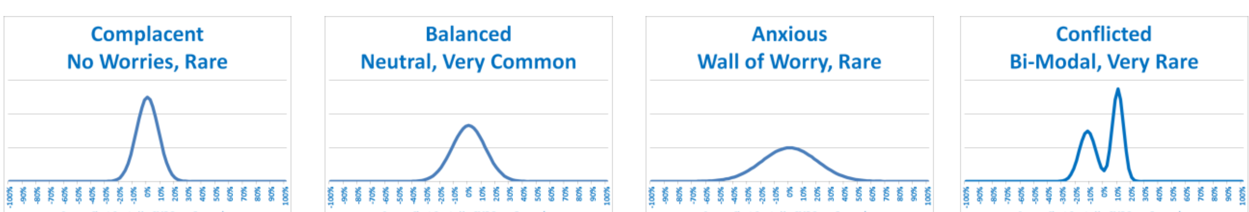

V. Case Studies

To illustrate our probability risk distributions, we will start by examining two examples from US equities, one involving a complacent distribution and one involving potential event risk. And, we will also examine an event risk distribution from the commodity markets, specifically in corn.

In late 2017, our probability risk distribution for US S&P500® (CME E-Mini Futures) shifted from "balanced" to "complacent". US stocks were being propelled higher in no small measure by the large and permanent US corporate tax cut which was increasingly likely to become law and, indeed, was passed by Congress and signed into law by the President in December 2017. Due to the corporate tax cut, market participants were expecting more stock buybacks and higher dividends, among other things. As it turned out, the complacency was somewhat misplaced. Early in 2018, the US-initiated trade war, with China, the European Union, Canada and Mexico, resulted in a sharp market selloff, temporarily higher volatility, after which the market that started to gain ground again with diminishing volatility.

