Deep Dive: Peak Volatility – How Hot Can It Get?

(Thorsten Wegener, 18 March 2020)

Chart 1: Phantom VIX 1987

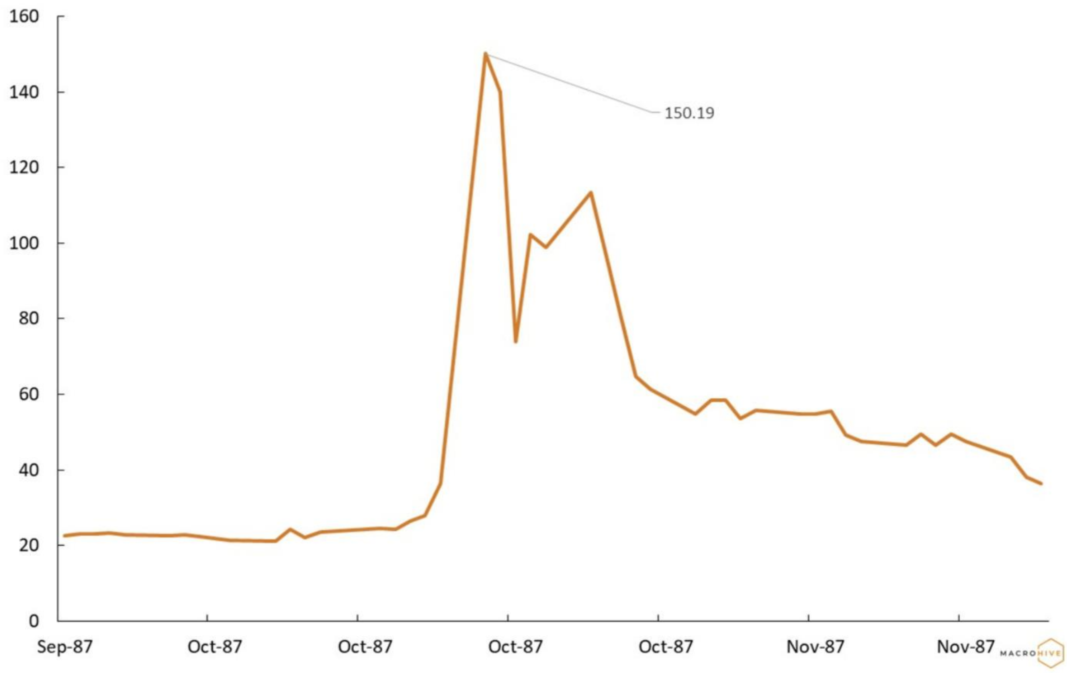

Source: Yahoo Finance, Macro Hive

Is there a concept to approach 'Peak Volatility'? And if not, why not? No, I am not mixing my words up over recent, let's politely call it *erratic*, market behaviour. To get an idea of what kind of truck could hit you might just help you to survive. Selling volatility to pocket the carry has been a very lucrative trade for many players for many years until it wasn't (isn't) anymore. In subdued, 'normal' markets, annual returns between 30% and 50% are by no means uncommon and can be achieved by shorting the front month futures, rolling them into the next expiration when the time comes – and all that while sipping tea until the time arrives for ritual self-disembowelment. Were it not for the occasional (but inevitable) unpleasant surprises, this strategy could keep a hungry hedge fund manager fed quite nicely. So, the question I have been encountering, unsurprisingly, over recent days goes like this: how high can volatility go?

Zero Volatility

As so often, physics provides an analogy. It's a well-known question: what is the maximum temperature a system can approach? Almost all students who have stumbled through a *Physics 101* textbook can answer the question's opposite, what is the lowest temperature any system can approach? It's 0° Kelvin or -273.15° Celsius – affectionately known as 'absolute zero'. Ridiculously simplified, this is the state where nearly all motion ceases to exist. Without motion there is no friction, and without friction, no heat.

Applying this analogy to security markets, the 'absolute zero' of a marketplace would be zero volatility. Theoretically, that would happen either because everybody agreed on the price of a security and trading only happened upon this universally agreed price without any deviation to the upside or downside. Or, more likely, it would be if enlightened central bankers all over the world stepped up and closed markets and stopped trading to avert pending doom. Of course, you could argue that in this case volatility is not really zero but merely in a state of suspended animation, waiting to resume its reign of terror once markets open up again. But for the sake of this argument, let's say we have defined the lower boundary of volatility with relative ease.

However, no aspiring volatility short seller is too much concerned about that question. If you want to sell volatility at *perceived* peak levels the real question is, how high can it go if the underlying rational about the market is wrong? Having been in these shoes, I can assure you: no one sells the top. That is usually the place you buy back and close your shorts out when the pain gets too much.

So, the next question deriving from this unpleasant realisation presents itself as a money management problem. Taking individual capital and margin requirements into account, peak volatility will determine up to which level I can 'average in' my shorts and stay solvent to benefit from an eventual decrease in volatility. Nobody can stay crazy forever, not even the market. Eventually it will calm, and those unfortunate souls who think 'this time it's different' win a Darwin Award. In other words, we have gone full circle and returned to the time-tested question: can you stay solvent for longer than markets are defying common sense?

