Managing Risk in the Era of Dissonance

March 2019

Blu Putnam, Chief Economist, CME Group

All examples in this report are hypothetical interpretations of situations and are used for explanation purposes only. The views in this report reflect solely those of the author and not necessarily those of CME Group or its affiliated institutions. This report and the information herein should not be considered investment advice or the results of actual market experience.

The era of dissonance is upon us. There are demography-related deflationary pressures as the post-WWII baby-boomer generation retires and debt-laden millennials struggle to pick up the slack. There are also technological disruptions – in retail, the growing evidence of the shift away from brick-and-mortar stores to internet shopping; in energy, the phenomenon of the shale revolution; in the consumption of news and means of communications, we have smart phones and social media that have fundamentally altered how individuals interact with each other and within society, just to name a few disruptive transitions that are in progress. While the causes of economic and societal dissonance may have their roots in demographics and technology, their implications include the sharp political divisions being witnessed all over the world. We have the will of deeply divided electorates reflected in a variety of policy debates with implications for trade, immigration, health care and tax policy, for example.

Our focus here is that the era of dissonance has been accompanied by major challenges regarding how market participants manage risk. Disruptive transitions in economics and markets have similarities with the concept in physics known as phase transitions. In a phase transition, there is a change from one state or environment to another. Think of water turning from a liquid into vapor upon boiling, or a liquid flowing smoothly (laminar flow) in a pipe and then flipping to a turbulent flow. With phase transitions in physics, the chaotic activity is at the border or transition zone between the two states. In economics and markets, for example, the transitions driven by demographics and technology and reflected in a divisive political environment have meant that event risk is much more common along the transition or fault lines, and event risk may often overwhelm more slowly evolving economic fundamentals. In this research, we wish to highlight and discuss a set of key challenges that emanate from the rise of event risk in the era of dissonance.

- Volatility may be a poor proxy for risk.
- High levels of uncertainty can co-exist with low levels of observed volatility.
- The rising probability of abrupt price movements dramatically complicates volatility analysis.
- Liquidity matters more than ever in an event risk episode.

I. Volatility May be a Poor Proxy for Risk

One can easily confuse uncertainty with volatility, but they are not the same. As used in finance, volatility is a statistical concept often measured by the standard deviation of returns of a security or product over a given period. It is also about the degree of up and down movements in prices that has been observed historically or is expected to occur in the future. While often used as a proxy for risk, our perspective is that volatility is only one dimension of financial risk and not necessarily representative of the risks that are more dangerous or more feared. Uncertainty revolves around fears; that is, the probabilities of future developments and events about which one is worried. So, let's try to untangle some of the confusion.

Behavioral finance teaches us that market participants often have asymmetric risk preferences. That is, a person's fear of losing a certain amount of money is greater than the pleasure of gaining that same amount. The point here is that there is considerable evidence that risk preferences are non-linear. This means a symmetric metric, such as the standard deviation which weights equally the up-and-down price movements, is not appropriate when market participants have a much greater fear of the loss side of the ledger.

Once one is focused on the size of possible losses, other challenges with the standard deviation come into play. A Russian mathematician, Pafnuty Chebyshev, developed a theorem in the 1800s that essentially said that if one knew the standard deviation, then one would have a very good idea of the range of outcome likely to occur, say 75% of the time. This makes the standard deviation sound quite useful. The problem is that Chebyshev's theorem also implies that only knowing the standard deviation provides one with very little information about what could happen in extreme cases. That is, if your fears are focused on maximum possible losses, then do not rely on the standard deviation as your guide as it may badly underestimate potential losses at the extremes even though it provides a good guide for how outcomes two-thirds or three-quarters of the time.

II. Shifting Political Probabilities Versus Economic Fundamentals

Now let's turn our attention to the ability of relatively high levels of uncertainty to co-exist with relatively low volatility and how the two relate to event risk. Our perspective is that the fact that market participants are worried and highly uncertain about a future event or outcome does not necessarily create price volatility, and the high level of uncertainty may even dampen volatility prior to the actual event.