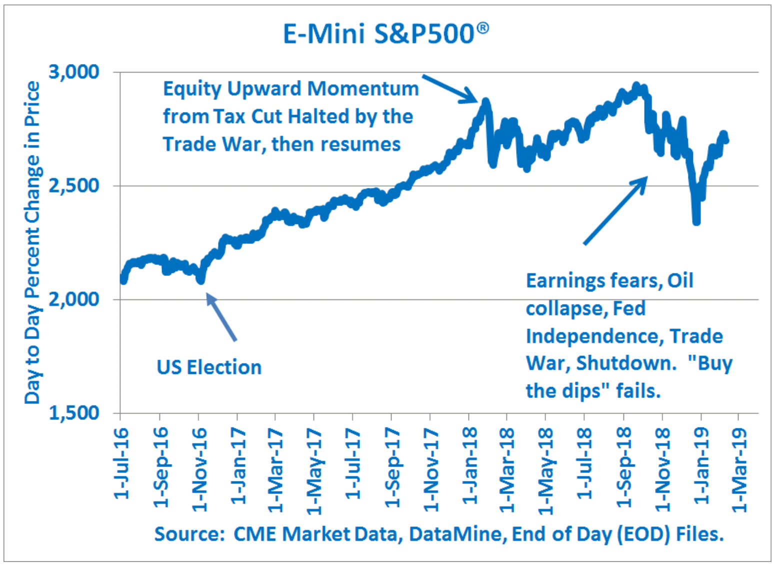

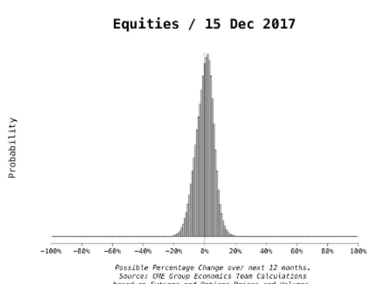

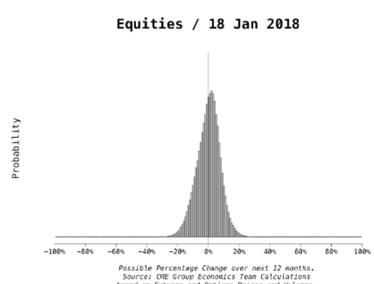

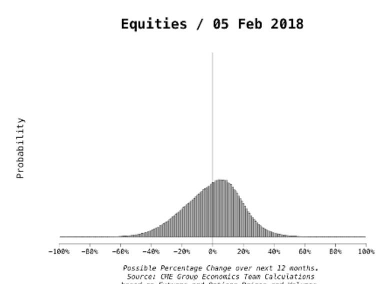

In late October 2016, our probability risk distribution for US E-Mini S&P500® futures shifted into the event risk state showing a bi-modal probability distribution. The cause was the upcoming US election, which was becoming increasingly polarizing and involved two candidates with strikingly different views on many issues, including a few of major interest to equity market participants, such as whether taxes would be cut or raised. Once the election outcome was known, within a week the probability risk distribution shifted quickly to an anxious state and then soon to a balanced risk state.

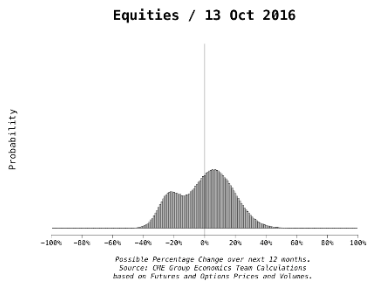

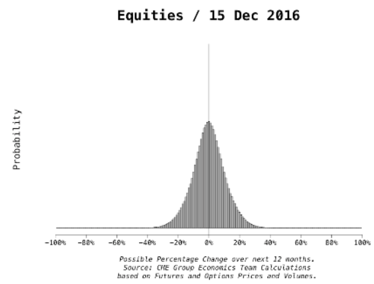

These two, polar opposite, probability risk distributions in the US equities futures market illustrate the method's ability to identity event risk (polarizing US Presidential and Congressional elections in November 2016) as well as to highlight a market state with a noticeable absence of fear, namely the "complacent" distribution, which occurred as the US corporate tax cut was being passed into law. Neither of these more extreme probability distributions lasted very long. Complacency in December 2017 and January 2018 gave way to fears related to the trade wars. And with the event risk from the US elections, the outcome resolved the debate about which scenario would be the winner, and the bi-modal distribution quickly resolved into a single-mode distribution.

Our next case covers a very interesting evolution of our probability risk distributions in the corn market in late 2012 and into the first half of 2013. The summer of 2012 had seen large swaths of the US corn belt experience severe drought. Late in 2012, after the harvest, market participants' thoughts turned to the 2013 crop, about which there was much disagreement. How much acreage would be planted after the drought year? Would 2013 see another drought or its disappearance? While not of the political version of event risk, corn market participants were worried about the drought and a two-scenario market developed for a while in February 2013 as one side of the market took the view that the 2013 crop would be much better than 2012's drought-constrained crop and other market participants worried about another poor crop. Our probability risk distribution was already in an "anxious" state late in 2012, shifted to "event risk" in February 2013, went back to "anxious" for most of the spring of 2013, before returning to the most common state, "balanced risks" in the summer of 2013.

