

Is VaR to blame for the downturn?

During the ongoing financial crisis, one risk metric that has received much criticism is Value at Risk or VaR. **Angus Hung** and **Michael Langton** put this into context and advise on how to get the best from risk metrics

Author and specialist financial derivatives Nassim Nicholas Taleb was recently quoted extensively in a New York Times article entitled “Risk Mismanagement”. He makes some valid points with regard to the usefulness of risk metrics at times of extreme market behaviour. But while VaR certainly has its laundry list of problems, Taleb takes VaR out of context by focusing on only one version it: the Gaussian based parametric VaR, which he rightly points out is severely constrained by the dangerous assumption that asset returns follow a normal bell-shaped distribution.

In fact, he even goes so far as to say that VaR was highly responsible for the current financial crises. This is rather disturbing, as his claims seem to have gained a wider currency, thus detracting from the infinitely more important issues behind the crisis. If we look back in history, we can see quite clearly that most “blowups” were not due to poor allocation decisions based on an over-reliance on risk measurement and optimisation models, but were about leverage, unchecked greed, operational disaster and outright fraud.

While VaR is a requirement for a bank, most traders and fund managers would laugh if you asked them if they took VaR very seriously. The reality, alarmingly, is that risk managers have hardly any clout when it comes to strong-arming a trader or liquidity.

Risk manager warnings are often ignored or overridden as senior management tends to focus purely on profitability, not risk. This is not a risk model problem, but a corporate governance problem. Instead of bashing risk managers, we should be giving them more independence, capabilities and authority to identify and limit excessive risk taking.

Long-Term Capital Management (LTCM) was leveraged 100 times at one point and Bear Stearns’ credit hedge funds over 40 times. A simple cap on gross exposure would have helped to avoid the problems they encountered with leverage. Of course, this would have interfered with a strategy that depended heavily on leverage to boost minuscule returns.

In the 1990s, Nick Leeson at Barings, the Orange County debacle, events in Mexico and Korea – all of these events had excessive leverage in common. The problems that lie within VaR are its inability to fully capture leverage and liquidity risk. Good risk managers are fully aware of this shortcoming and, as a result, VaR is only one in a whole repertoire of tools, both quantitative and qualitative, that risk managers use to get a sense of the risks they are taking on.

Taleb gives the impression that risk managers are only managing risk according to Gaussian principles, where probabilities are assumed to be normally distributed. There is more to the story than he lets on. Interestingly enough, Taleb seems to be a big fan of Monte Carlo simulations (a method that does not need to assume normality in asset return distributions) as seen in his use of Monte Carlo in the book 'Fooled by Randomness'.

Taleb suggests Monte Carlo simulators allow us to learn from the simulated future which is superior to learning from the past, because the past has a survivorship bias, and we also tend to denigrate the past by claiming misfortune had by others will not happen to us. Most sophisticated risk managers use Monte Carlo very much in the same way he does.

To his credit, though, managers at LTCM were proponents of parametric VaR, which severely underestimated risk and was largely responsible for its catastrophic collapse. But even in the case of LTCM, leverage was the real culprit, just as it is in today's financial crisis. VaR or no VaR, being leveraged 100 times can only lead to eventual disaster should markets move against you and liquidity dries up. No fancy risk metric is needed to logically arrive at this conclusion, just common sense. The problem with LTCM is that it relied too much on VaR, and was lulled into a false sense of security pertaining to risk by not properly addressing the inadequacies of VaR via other means, such as stress testing.

Philippe Jorion, considered the foremost authority on VaR, says VaR is like having a wobbly compass in a dense forest. It can point you in the right direction, but it will never give you the exact coordinates of where you need to get to. Instead, you will still need to be acutely aware of your environment, being mindful that unforeseen pitfalls can occur along the way. VaR won't get you there alone. It needs to be coupled with other tools. In addition, you need to pay close attention to your instincts while, at the same time, incorporating a healthy dose of scepticism.

