Risky Business II

The Coming Volatility

WoodedPark Strategies, LLC

Adam Hoffman & Michael Patterson
October 24, 2018

Volatility is coming: in the stock market, in commodity markets, in interest rates, and weather. Sooner or later, volatility is coming. That is all but guaranteed. But when will it arrive?

We'd all like to know when, but not one of us does with precision. Financial firms attempt to answer the *when* question by forecasting volatility, and markets themselves offer forecasts. Market participants can trade volatility-based futures and options contracts, and those contracts offer yet another view of future volatility. Given the importance of volatility in financial and commodity markets it's natural that investors of all types are interested in volatility and its forecast.

In this article, WoodedPark's second in a series on risk, we'll focus on the topic of volatility. We'll examine some of its properties and some of its shortfalls when used to measure risk. We'll discuss several methodologies for measuring volatility. Finally, we'll discuss some techniques for forecasting volatility. We will keep the math to a minimum.

Volatility

In our last article, *The Energy Industry's Relationship with Risk*, we noted that risk is often measured as volatility. We referenced Merriam-Webster to define volatility as *a tendency to change quickly and unpredictably.* In that same article we provided the mathematical formula for volatility σ of a price series p:

$$\sigma = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \overline{x})^2}$$

In the above equation, χ_i denotes the logarithmic return of a security's price p_i. The subscript i denotes different measurements, for example, the price p on different days. The continuous return χ_i is calculated as $\ln(p_i / p_{i-1})$. $\bar{x}$ is the mean of the return, and *N* denotes the number of returns to be measured. This formula can be used to calculate the volatility of stock prices, commodity prices, temperatures, corporate earnings, or any set of time series data. By convention volatility is listed as an annualized number, meaning that if χ denotes daily returns, σ is to be multiplied by the square root of 365, or the square root of ~250 if the measured prices are recorded only on business days.

The formula to calculate volatility is straightforward and is available on many software programs. However, the formula requires specification of several parameters, and that specification can be subjective. The parameters include the following:

1. The frequency over which to compute returns. The frequency could be monthly, daily, hourly, or every minute.

2. The window *N* over which returns are squared and summed. For daily returns *N* is often set to reflect one month of business days (*N*=21) or three months (*N*=63).

3. The number of periods that comprise a year (the 365 or 250 number of days mentioned above to annualize volatility). For financial instruments traded on weekdays, 252 periods is commonly used. However, with global

marketplaces becoming more integrated and with the advent of electronic trading, this number may eventually standardize to 365.

The above formula for volatility employs log returns, and people often ask, 'Why not just use the price series itself?' The answer is that some prices, e.g., stock prices, trade for pennies while others trade for hundreds of dollars. Rather than have volatility misleadingly scale with the price, we normalize the price with returns. As an alternative, one can compute volatility with a percent change in prices.

Because of the differing parameters, any given data set can produce different volatility measurements. It is important to compare volatility measurements carefully.

We're finished with the math for this article, and from here on it gets easier.

Volatility Characteristics

Volatility has several features that warrant mention. First is that volatility is often equated with risk; this view is incorrect. We define risk as *the measurable probability that in the pursuit of an objective, an unexpected event causes loss of revenue, reputation or viability.* The volatility calculation provides only a measure of risk, one of many such measurements and imperfect in numerous ways.

Second is that while risk refers to a loss, volatility measures both loss and gain symmetrically[1]. According to the volatility calculation, you are as likely to gain as lose.

Thirdly, volatility provides to us an expected loss or gain, but it fails to address the *when* question. Annualized volatility tells us what we might expect over the next year, but the losses or gains could happen tomorrow, a year from now, or not at all.

Given all the subjectivity and imperfections, it's a wonder we use volatility at all. But volatility is one of the better understood risk measurements across many industries, and it is widely employed.

Volatility Example

Let's look at an example. Rather than focus on an energy contract, let's look at a specific financial contract: the widely traded S&P 500 Index (ticker: SPX). We'll look at this contract because (a) we have a long history of its data, and (b) there are multiple ways of measuring its volatility.

