2008: Volatility in action

Kambiz Kazemi, CFA, Polar Securities

Financial markets have experienced unforeseen levels of uncertainty since the start of the current credit crisis, and more specifically since August 2008. Levels of realized (i.e. past) and implied volatility (i.e. market expectation of future volatility) on many securities across asset classes reached all-time highs between September and December 2008. For instance, in the case of equities, we have to go as far back as the 1929–1933 period to find the S&P500 index exhibiting similar levels of volatility (Chart 1).

Historic data: S&P500
3 months realized volatility (1928-2008)

Chart 1

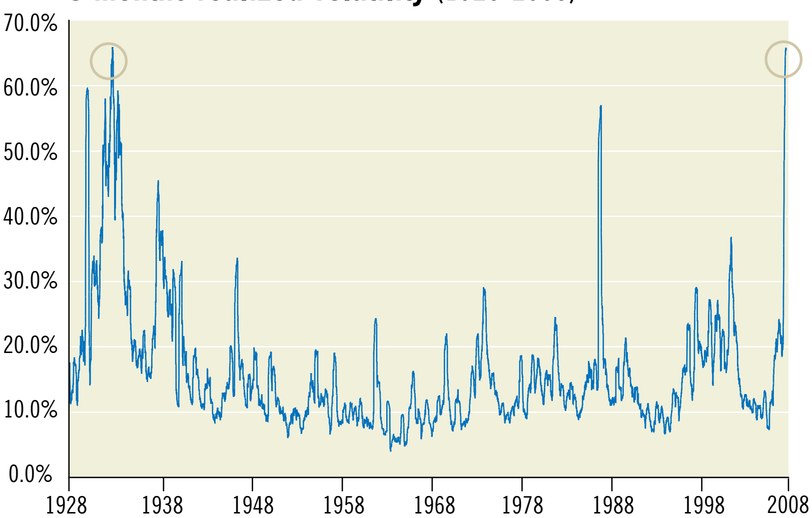

Most would agree that the massive deleveraging of the financial system at the global level accentuated the abrupt price movements and caused a fundamental re-adjustment of the perception of future risks and uncertainties. This in turn resulted in both high levels of realized volatility (due to wide price movements) and an increase of risk premia, of which implied volatility is a direct measure (i.e. expectations of future volatility is the "price of risk"). As an example, Chart 2 illustrates how the "price of risk" over a one-year horizon for the EuroStoxx50 index (Europe's most-traded equity index) tripled between the end of September and beginning of December 2008. The same measure for crude oil remains at its high at the time of writing.

1 yr implied volatility of EuroStoxx50
& Crude Oil (i.e. price of risk)

Chart 2

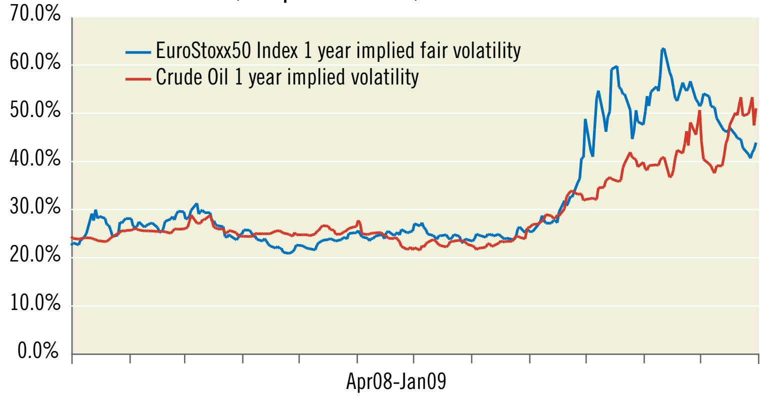

In the sense that 2008 was a real-time stress test environment for all of our theories and models, it provided us with an exceptional opportunity. The amplitude and speed at which both realized and implied volatility changed during 2008 took most by surprise. Virtually all of the models on which an arsenal of financial and risk management tools rely proved to be insufficient in foreseeing and accounting for the magnitude of the events we witnessed. Assumptions that were relied on to define the world of financial markets often proved to be invalid or incomplete. In hindsight, our industry's general overconfidence in our theory of financial phenomena, coupled with an over-reliance on a set of models often touted as the "true" way of viewing risk and derivatives, were probably in part a catalyst of the recent market turmoil.

