Each of the three methods of calculating value-at-risk has its pros and cons. *Brian Shydlo* of Sirius Solutions examines them

The three flavours of VaR

★ While most energy trading firms calculate a value-at-risk number every day, few people within an organisation understand the mechanics behind the calculation. Even people who can list the three methods of VaR – parametric, Monte Carlo and historic – are not always familiar with the underlying techniques. This article introduces the mechanics of each method, and examines their individual advantages and disadvantages.

Defining VaR

VaR is a measure of the risk of an adverse effect from a movement in prices on the value of a firm's positions/trades. It is a measure of market price risk as opposed to credit risk or other forms of risk. A one-day VaR calculation at a 99% confidence level, for example, would generate a number for which the odds of losing more than that amount on any one day are exactly 1%.

All three methods used to calculate VaR are designed to produce a value matching that definition and, as result, all are subject to the same general benefits and criticisms. For example, all flavours of VaR have the feature of condensing market risk into a single statistic and all suffer from the limitation of identifying only the maximum likely loss and not the maximum possible loss.

In order to provide an 'apples-to-apples' comparison, this article assumes an identical starting point in terms of raw data then walks the reader through each method step-by-step to convert that data into the final VaR number (see figure 1). The only starting data needed is a set of trades, their current market values (MTMs) and a one-year (256 business days) history of market prices. The first step in all methods is to convert the 256 prices of a commodity to 255 returns (you get one less than the number of prices), where the return is the change in price from one day to the next.

Parametric VaR

Parametric VaR is the original method of calculating VaR. It uses a relatively simple formula based on certain parameters/inputs. The inputs to the formula are positions, volatilities (a measure of how widely prices vary) and correlations (a measure of how the prices of two different commodities move together).

The first step in calculating parametric VaR is converting the 255 returns into volatilities and correlations. The next step is converting deals into position equivalents. The simplest parametric VaR formula requires just the positions of trades (the 'delta'), though more accurate formulas take both the positions and a measure of how positions change as prices change (the 'gamma'). The final step is to plug the inputs into the formula, which produces the VaR number.

Monte Carlo VaR

Monte Carlo VaR, named for the Monte Carlo algorithm it uses (itself named for a city famous for casinos), begins – as does parametric VaR – by calculating volatilities and correlations on a set of returns.