Infinite Volatility

Back to physics and our initial question about the highest temperature that a system can approach or the maximum volatility a security market can spit in your face. Enter Albert Einstein and his equations that proved the existence of black holes.

If you keep adding energy to a system, its mass will increase because the speed of light is constant. Say what? $E=mc^2$. Energy (E) and mass (m) are interchangeable. Mass can be converted into energy and vice versa, so if you keep adding energy to a system its mass will increase and eventually start to contract under its own gravity the 'heavier' it gets – until, that is, it goes 'poof' and a black hole appears. I.e. the market stops working.

As the laws of physics break down in a black hole, the last moment at which we can calculate maximum mass of this object is just before it disappears into a so-called singularity. As our object contracts under its own gravity it will get hotter and hotter, but we are still able to calculate the maximum heat at the point before all equations break down in a pool of darkness.

This is called the 'Planck Temperature' and it is roughly 141,000 billion billion billion° Celsius. Of course, we could keep on adding energy to a system and there might well be higher, even infinite temperatures possible, but we would not be able to observe what happens as it is beyond the event horizon of the black hole. Can we use this approach in a thought experiment to calculate an upper limit for volatility, just before it disappears beyond a horizon of inaccessibility? Let's try.

Volatility – The Future's Past

After leading you down the exciting path of Planck constants, black holes and event horizons to make up a highly inappropriate analogy, we must unfortunately step back for a moment to define volatility before we can think about a potential event horizon for market volatility. My favourite definition for volatility is a simple one: volatility is the measurement of how much stuff moves around. It's not very scientific, but it is very appropriate. Using some statistics, you can ascertain three kinds of volatility – historic, implied, and future realized volatility.

Historic volatility is the easiest to tackle. Data is widely available, and you can easily compute by how much security returns moved on average around the mean, average growth path of that specific security. No ambiguity here. And, it often is a good guide for the future.

Implied volatility, on the other hand, is a real puzzler. It is the collective guess of market participants about how a specific security will move around its mean. To quantify this picture the marketplace delivers option prices, which in a free market are being driven by the risk preferences and individual divinations about the future expressed in those bids and offers. Plugging these prices into the famous Black-Scholes equation allows you to compute the implied volatility. The VIX Index, for example, is calculated by using the prices traded for 'at the money' SPX options with a maturity of one month and, after having plugged in all parameters, solving for the resulting implied volatility. This output is normalized to adjust for different option maturities.

A VIX Index of, let's say, 40% conveys the information that market participants, having voted with their own funds through buying and selling options, expect the SPX to not exceed a positive or negative performance of +/- 40% in a year. This represents a one standard deviation move, in other words a probability of roughly 68% to stay within these +/- 40% boundaries after one year. Two standard deviations represent a probability of 95% to stay within a +/- 80% corridor, and three standard deviations represent a 99% probability not to trade above or below 120% of its current level. As prices can't go negative (that is what we once thought about bond yields…) the resulting price distribution is skewed to the upside, or in geek speak it is log-normally distributed.

The pot of gold underneath the rainbow is, of course, the future realized volatility of an asset. Turned around, this can also be seen as a calculation of the historic volatility once the future has run its path. It's the number that tells you what will have happened, and if you are an options trader and get this estimate right, you will make money without having to be bothered about market direction. The equivalent hedge portfolio will compensate the trader for the trader's wins/losses on the other side of the equation. As long as the trader sells a bit above this future realized volatility or buys below this level the P&L will be positive. However, even if a trader has got these levels right, how much can they get beaten up by the markets in the meantime? Is there a maximum volatility we can expect in a highly frantic market, just before we disappear behind the event horizon?

Is There an Event Horizon for Market Volatility?

How high can volatility rise? This should not be countered with the question, how long is a piece of string? In fact, we can approach this tricky question with some parlour tricks.

Remember, realized volatility is measuring the moves that have happened and will have happened. Naturally there can be only one, the security in question will have moved as it did. Implied volatility, on the other hand, represents the market's best guess about this future path our security will have taken, and there are countless different guesses for different strikes and maturities. Ever wondered why options on one underlying can have different implied volatilities at different strikes, when there is only one possible path for the future? So did Merton, Black & Scholes, but skew and term structure are a topic for a different article…

I am procrastinating. Let's start with a look into the black hole of infinite volatility. I will, in my deliberation, focus on the left side of the tail, as maximum volatility tends to coincide with falling markets (if you don't happen to be short Tesla).

Infinite Volatility

Can implied volatility on the SPX, represented in this case through the VIX, reach infinite levels? Sure, as long as you can find a moron who is willing (probably in a fear-induced waking coma) to pay any price for an option, then implied volatility can be whatever our moron trader in this example wants it to be. But like the possible temperatures in a black hole, those are beyond the realm of the observable universe and, if observed, picked apart at the speed of light.