In what follows, we will briefly review selected aspects relating to our understanding of volatility, risk and correlation in light of 2008's volatility regime and attempt to take away lessons for the future.

The future: Who knows?

Forecasting volatility has been an area of focus of academics and practitioners over the past several decades. In fact, given the close relationship between volatility and risk, if one had a reliable forecast of volatility, not only could it be used for trading purposes, but also for managing risk at large.

The natural first step is to test if past (realized) volatility is a good indicator of future volatility. The short answer is no.

continued... pg.17

'Investment Theme' continued from pg. 16

Many forecasting approaches (long-term memory, adding jumps, using ARCH and GARCH, etc.) were tried and did not exhibit any forecasting capability that could be used in practice.[1]

So how about implied (market expectation of) volatility? Is it a reliable measure of what actually lies ahead of us?

Authors studied the volatility of the S&P500, S&P100, T-Bills, foreign exchange, U.S. treasuries, etc., over various time horizons. What did they conclude? Market expectation of volatility often has not provided a reliable picture of the actual (realized) future volatility. In other words, markets have had a poor track record in assessing future risks.

As an example, the recent market turmoil between September and December 2008 was definitely not foreseen by market participants, even three months before it happened. Chart 3 shows that in September, the market expected the following three months to have 30 percent volatility whereas this period ended up having a volatility of over 60 percent.

Since December 2004, the market's three-month expectation of volatility for the S&P500 has proven to be off by more than 10 percent, almost 77 percent of the time. In periods of high volatility, market expectations seem to deviate the most from what follows in reality: up to 250 percent error in predicting the level of September to December 2008 realized volatility.

It appears that implied volatility (expectation of future volatility) is not a reliable measure of actual future volatility, but rather an expression of the perception of risk at the present time. That is why VIX (CBOE's volatility index) is often referred to as the "fear gauge." Chart 4 illustrates this point. The EuroStoxx50 index level is closely correlated with the three-month implied volatility (85 percent r-squared). In a sense, market participants tend to project their perceived risk of the present situation into the future.

Take away: We should avoid relying on the market's appraisal of risk as a forward-looking measure, especially when it comes to risk management.

"Normal"? Probably not

After August 2007, we often read statements of the type "this is a 1 in 50 or 1 in 100–year event." Sadly, such events have probably occurred two dozen times since. In recent years, many have been vocal and have written about the fact that we do not live in a Gaussian world (where a security's price changes are normally distributed). However, the first to point this out was probably Benoit Mandelbrot (sometimes referred to as "father of fractals") in a seminal paper published in 1969, where he studied the price changes of cotton in the United States back to the late 1700s. Many traders (especially those in foreign exchange and commodities) have a strong intuitive sense of this reality since they have witnessed it in practice. Much work has been done on incorporating this observation into the price of derivatives. A wide range of models accounting for jumps and discontinuities were developed and deployed. Yet, 2008 proved that market participants were overly relying on models and assumptions and did not question

continued... pg.18

EuroStoxx50 3 months realized volatility vs. implied (expected) volatility 3 months earlier