The chart below shows the S&P 500 Index, from 1990 to present. The blue line is the Index, and the orange line its 30-day moving volatility. Volatility was computed using our formula above, with N=30, and annualized with 252 periods. Note that this volatility calculation is based upon the S&P 500's closing (end of day) price and ignores intra-day price movements.

---

[1] One of the more interesting features of volatility is how it affects the *price* of things, such as a stock or commodity. Those prices change in a random manner but will never become negative. In mathematical terms, when price changes are normally distributed the price itself becomes lognormal. In non-mathematical terms, returns are not symmetric; potential gains are unlimited, but losses stop when a price reaches zero. Nassim Taleb has an eloquent discussion of this volatility characteristic in his book, Dynamic Hedging.

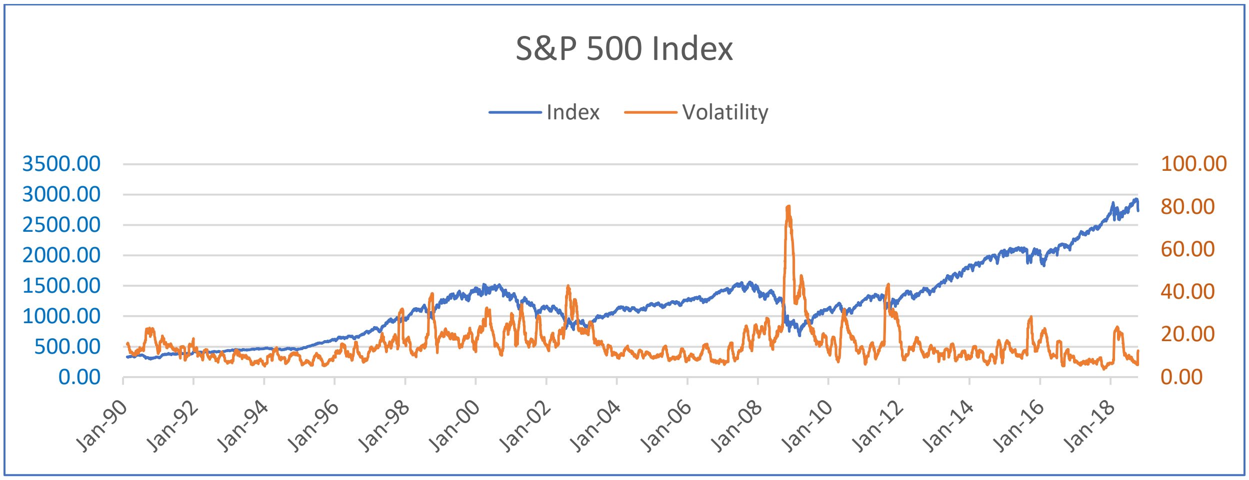

Figure 1. S&P 500 Index and its Volatility

The above chart illustrates a few features of volatility. First, with only an occasional exception volatility of this index hovers around 10-20%. When volatility does spike it normally does so as the result of a price movement down, not up. Restated, when volatility is large the market is probably losing, not gaining. There's a humorous saying in finance world that *the market takes the stairs up and the elevator down*. Studies have shown this to be true: volatility is reduced in rising markets and increased in falling markets. As an example, observe the (generally) low volatility in the 2010-2018 time-frame in the chart above as the S&P 500 posted steady gains.

As we have mentioned several times now, the above chart has one significant limitation: the volatility measurement is past-tense. That is, you cannot measure volatility until you've lived it. There are ways, however, to estimate future volatility.

Volatility Term Structure

Let's cover one last feature of volatility: its term structure. As discussed earlier, the volatility calculation involves several parameters, and their specification is subjective. One of those parameters, the window *N*, has a large effect on the resulting volatility values, as presented in the following chart:

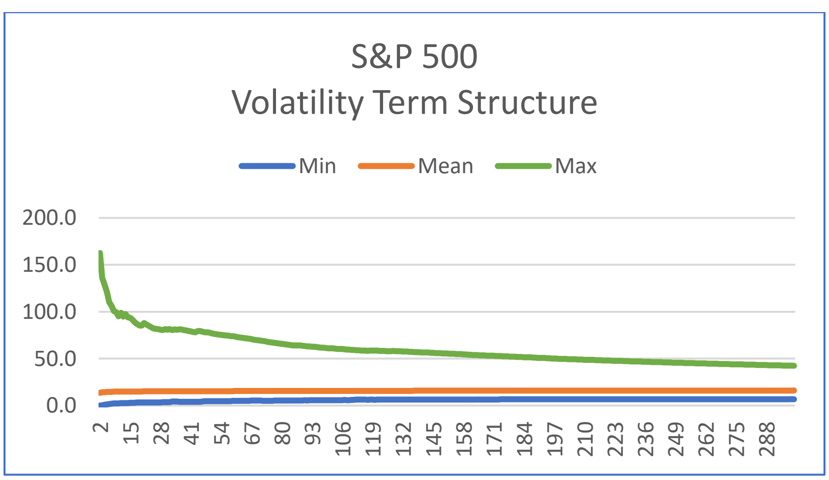

Figure 2. The Term Structure of Volatility

The above chart shows the minimum, maximum and mean volatility of the S&P 500 (the exact same data set used before) for $1 < N < 300$. For small values of N, the maximum volatility surpasses 150%. As N increases, the maximum volatility tends toward the mean, as does the minimum volatility. Most financial data exhibit these trends.

What the term structure tells us is that volatility has a feature called *mean-reversion*. When volatility is low or high, given time it will revert to its mean.

The VIX Contract

The Chicago Board Options Exchange (CBOE) published the financial industry's first volatility index in 1990. Known as the VIX, this index provides an estimate of the S&P 500 index's volatility for the coming thirty days. By convention, the VIX is reported in annualized terms. The VIX is a computed index, and it is important to know that the index is computed from *options* on the SPX contract, not from the SPX itself. As such, the VIX is an indirectly computed estimate of S&P 500 volatility.

Option premiums increase with perceived risk, and as premiums rise so will the VIX. Some investors thus refer to the VIX as the *fear index*. Let's look at the VIX for the same time-frame as the above chart:

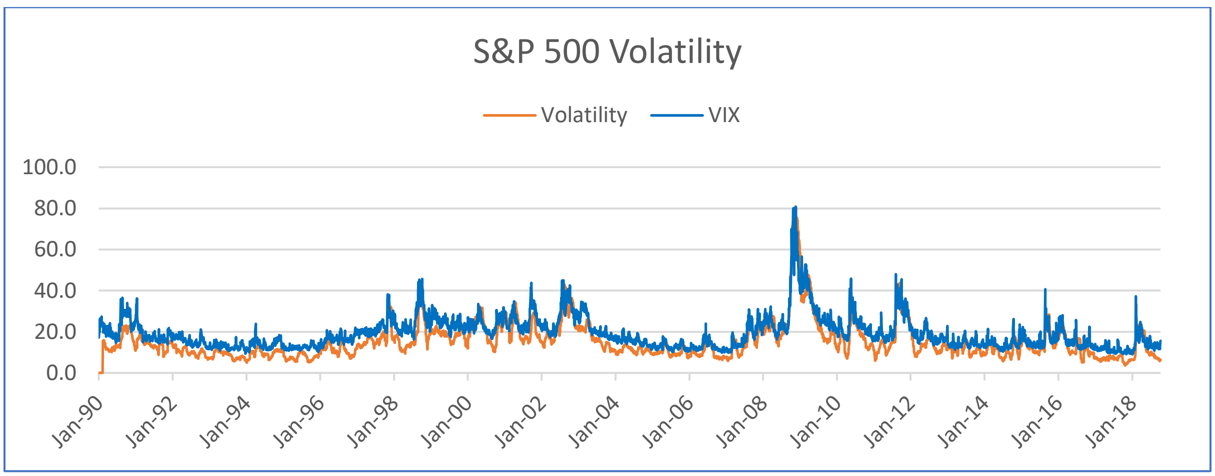

Figure 3. Computed Volatility and the VIX Index

In the above chart, computed volatility is again shown in orange and the VIX is shown in blue. It's clear that the two data sets are highly correlated, but the VIX runs higher than the computed, or historical volatility. Why is that? It's because historical volatility (orange line) is a past-tense measure of volatility whereas the VIX (blue line) is an estimate of future volatility. The option market makers who buy and sell contracts on the VIX require a premium to cover risks they cannot foresee, thus the VIX runs higher than volatility itself.