The challenge is that once market participants are focused on an event or policy debate that has the potential to go in starkly different directions, there is a shift in analysis toward the relative probability of either outcome occurring. That is, news that changes the probabilities of the outcomes will move the price immediately, but the price movement is not likely to be nearly as much as when the outcome becomes known.

Let's take an example from the world of mergers and acquisitions. Suppose company A has made a bid for company B at a price of $120, representing a 20% premium over the previous market price of $100; that is, the bid represents a substantial premium to gain full control of the target company. Now, suppose that the regulatory authorities decide to review the potential merger and that

there is a reasonable probability that the merger might be rejected. Let's say that after the merger is announced but before the regulatory authorities make their decision, there is a 50/50 probability the merger will be approved.

Effectively, there is one set of expectations for the stock price of company B if the merger is approved and that is $120, with a very small variance, since an adjustment to the bid in this case is not assumed to be very likely. And if the merger is rejected by the authorities, then the stock price of company B would likely fall back to its pre-merger announcement price of $100, give or take, and there would be considerable uncertainty as to the future of company B (high variance) since a new suitor might appear or the company might be perceived as damaged goods and the stock price spirals downward. The expected value (mean) of the combined distributions depends on the probability of scenario A (merger approved) or scenario B (merger rejected). At 50/50 probability, the expected mean for the combined probability distribution is 110.

Now, some new information arrives. Maybe it is a rumor that the regulatory authorities are leaning against merger approval. If the odds against an approval of the merger go down, the probability-weighted expected mean of the combined distribution will also decline. Note that the expected mean shifts in these cases of two far-apart scenarios when news alters subjective judgments of the relative probabilities. What is nearly certain is that the market price will not stay where it is once the regulatory authorities make their decision, and the market clearing price will move immediately to a 100% probability of the winning outcome. The status quo is decidedly unstable in this sense, but markets are not necessarily displaying much day-to-day price volatility until the event occurs.

From a practical perspective the argument here is that when event risk is present, then prior-to-the-event outcome prices move mostly on news that moves the relative probabilities of one outcome versus the other. This means that the typical economic fundamentals are not very much in play, while a different type of calculus is driving the relative probabilities.

We saw this during the Brexit referendum in June 2016. Prior to the vote, the British pound was trading at around USD 1.42 per GBP. Daily movements were relatively muted as opinion polls and analysts' conjectures on the upcoming vote shifted one way or the other. When the vote was announced and it was clear that the "Leave" camp had won the day, the British pound dropped sharply to around USD 1.30 per GBP and then even lower. It was similar for U.S. equities in the Presidential election of November 2016. Prior to the election date, many analysts argued that a Republican victory would lead to a large corporate tax cut, and support equity prices. Equity prices zig-zagged ahead of election day in a relatively tight range, but as results arrived on election night and it became clear that the Republican Party had swept the Presidency, Senate and House of Representatives, then the corporate tax cut probability rose sharply and U.S. equities surged. In both cases, observed volatility was subdued before the event even with very large uncertainties present, and prices surged once the outcome was known.

III. Appreciating Price-Gap Risk

There is another important observation related to what happens with market activity around event risk. There is a tendency to speak of volatility in a general sense, but with event risk it becomes

quite important to separate a one-time, abrupt price movements from a shift in the overall volatility regime. In financial markets, the term volatility, as we have discussed earlier, is focused on the standard deviation of the returns of a security or product. While a one-time, abrupt price change will impact the calculation of the measured standard deviation, it does not necessarily signal a shift to a new and higher volatility regime. Instead, our analysis suggests that the one-time, abrupt price change is simply the reflection of the market moving immediately to a 100% probability of the now known winning outcome, and that the new volatility regime may well be lower, the same, or higher than the old, pre-event regime.