Chart 3

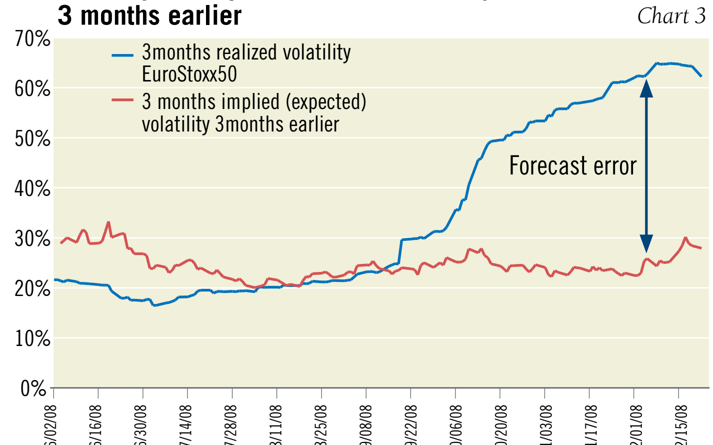

EuroStoxx50 level vs. EuroStoxx50 3 months fair volatility (2008)

Chart 4

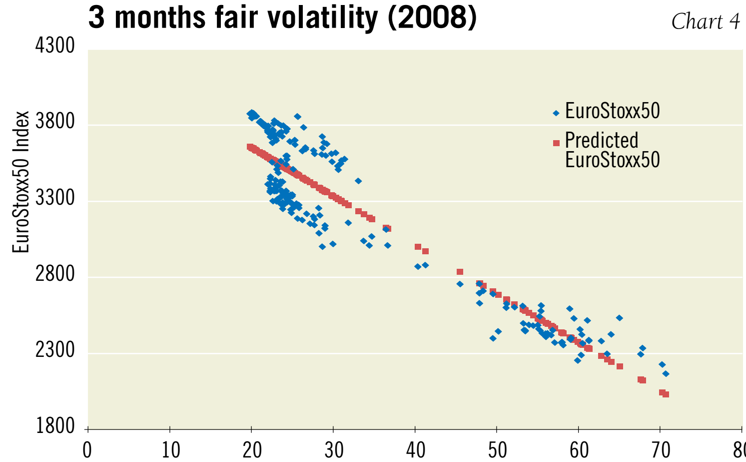

'Investment Theme' continued from pg. 17

them nearly enough. Chart 5 shows how 2008 was far from being a representation of a "normal" Gaussian world.

Reality vs. theory (S&P 2008) the "tail risk" *Chart 5*

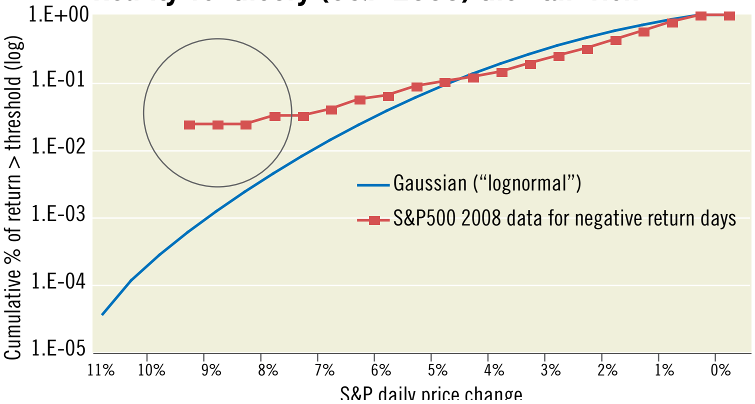

Equity derivatives markets, unlike currency or commodity markets, have been somewhat immune to this "tail risk," since they had seldom witnessed the recent extreme moves since the advent of volatility trading. There is no doubt that going forward, traders and risk managers will be much more attentive to "tail risk" than they have ever been in the past. Take away: Beware of assumptions and models. There is no absolute truth when it comes to translating the realities of the market into a working model.

Uncorrelated, you said?