Implied Volatility

The VIX provides a measure of the volatility of a single index: the S&P 500 index. But there are countless other contracts which can be bought or sold, and they don't have a published index like the VIX. Many financial and commodity contracts, however, also have corresponding options contracts. The price or value of every option contract is a function of several variables, including the volatility of the underlying instrument. Once market participants trade an option contract, the price of the contract is known and it is then possible to compute a corresponding volatility called *implied* *volatility*. This measure of volatility is another estimate of future volatility for the underlying financial or commodity contract.

Forecasting Volatility

Mathematical models for forecasting volatility first appeared in the late 1970s. These models attempt to replicate several key features of volatility: a tendency to cluster, a tendency to be autocorrelated, meaning that future values depend upon past values, and mean-reversion.

The model names can be imposing: exponential weighted moving average; autoregressive; and generalized autoregressive conditional heteroskedastic. Beginning in the mid-1990s, artificial neural network models appeared and have shown some success in forecasting volatility. Generally, models specify a mathematical formula for volatility which contains a set of unknown parameters. In completing the model one attempts to determine parameter values which minimize the errors that a model produces for an historical time-frame. That model is then applied to the future.

We'll skip the math for the sake of brevity. But if you are interested in learning more, a very readable paper titled *Forecasting Volatility in the Stock Market* reviews the underlying math of several models, and then compares each model's ability to forecast S&P 500 volatility. While each model presented in the paper shows some success in tracking measured volatility, each exhibits a time lag. That is, the models can track a jump in volatility, but only after it has begun. As such, the models don't answer the *when* question.

The Importance of Volatility

You might be asking yourself, 'Given all the imperfections of volatility measurements and forecasting, of what use is volatility?'

If you invest in stocks, bonds, commodities, real estate, or any other asset you may already have an answer to that question. Before you invest in any asset it is wise to view its historical volatility to help determine if you want to be exposed to that asset's risk. Even if you invest only in a widely diversified index such as the S&P 500, you'll want to watch its volatility: if the SPX rises from its usual 10-20% range of volatility, you'll want to watch more closely.

If you happen to invest in derivative contracts (futures, options and swaps), then volatility is an essential part of your life. Many options traders follow volatility more closely than they do prices. In fact, volatility is a key driver of option prices.

And if you want to invest in volatility itself, you can. You can purchase futures and options on a contract such as the VIX.

Government entities that regulate financial firms often use volatility-based measurements to monitor a firm's activities. This topic will be discussed further in our next article. Lastly, firms use volatility in both market portfolio theory and in risk management.

For all their deficiencies, volatility measurements and forecasts are ever-present tools used in the finance and commodity industries.

Conclusion

We never answered the *when* question. Volatility is coming, we just don't know when.

A more accurate statement would be that volatility is ever present. Stocks, bonds, interest rates and commodity prices change constantly, and thus are constantly volatile. There's a volatility measurement for every day, if not every hour or minute. We tend not to concern ourselves with these price changes, until a big one arrives. And the big changes are usually downward ones. If only we knew when those changes were coming, we'd get out of the way.

Warren Buffett said, 'Predicting rain doesn't count. Building arks does.' Preparing is better than predicting. And that is what risk management is about. Volatility is coming, we just don't know when or how severe it will be.

In our next article we'll talk about a more highly refined measurement of risk. We'll talk about VaR (Value at Risk) and its significant limitations, and after that we'll expand the discussion to more advanced measurement techniques. We thank you for your time in reading this article, and we welcome your feedback.

Adam Hoffman, CEO of WoodedPark Strategies, is a commodity trader and risk manager with over 20 years of experience managing organizational risks related to capital, personnel and technology. Adam can be reached at adam.hoffman@woodedparkstrategies.com.

Mike Patterson, Lead Technologist at WoodedPark Strategies, is an analyst and technologist in the energy industry. Mike can be reached at michael.patterson@woodedparkstrategies.com.

WoodedPark Strategies, LLC is a privately held risk management advisory firm located in Houston, TX.