To highlight Taleb's primary gripe with VaR (specifically Gaussian parametric VaR), we conducted a back-test whereby for each day during December 2008, we calculated one day parametric and Monte Carlo VaRs at both the 95% and 99% confidence levels for a portfolio with an Asian long/short focus trading primarily in cash equities, index futures and options (puts and calls). It is worth noting that a back-test should be longer than one month, usually a year or so. But even with a month's worth of data, we can clearly see some significant results.

Fat Tail Risk

As seen in table 1, the results of the back-test will show that parametric VaR does indeed fall short in capturing fat-tail risk due the number of breaches that occurred when comparing the VaR forecasts vis-a-vis the daily static profit and losses of the portfolio. Taleb would use this as a stick to beat risk managers and software vendors with. But if we also look at Monte Carlo VaR, while not perfect, we will see that it fares much better, even in these highly volatile and turbulent times, at capturing fat tail risk.

Table 1: Back-test showing Monte Carlo VaR versus Parametric VaR

Static P&L date	Monte Carlo VaR				Parametric VaR				
	Static P&L	Monte Carlo VaR 99%	Monte Carlo VaR 95%	Breach 99%?	Breach 95%?	Parametric VaR 99%?	Parametric VaR 95%?	Breach 99%?	Breach 95%?
1/12 2008	36,540	-58,118	-39,167	No	No	-21,448	-15,188	No	No
2/12 2008	-33,266	-82,539	-58,583	No	No	-44,687	-31,645	No	Yes
3/12 2008	22,967	-54,843	-38,228	No	No	-34,901	-24,715	No	No
4/12 2008	2,878	-100,570	-71,490	No	No	-55,009	-38,955	No	No
5/12 2008	4,199	-124,273	-79,064	No	No	-71,679	-50,760	No	No
8/12 2008	-108,232	-113,360	-70,995	No	Yes	-58,512	-41,463	Yes	Yes
9/12 2008	-42,409	-57,693	-43,201	No	No	-46,451	-32,894	No	Yes
10/12 2008	-587	-61,883	-40,042	No	No	-52,943	-37,492	No	No
11/12 2008	-14,277	-132,991	-93,182	No	No	-116,753	-82,679	No	No
12/12 2008	-113,919	-173,213	-123,450	No	No	-143,422	-101,565	No	Yes
15/12 2008	-11,118	-65,603	-41,384	No	No	-42,406	-30,030	No	No
16/12 2008	15,280	-42,615	-27,504	No	No	-27,884	-19,746	No	No
17/12 2008	13,395	-51,895	-31,683	No	No	-30,022	-21,260	No	No
18/12 2008	24,150	-131,527	-75,511	No	No	-79,608	-56,375	No	No
19/12 2008	-30,341	-163,048	-109,746	No	No	-127,823	-90,518	No	No
22/12 2008	-12,808	-185,037	-125,695	No	No	-116,076	-82,200	No	No
23/12 2008	-28,833	-97,075	-63,610	No	No	-48,996	-34,697	No	No
24/12 2008	15,797	-129,951	-71,602	No	No	-79,268	-56,134	No	No
29/12 2008	33,781	-141,025	-93,478	No	No	-48,165	-34,108	No	No
30/12 2008	-5,829	-150,232	-93,872	No	No	-72,744	-51,514	No	No
31/12 2008	-15,904	-65,951	-37,883	No	No	-38,183	-27,039	No	No

Statistically, at the 95% confidence level, we would expect one breach in every 20 trading days. Using the parametric approach, there were four breaches in just under 21 days – this is absolutely unacceptable.


Furthermore, at the 99% confidence level, we should only see one breach in every 100 trading days. Already, in just 20 days, there was one breach using the parametric approach. In this instance, a risk manager should not trust the parametric VaR figures because the results clearly show that this method has failed the back-test miserably.

It is highly likely that the reason there were so many breaches using the parametric approach is because it assumes that historical daily returns tend to be normally distributed, ie following a bell-shape pattern over time – this is the supposed eureka argument Taleb has been heralding to the media.