Another cornerstone in derivatives pricing and hedging is correlation. Many structures sold to retail investors, such as principal guaranteed notes, as well as many so called "exotic" derivatives traded by dealers and professional investors, make assumptions on the nature of correlation between securities. It so happens that correlation and joint-probability distributions are also at the heart of pricing and evaluating the risk of mortgage-backed securities (MBS), collateralized debt obligations (CDOs), and many other instruments that relate in one way or another to mortgages. To value these instruments and measure their risk, one needs to define how the components of a basket behave relative to each other. For this, one has to make assumptions on the likelihood of each component defaulting given the survival or default of other components of the basket. Rating agencies are therefore major users of a wide range of statistical and theoretical models in this domain. What the events of 2008 showed is that correlation among securities in extremely volatile markets can reach levels that widely surpassed those that the models and theories had assumed. This probably contributed to the breaching of risk limits, and perhaps resulted in a forced liquidation of some assets. The liquidation frenzy acted like a "feedback loop" adding to the volatility of the markets across asset classes.

Chart 6 illustrates how during the market turmoil of the fourth quarter of 2008, correlation between daily returns on securities, which intuitively or historically was negative or low (for instance, sugar or oil and the S&P), increased and converged towards 1.

3 months rolling correlation of S&P500
with various securities *Chart 6*

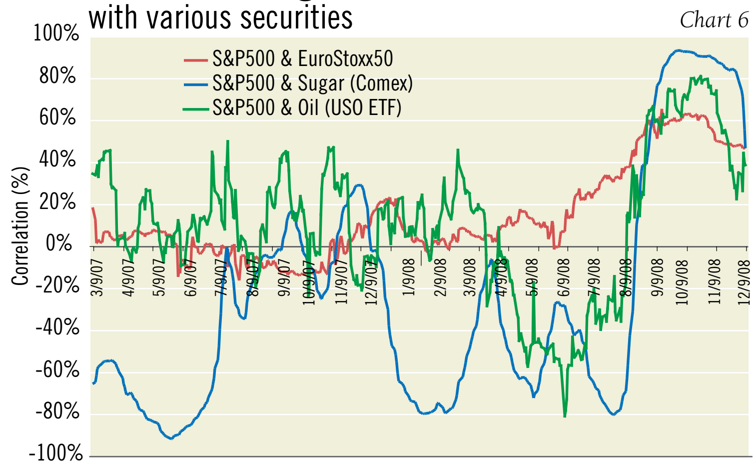

Now, let's take this phenomenon and apply it to a basket of mortgages. If everyone becomes "correlated" (i.e. defaults) under strain, more than what we ever assumed originally, then we have a glimpse of the repercussions: erroneous pricing resulting in unforeseen losses, additional risk eating away into the risk limits, etc. Take away: Beware of "common wisdom," and yes, correlations can go to 1. Make sure you account for it.

After the storm

On an intuitive basis, it seems that volatility "hits" the market. Periods of low and moderate volatility are brought to an abrupt end, often due to an external catalyst. The 1987 crash, the 1998 crisis, or the bursting of the tech bubble exhibited this type of jump in volatility. These shocks were followed by a gradual relaxation of volatility.

This shock/relaxation characteristic of volatility has an interesting resemblance to various natural phenomena.

continued... pg.19

'Investment Theme' continued from pg. 18

Let us think of earthquakes: An initial shock is often followed by a series of aftershocks, but the overall trend is for things to revert back to normal over a time period. Charts 7 and 8 illustrate the similarity of the shock/relaxation between S&P500 returns (two-minute intervals, September to December 2008) and a seismogram of a 2006 earthquake in Pakistan.