An easy example to visualize this: as of writing, the SPX is trading at roughly 2,500 points (it could be 200 more or less in 15 minutes, who knows). How much would a (half-)sane trader be willing to pay for a bet producing the smallest possible profit of 1 cent? Surely not more than 2,499.99 bucks, otherwise the trader would lose money even in the case of a certain doomsday scenario of 'SPX Zero'. If they would pay more, even a black hole would fail as an analogy for the incomprehensible. If they want to put this hedge into place, pay $2,499.99, and we assume zero interest rates, the implied volatility for a 30-day hedge tops out at 3203.3%. Not a huge consolation if you consider shorting volatility at levels around 70% and you are left with the question, how high could it possibly rise until I would no longer be able to average up my shorts, once we reached 140% or 280% or 560% in the VIX. This would be equivalent to our event horizon. Let's move a little bit closer to reality.

Batshit Crazy Vol. (it's a term)

Volatility is defined as the one standard deviation covering the probability distribution for 68% of its width. Ever frightful of the worst scenario, it would be prudent for our trader to hedge their bets by going all in. Let's hedge all the way up to three standard deviations, covering 99% of the probability distribution. By pretending the market could drop to zero within one year within one standard deviation instead of the usual three or more, we arrive at an implied volatility of 289%. Now we have left the event horizon of very unlikely (yet not zero probability) events and moved further out into the orbit of remote possibilities.

Just Plain Crazy Volatility

The next step would be to fall back to our standard model, operating with one standard deviation covering 68% of the distribution, resulting in three standard deviations covering the whole length of the SPX going all the way down to zero, but not negative, which should be considered a sensible assumption. After all, we still rely on a bell-shaped distribution of returns, never mind the level of volatility. On a side note, once we abandon this assumption you might as well go for a nice cold Corona, because my whole beautiful thought experiment goes out of the window.

Anyway, returning to our 'reasonable' assumption that three standard deviations cover the SPX down to roughly zero levels, we arrive after a back-of-the-envelope calculation at an implied volatility of 97%. This translates into 171 days with a daily expected volatility between +/- 5.05%, 75 days of daily variations of up to, but not larger than +/- 10.11% and around 2 or 3 days of mayhem with moves of up to +/- 15.16% out of 252 trading days.

Maximum Volatility

Ignoring the insane case of an implied volatility of 3,203.3% (being specific – offering the .3% adds a nice touch) at the volatility event horizon, we should reasonably assume that maximum volatility could be localized somewhere in between batshit-crazy levels of 289% and just-plain-crazy levels of 97%, probably closer to 97% than to 289%. Markets will be shut down long before we approach 200%.

What is left is to compare the most appropriate real-life scenario when we were very much in the proximity of maximum volatility levels while staying with our admittedly very deductive thought experiment.

'87 Baby

A reconstructed Phantom VIX[1] spiked at 172.79% intraday on 20 October 1987, one day after an unprecedented one-day drop of 22.90% in the SPX – in itself a 22 standard deviation event. However, the Phantom VIX level I used for my calculations was 140% by closing, roughly 10 points lower than on the day of the crash itself (150.19%).

Translating this prediction, an implied volatility of 150.19% generates expected daily moves over the next 252 trading of:

171 days with daily volatility between +/- 7.86%, 75 days of daily variations of up to, but not larger than, +/- 15.72%, and between 2 and 3 days of epic moves larger than the initial crash, with trading ranges of up to +/- 23.58% out of 252 trading days.

What happened? Instead of having 2 to 3 days with moves between +/- 15.72% and +/- 23.58% there was none. The expected 75 days with moves between +/- 7.86% and +/- 15.72% turned out to be 2, and over the remaining 250 days the market traded between +/- 7.86%.

In fact, the realized volatility, the path of returns the SPX actually followed after the '87 Crash for the next 21 trading days, turned out be 56.49%, and for the whole 252 trading days following peak volatility, it came in at a miserly 25.11%, just 5.3% points above its long-term average.

We haven't seen these levels since '87. Even the dicey time around 2008/09 kept peak volatility at around the 80% level – far below the madness of '87. A trader selling panic level volatility through SPX options while staying delta neutral at closing levels to hedge their exposure would have earned a nice nest egg for a rainy day, and those who stayed calmer than the rest of the crowd would have saved a fortune by not closing out their positions at this point in time. Theoretically, volatility can still increase substantially but not beyond the point of logic.

The Bottom Line

My heart goes out to the savvy trader who is prepared to sell volatility when everybody wants to own it, or not to buy it back when the cold breath of risk management is touching your neck, giving you goose bumps. It doesn't help either having board members suddenly discovering their fondness for the trading floor in a crisis, giving good advice while staring over your shoulder. But what helps to guide one's approach to allocate risk capital is an understanding of what this illusive figure of, for example, 150% volatility really means and what the future might hold in a worst-case scenario. If you are still in the game after volatility spiked to 150% you might just survive volatility levels of 300% for a day or two, knowing what can be reasonably expected this side of the volatility event horizon.

[1] At the time of the methodology switch from OEX options to SPX options, the CBOE created the VXO to provide continuity with the historical method of calculating the "old VIX" prior to 9/22/03. Let's stick with Phantom VIX – that sounds cooler.