Empirically, it has been shown that parametric VaR is good at capturing and approximating the risks of assets that exhibit linear delta-1 payoff functions such as cash equities, swaps, futures, etc. Even in the best of times, under less volatile market conditions, parametric VaR tends to underestimate risk when the portfolio includes assets with non-linear or asymmetrical payoff functions such as options (puts/calls) and other contingent claims.

Indeed, the hypothetical portfolio used in table 1 includes options (primarily risky written calls) in its portfolio where they exhibit asymmetric return distributions with fatter left tails. Therefore, parametric VaR, which assumes symmetry in the right and left tails, would necessarily underestimate risk in this case because of its thinner tails, thus resulting in a high number of breaches in static profit and loss (P&L) as against the VaR forecast.

However, we can also see that Monte Carlo fared much better than its parametric counterpart. This is because it is able to capture the risks associated with non-linear payoff functions inherent in option strategies because, unlike the parametric approach, it does not assume normality in asset returns. It can be seen above that there were no breaches at the 99% confidence level and only one breach at the 95% confidence level – well within statistical expectations.



Above we show graphs of both a parametric normal distribution and an asymmetric distribution. As we can see, the asymmetric distribution has a fatter left tail, which means there is a possibility of incurring larger losses.

This asymmetric distribution could be reflective of the portfolio used in Table 1 which contains written call options, which would explain the existence of fatter tails due to unlimited exposure to the downside due to upward price movements which could potentially more than offset the premiums collected for writing the call options. As such, parametric VaR would underestimate VaR for this portfolio because it would wrongly apply the normality assumption when calculating VaR.

Monte Carlo, on the other hand, can handle asymmetric return distributions, thus being better equipped to estimate fat-tail risk. If Taleb conceded that he wrongly pigeon-holed his attack on VaR to the Gaussian parametric method, he would still raise another complaint. The complaint would be that VaR, accurate or not, does nothing to tell you how much you could potentially lose past a certain confidence level, say 99%.

Assume we calculated that the one day Monte Carlo VaR at the 99% confidence level was equal to \$100,000 (€79,000). In layman's terms, this means that we can be 99% certain that we will not lose

more than \$100,000 over the next trading day. According to Taleb, what it does not say is how much we would be susceptible to losing for that other 1% of uncertainty.

Taleb would be absolutely correct to make this criticism of VaR. For that 1%, we are indeed uncertain about the amount we could possibly lose, ie we could lose \$150,000 or \$500,000 or some other unspecified amount – whatever the case may be, we have no way of knowing and our typical VaR calculations would do nothing to tell us otherwise about that critical 1% of the distribution. This is where the concept of Conditional VaR (CVaR), also known as expected shortfall, can try to fill in the gap.

Conditional Value at Risk

For CVaR, we again turn to the use of Monte Carlo simulations which, at the risk of sounding technical, utilise geometric Brownian Motion within the Weiner framework while using an Ornstein-Uhlenbeck mean reversion process for volatility. The risk factors during the random simulation are equity, interest rates, currencies, and volatility of volatility.

Without getting into further specifics, assume after having identified and mapped the main risk factors and calculating the covariance constraints among such risk factors based on historical data, we then generated 1,000 simulations. Each simulation represents one outcome out of 1,000 regarding the possible paths to gains and losses the portfolio could take over the next trading day. Once the simulation is finished, these outcomes can then be arranged into a histogram format – showing the worst losses in the left tail and the best gains in the right tail. The 99th percentile for the left tail of the histogram would be interpreted as the Monte Carlo VaR at the 99% confidence level and would represent the 11th worst loss in the overall simulation.

Continuing with the example above, recall that we assumed a Monte Carlo simulation generated a VaR at the 99% confidence of \$100,000. Because we have a finite set of simulations, ie 1,000 simulations, there are still 10 other outcomes in the overall histogram distribution that resulted in a loss worse than the \$100,000 reported at the 99% confidence level. Assume that these losses in the left tail of the distribution are shown in Table 2.