If one is relying on implied volatility readings from an options pricing model based on the straightforward Black-Scholes-Merton models of the 1970s, then event risk brings a problem for the interpretation of implied volatility. The original Black-Scholes-Merton models assume that abrupt price changes cannot occur. In their models, price moves continuously in tiny increments up or down, and there are no price discontinuities in their technical language. In the real world, of course, abrupt price changes definitely occur and they are even more likely to occur when event risk is present. Option prices may include a premium for potential for a one-time, abrupt price movements in the presence of event risk. This means that the implied volatility assessment from options prices, assuming no price gap possibility, will typically over-estimate volatility for the post-event period. For those operating in the world of practical risk management, this distinction between price-gap risk and a shift in the volatility regime means that an options hedging strategy based on delta-hedging will not work properly if the price gap occurs.

As an example, the two probability distributions in Figures 1 and 2 have the same mean and standard deviation, but they indicate very different risk management challenges. Figure 1 is a normal, single-mode, bell-shaped distribution, while Figure 2 is bi-modal distribution which was created by mixing two normal distributions with highly divergent means and different standard deviations for each scenario. We want to highlight two risk management challenges of note when event risk is present to the degree that it would be reflected in a bi-modal risk-return probability distribution.

Figure 1: Risk-Return Expected Probability Distribution: Normal Distribution

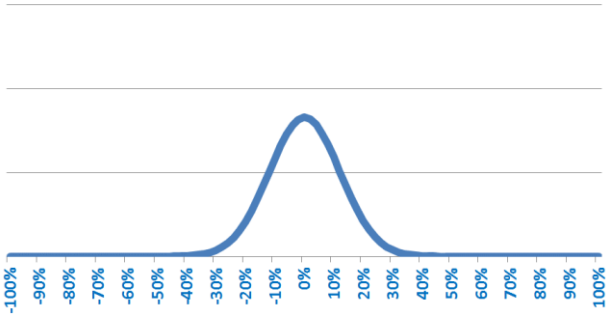

Mean of Distribution = 0%, Standard Deviation = 12%
Source: Chart Created by CME Group Economics.

Figure 2: Risk-Return Expected Probability Distribution from a Two Scenario Mixture

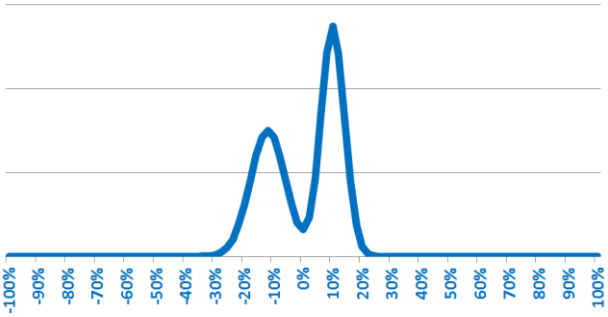

Mean of Distribution = 0%, Standard Deviation = 12%
Source: Chart Created by CME Group Economics.

To summarize this discussion, let's examine the type of market expectations captured by a typical bell-shaped curve. With a bell-shaped risk-return probability distribution, one is essentially saying that there is a consensus view around the expected return with incrementally different possibilities of

lower probability creating the bell-shaped risk distribution. While the market may shift in increments to a new consensus based on new information, the expected mean shifts are continuous and do not constitute price gap risk, and the expected volatility remains more or less the same. In the case of a bi-modal risk-return probability distribution, there is the possibility of an abrupt shift in the expected mean (price gap risk) and a shift to a new volatility regime.

To make matters a little more complicated, we have also observed some evolution in the nature of event risk. The examples given earlier from Brexit and the U.S. Presidential elections came with very specific dates on which the outcome would become known. Not all event risk works this way, and probably most event risk does not work this way. Event risk often comes with a policy debate but not necessarily a specific date. The U.S.-China trade war would be an example of event risk without a specific date. Market prices have typically reacted quite abruptly to the imposition of tariffs or retaliation for tariffs. But the tariff announcements or the retaliation have sometimes come as surprises to the market. And, while deadlines for deals have been set, they are not necessarily binding, as delays can be announced.