The next step is to generate 10,000 fictional

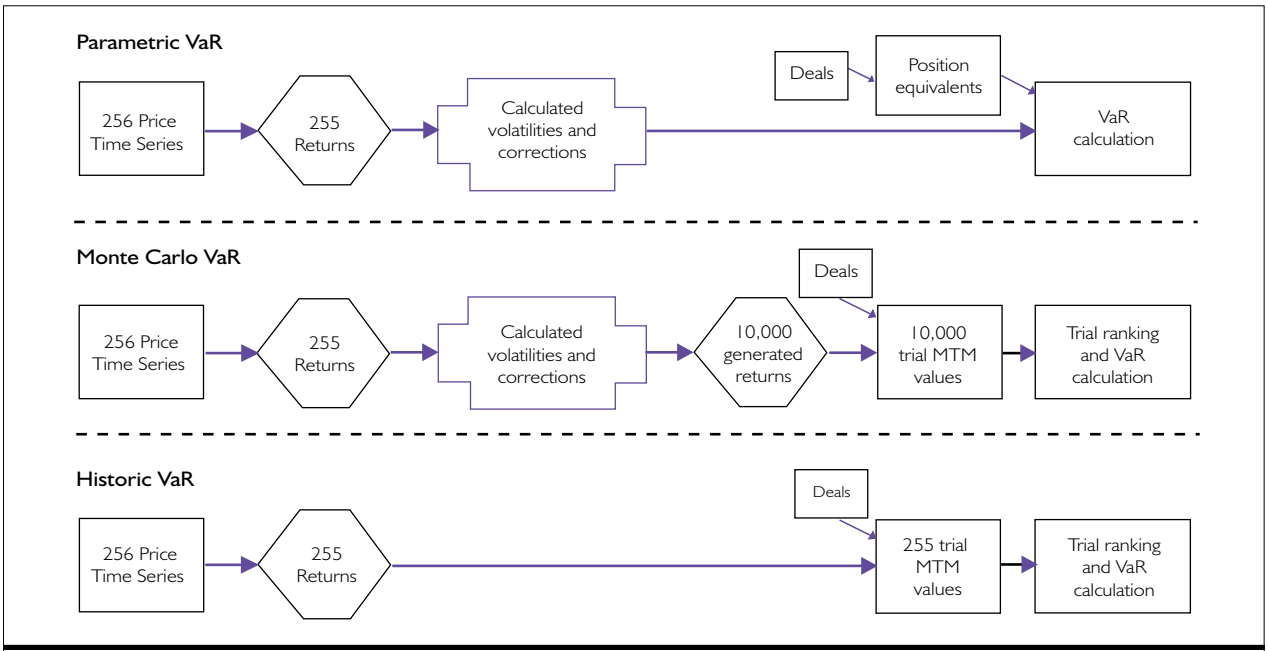

F1. Comparison of the methods of calculating VaR

Source: Sirius Solutions

returns based on those volatilities/correlations such that they have the same volatilities/correlations as the original dataset. Then the current market prices are moved/shocked up/down for each of the 10,000 fictional returns. Each deal is then revalued against the shocked returns.

Each of the 10,000 revaluations is known as a trial, so there are 10,000 trials. The 10,000 net MTM values, one for each trial, are ranked from worst to best. The 100th worst trial (100 is 1% of 10,000) is selected and the VaR is calculated as the difference between the current valuation for all deals and the valuation from the 100th worst trial.

Historic VaR

Historic VaR takes the actual returns from the prior year and, in a technique identical to Monte Carlo VaR, produces 255 trials with the valuations of the deals ranked from worst to best. Since 1% of 255 is 2.55, it is common to take the average of the second and third worst trials to calculate the 1% historic VaR.

Pros and cons of parametric VaR

Parametric VaR has the benefit of the quickest calculation time. It is also relatively simple to calculate. On the minus side, the requirement of reducing deals to their position equivalents can mean that both an option 'buy' and an option 'sale' can have the same VaR number, even though when you buy an option you have limited risk while when you sell an option you can have unlimited risk.

Another negative is that parametric VaR formulas assume that the log of prices are normally distributed (follow the bell curve) even though returns often far exceed the levels predicted by the normal distribution, which leads to misleading VaR numbers in times of extreme market movements. For example, figure 2, overleaf, shows two sets of data which have been engineered to have identical averages (returns) and standard deviations (volatilities). To the parametric VaR formula, both sets of data are the same although to the human eye they are markedly different. By converting actual returns into volatilities/correlations you lose the uniqueness of the original data.

Pros and cons of Monte Carlo VaR

Monte Carlo VaR fully revalues deals at each of the 10,000 trials and therefore is relatively better at measuring risk for options because it will produce a higher risk number when you have a net short option position.