Table 2: Measuring CVaR

Ranking of simulation's worst losses	Histogram observation (out of 100)	Loss (US\$)	Monte Carlo VaR confidence level
11th	990	100,000	99%
10th	991	103,000	99.1%
9th	992	105,000	99.2%
8th	993	111,000	99.3%
7th	994	119,000	99.4%
6th	995	122,000	99.5%
5th	996	123,000	99.6%
4th	997	126,000	99.7%
3rd	998	129,000	99.8%
2nd	999	134,000	99.9%
1st	1,000	142,000	100%

Given the above values, in order to calculate CVaR, we simply take the average of the 10 losses that exceeded the 99% confidence level during the running of the Monte Carlo simulation, ie the average of the values for histogram observations 991 to 1,000. This would lead to a result for CVaR of \$121,400. Ultimately, with CVaR, we can derive a statistically more robust number that gives us important information regarding the expected possible loss we might incur in the area of the left tail that goes beyond a VaR confidence level.

While Monte Carlo VaR combined with its corollary CVaR may provide greater insight into the riskiness of a portfolio, especially more so than parametric VaR, Monte Carlo is no panacea and also has its caveats, which are:

- Computational complexity increases with the number of risk factors and trials.
- Random simulations are not necessarily predictive of the future. Importantly, adding more trials does not necessarily improve the simulation.
- Monte Carlo is built on the robustness of the embedded covariance matrix (ie which gives the covariance between factors). If the matrix is not applicable in terms of the future and/or fails under stressful conditions (this can be back-tested), the simulation is flawed.
- The volatility parameter is usually kept constant over the simulation horizon. While this may be less consequential for a one day forecast of VaR, it becomes more so for longer forecast horizons, say 10 days. In the markets, volatility usually does not stay constant, but fluctuates, and sometimes largely, through the passage of time. There are ongoing attempts to address this difficult issue by modeling stochastic volatility using autoregressive methods such as GARCH.

To address the above weakness, many risk managers implement historical scenario analysis to gauge how the current portfolio would have performed over a very extreme market event. There is some truth to the saying that history does repeat itself, and this is the guiding principle when using historical analysis.

Of course, there will always be new market events and the guiding principle here is that the past is not very predictive.

Compounding the problem is that even if a severely adverse market event occurs and shows striking similarities to another adverse event, we still would have no way of knowing beforehand when the adverse event was going to happen.

Although we may never know when the next big crisis will present itself, the risk manager and fund manager can still take steps to try and protect the portfolio. In this case, they are now required to truly think about unforeseen catastrophic events, no matter how remote, and how they could potentially impact the portfolio in case they actually do occur. This is where sensitivity or what-if projection analysis can help out.

Using what-if analysis, the risk manager can create user-defined stress tests where he would, in effect, shock various market risk factors by varying stepped degrees in order to try to unearth possible risks that were previously unknown or hidden. For instance, he may want to see what the impact on his portfolio would be if he shocked a particular yield curve in a non-parallel fashion, while concurrently creating spikes in volatility and shifting both currencies and credit spreads along a continuum of percentage intervals.

The end result will hopefully display possible P&L outcomes under hypothetical stressed situations. Coming up with meaningful what-if stress tests is perhaps one of the most difficult aspects in preventive risk management as gauging their effectiveness only becomes apparent after the fact.

The reader should also know that risk managers are looking at Extreme Value Theory (EVT), which specifically focuses on the fat tails seen in distributions – the area Taleb is most concerned about – in order to gain a better understanding of risk. In conclusion, while using financial risk metrics such as VaR responsibly can help one to understand the risks facing a portfolio; they should not be considered a crystal ball.

Furthermore, such metrics are no use without a good dose of qualitative skepticism, transparent and stringently adhered to compliance measures (eg stop loss, exposure and leverage limits, liquidity limits, etc) and a fundamentally sound corporate governance framework that allows the risk manager to effectively monitor, report and limit, if necessary, the fund manager's risk taking.

**Angus Hung and Michael Langton are directors of
Hong Kong-based risk management consultants QRMO**