IV. Liquidity in the Era of Dissonance

Event risk and the possibility of abrupt price changes put liquidity in the headlights for risk managers. Whether the event risk comes with a known date (e.g., elections) or a general time frame with unknown dates (e.g., U.S.-China trade war), risk managers are going to focus their strategies on trade execution in the most liquid markets which are correlated to their risk exposures. And, for sure, event risk does not necessarily respect typical trading hours.

What we have observed with liquidity during an event risk episode is that aggressive buyers and sellers, as well as liquidity providers are all drawn to the liquid marketplace. Take equites, for instance, during the night of the U.S. Presidential election in November 2016. The likely outcome became known in the middle of the night – U.S. time. If one owned a portfolio of single stocks, there was not necessarily much liquidity in any given stock and an aggressing seller or buyer would likely pay a very high bid-ask spread to get a transaction executed immediately. By contrast, stock index futures are operated on a nearly 24-hour electronic trading platform. CME S&P500® E-Mini index futures provided excellent liquidity with very narrow bid-ask spreads even as the market went from 5% down to sky-high. We have seen the same thing happen in other markets. During normal trading environments, a risk manager may have choices as to what venue and product to choose, but in an event risk episode, only the more liquid markets can offer narrow bid-ask spreads, and they attract the bulk of the trading activity.

CME Group's Data Science team has examined the liquidity of various futures markets. To view current and historical CME Group product bid-ask spreads, and other measures of liquidity, access the free CME Liquidity Tool available at: cmegroup.com/liquiditytool.

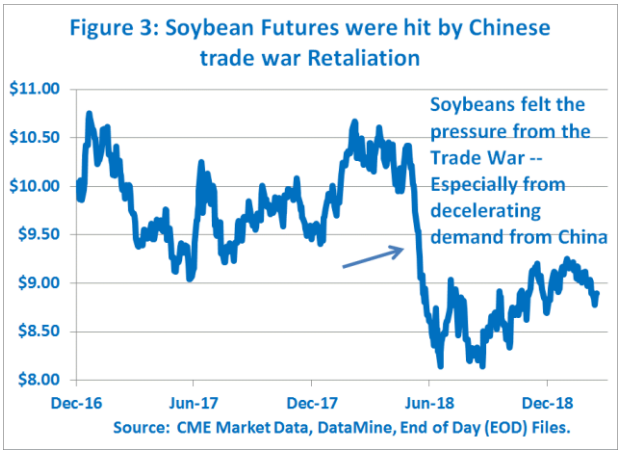

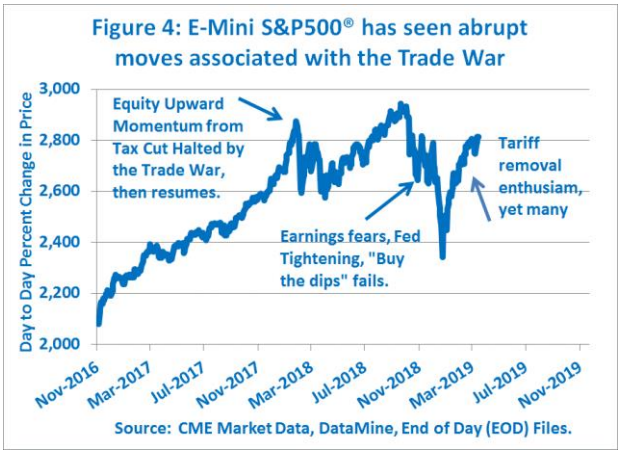

V. Bottom Line

To summarize, markets that are subject to event risk, which appears much more commonplace in this era of dissonance, require special risk-management attention.

- Volatility may be a poor proxy for risk when event risk is present.
- High levels of uncertainty can co-exist with low levels of observed volatility during the pre-event period.
- The rising probability of abrupt price movements dramatically complicates volatility analysis, may make standard interpretations of implied volatility inappropriate, and may impair the efficacy of common options hedging strategies, such as delta hedging.
- Liquidity matters more than ever in an event risk episode, and risk managers will naturally shift trading to the most liquid markets that are correlated with their underlying exposures.