Historic VaR relies on a relatively small number of trials and so can be highly

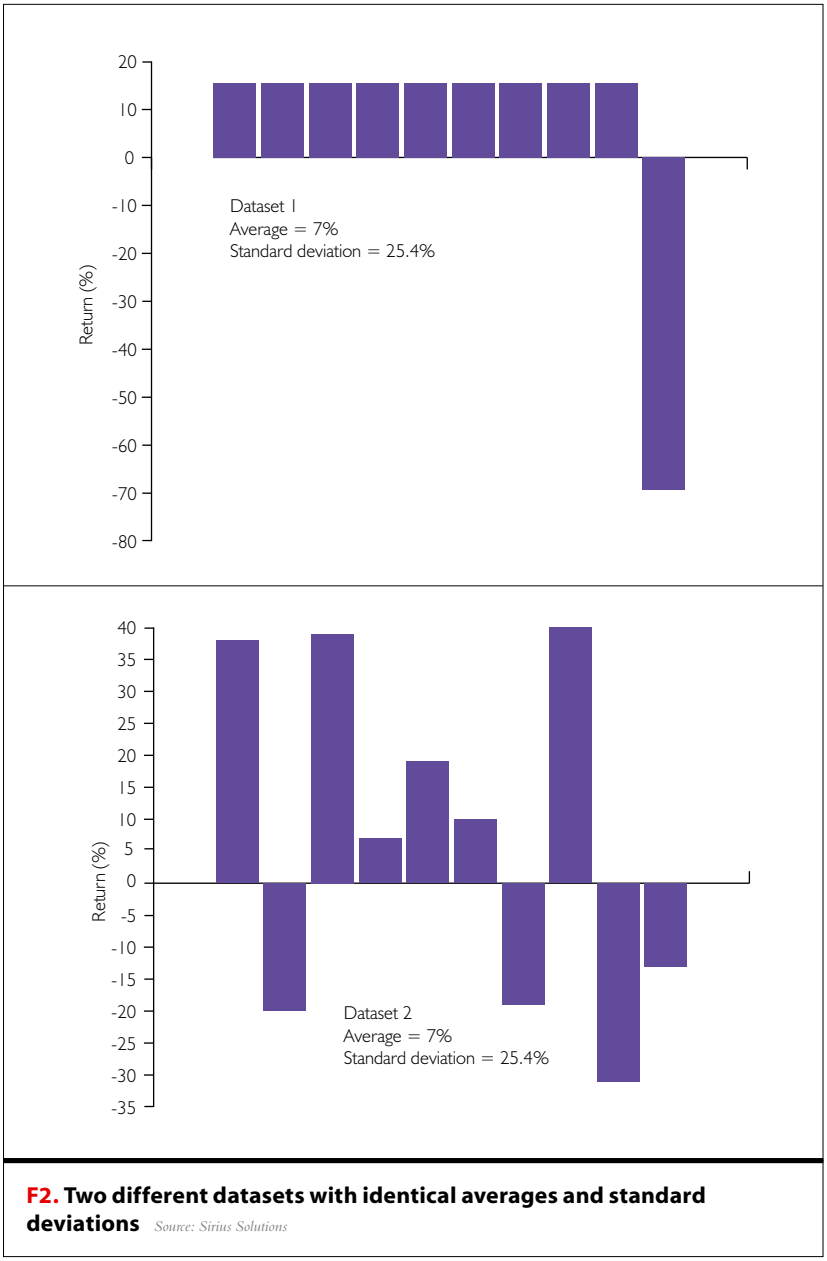

F2. Two different datasets with identical averages and standard deviations *Source: Sirius Solutions*

influenced by just two unusually bad days. Monte Carlo VaR is designed to smooth out the risk calculation to limit the impact of a few unusually bad days. Historic VaR is like watching standard TV while Monte Carlo VaR is like watching standard TV that has been modified using an upscale converter to display it on a high-definition television. While it is not as good as a real high-definition TV signal, it looks better than normal TV at the cost of increased calculations and assumptions about how to smooth the data.

Pros and cons of historic VaR

Historic VaR is arguably the simplest form of VaR because it does not require volatilities/correlations. Historic VaR also has the benefit of using the actual distribution of returns and does not make unrealistic assumptions about the distribution of returns being lognormal.

A historic VaR framework can double as a way to provide a stress test report based on specific dates of price moving events such as Hurricane Katrina. In other words, instead of taking the worst of 255 daily trials, you would use a specific date in history and shock current prices bases on the returns from that significant date.

Perhaps the biggest benefit is the ability to combine historic VaR calculations easily across multiple trading systems as you can line up each system's 255 trials by date and then add them up and rank them to produce an overall VaR number.

It is possible, but much harder to consolidate parametric or Monte Carlo VaR across multiple systems. For this reason, many energy trading firms have adopted historic VaR as their standard method of reporting VaR.

Conclusion

Market risk professionals rely on a VaR calculation for a number of reasons. These include: managing their firm's risk (often complemented by other risk reporting); communicating to management; disclosing risk in accordance with applicable regulation; and calculating the level of risk used in calculating compensation for traders who are judged on their ability to produce profits for a given level of risk.

Because of these reasons, it has always been important to understand what a VaR number means and its assumptions and limitations. Understanding which specific method of VaR is used by one's firm enhances one's understanding even further. ER

Brian Shydlo is a director at Sirius Solutions, specialising in energy trading and risk management software. He has 15 years' experience, which has included multiple implementation and upgrade projects.
Email: bshydlo@sirsol.com