SPX 2 min returns Sep13 to Dec30 2008 *Chart 7*

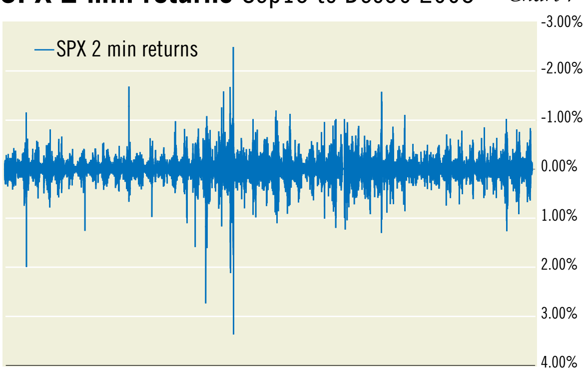

Seismogram of Oct 7, 05 Pakistan earthquake Magnitude 7.6 *Chart 8*

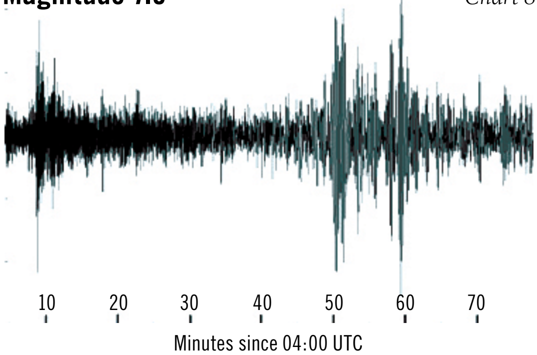

In the case of earthquakes, one can use the Omori law to estimate how many aftershocks, bigger than a defined threshold, are likely to happen in a given time after the initial earthquake. What if we tried to see if Omori's law is applicable to volatility shocks in financial markets? Lillo and Mantegna[2] showed in a short paper that the Omori law provides a good fit for the behavior of very short-term volatility (order of one minute) in the days following the 1987 and 1998 market crash episodes[3].

For the purposes of this article, we reviewed the behavior of the S&P500 (using the S&P futures in order to have quasi-continuous data). We selected the 3 October to 30 December period. Interestingly, we found that Omori law also provided a good fit for the post-shock period in the fourth quarter of 2008. It is comforting that our study found a coefficient p (which characterizes the speed at which the intensity of shocks diminishes), near one, which is in line to the one obtained when studying the 1987 and 1998 shocks; this suggests that volatility shocks have a common behavior.

Take away: While we cannot predict when a volatility shock will jolt the financial markets, we have a reliable tool that can be used to characterize how the markets will "relax" after such shock. From a practical perspective, for the S&P500, it appears that the extreme moves (corresponding to the "tails") subside noticeably over a 60-trading day period following the initial shock. In a world where we do not have any viable predictive tools, such information can be quite valuable for traders and risk managers.

Art or science

Over the last few decades, the portrayal of financial disciplines by academic institutions and the media have–purposely or not–resulted in a strong belief that these disciplines are akin to exact sciences. We believe that far too many market participants use a theoretical framework; this, after all, is a product of our own ingenuity, much as the "definite and absolute theory of finance."

The reality is that financial markets only translate the aggregate behavior of a great number of participants, and as such, it is unlikely that a set of assumptions or models will ever capture the whole picture. Unfortunately, most often this hypothetical "truth" proves to be wrong or incomplete only when the extreme scenarios materialize. By then, however, it is often too late to remedy the situation. Given the complex web interconnecting global financial markets, consequences of this perception of finance as a "pure science" can be dire. Financial disciplines will always rely on mathematics and be inspired by physics; nonetheless, they probably pertain as much to the realm of social sciences as they do to that of exact sciences. While models and a solid quantitative framework are necessary for risk management and derivative pricing, common sense and qualitative vetting are an absolute must if we seek to avoid building a house of cards ready to be blown away by the reality of financial markets.

[1] While studies show that past volatility has some informational content about furure volatility from a statistical perspective, this has not translated in any type of practical forecasting tool.

[2] *Power law relaxation in a complex system: Omori law after a financial market crash*, F. Lillo and R.N. Mantegna, Phycisal review. E, 2003, vol168 (2) no.1, pp. 016119.1-016119.5

[3] Borland and Bouchard went futher in developing a model using a similar principle with the addition of feedback and mechanisms. See: *On a multi-timescale statistical feedback model for volatility fluctuations*, L. Borland and J.P. Bouchard, arXiv:physics/0507073 v1, 10 Jul 2005