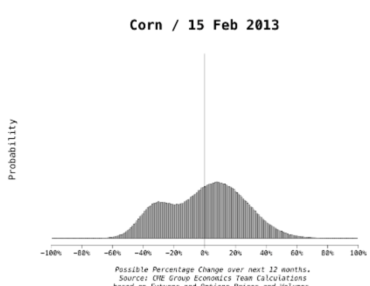

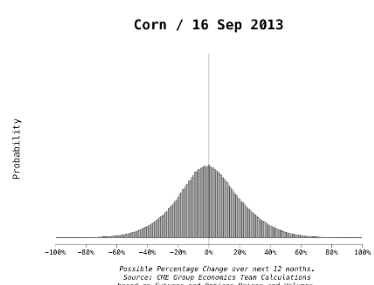

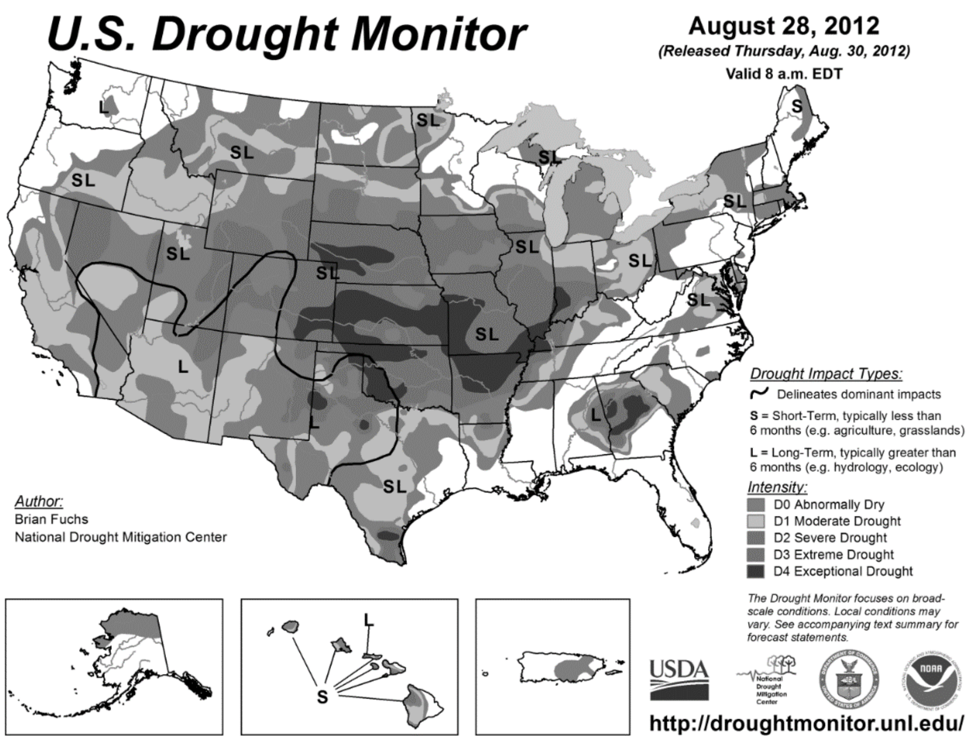

U.S. Drought Monitor
August 28, 2012
(Released Thursday, Aug. 30, 2012)
Valid 8 a.m. EDT
Drought Impact Types:
Delineates dominant impacts
S = Short-Term, typically less than 6 months (e.g. agriculture, grasslands)
L = Long-Term, typically greater than 6 months (e.g. hydrology, ecology)
Intensity:
D0 Abnormally Dry
D1 Moderate Drought
D2 Severe Drought
D3 Extreme Drought
D4 Exceptional Drought
The Drought Monitor focuses on broad-scale conditions. Local conditions may vary. See accompanying text summary for forecast statements.
Author:
Brian Fuchs
National Drought Mitigation Center
USDA
http://droughtmonitor.unl.edu/

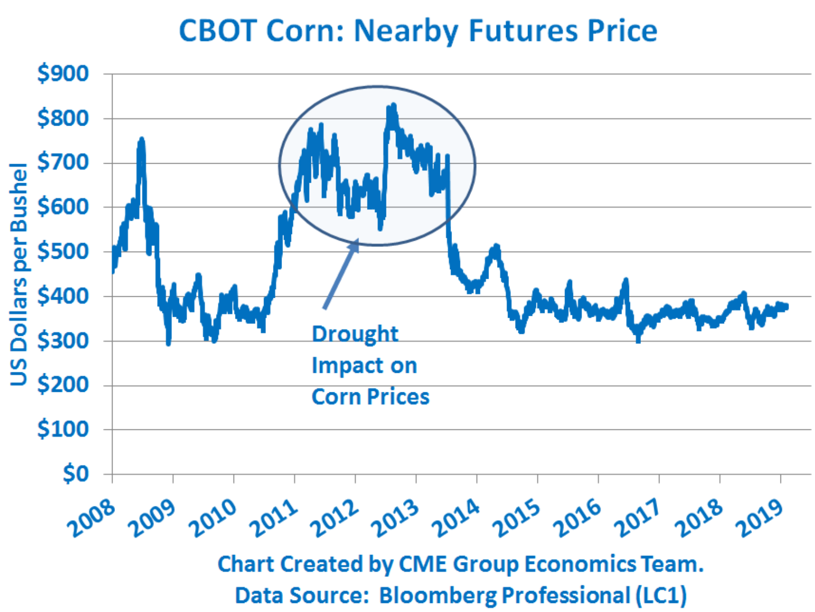

CBOT Corn: Nearby Futures Price
US Dollars per Bushel
$900
$800
$700
$600
$500
$400
$300
$200
$100
$0
2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019
Drought Impact on Corn Prices
Chart Created by CME Group Economics Team.
Data Source: Bloomberg Professional (LC1)

While these case studies are presented purely as illustrations, our research methods allow for the rarest of market states – event risk with a bi-modal probability distribution – to occur in all of the product classes we have studied so far, which includes US Treasury notes futures, equity index futures, Euro FX (versus USD), gold, oil, natural gas, soybeans and corn. And, we believe it is important to monitor our risk states, especially when they shift from one category to the next.

We do not expect the most common state – "balanced risks" occurring as much as two-thirds to three-quarters of the time, depending on the product, to provide any critical information that one would not acquire looking only at implied volatilities from options markets. We do think, however, that when the probability risk distribution shifts into a less typical state – "complacent", "anxious", or especially "conflicted event risk" – that risk managers should go on high alert.

We also warn that while our naming conventions describe the risk distributions, they may not describe what happens. "Complacent" states may well be followed by volatility when some new and unexpected risk factor takes priority. "Anxious" states may or may not overstate fears, as equity analysts talk about when the say "a market is climbing a wall of worry". "Conflicted or Event risk" states do not last long, as they tend to be resolved back to a one scenario, single mode distribution when the event occurs, and the outcome becomes known or when market participants become more confident that a one scenario outlook with appropriate skepticism is more appropriate than a conflicted two-scenario approach.

VI. Next Steps

An important milestone for our research will be reached in September 2019 when CME DataMine and 1QBit make the formal and final version of our curated daily data sets available based on our probability risk metrics, under the banner of Market Sentiment Meter. Daily data history goes back to January 2012. The data sets cover eight exchange-traded futures and options products:

- CME E-Mini S&P500®,
- CBOT US Treasury 10-Year Note,
- CME Euro FX,
- NYMEX WTI crude oil,
- NYMEX Henry Hub natural gas,
- COMEX gold,
- CBOT soybeans, and
- CBOT corn.

Each of the products will be available by yearly subscriptions, with discounts for ordering multiple products.

Our next step in the research process is to analyze our curated data sets to see what insights can be gleaned about future market activity, especially in the one to three months ahead time frame. Our research is focused on three types of shifts in market activity:

- Price gap risk: Is there a meaningful price change (up or down, that is directionally independent)?
- Vega risk or a volatility shift: Is there a shift in observed market volatility?
- Directional risk: Is there a meaningful percent change in the price in one specific direction?

The research is complex, because while the data sets may appear as time series daily data, what is really at issue is whether a given “episode”, of varying length, of a market sentiment state, such as bi-polar, anxious, or complacent, tells us anything about our research questions listed above relating to future market activity. Such tools as are now available with machine learning and artificial intelligence (e.g., a specialty of 1QBit) are likely to be exceptionally useful in this research. We note though, as financial risk disclosure statements always remind us, that past performance is not necessarily a guide to future performance.

As we make progress in our research, we will be making various research papers available for comment. And we hope a number of academics and practitioners will explore our free trial data set and share some of their insights.

Appendix – Trial Package

1) **CSV file** of E-Mini S&P index futures, daily data from 2012 through June 2019, including original futures and options price and volume data used in the analysis, market momentum metrics, market volatility metrics, high-low price spread metrics, put/call option volume metrics, and the mixture probability distribution including its sentiment states and various descriptive metrics.

2) White Paper: “Calibrating Event Risk with the Market Sentiment Meter”, CME Group, July 18, 2019.

3) “Commodity Risks: Describing the Unobservable”, Global Commodity Applied Research Digest, summer 2019, JPMorgan Center for Commodities, University of Colorado at Denver.

4) “Managing Risk in the Era of Dissonance”, CME Group & Hedge Fund Magazine, March 2019.

5) “The Changing Nature of Event Risk”, AIMA (Alternative Investment Management Association), January 2019.