

Treasury and Capital Markets Analytics

Understanding Value-at-Risk (VaR)

A Historical Simulation using Vector Risk

CONTENTS

HOW IS IT USED?	3
HOW IS IT MEASURED?	3
GETTING TECHNICAL WITH VAR	3
CALCULATION METHODS	4
VAR SIMULATION SCENARIOS	5
MORE COMPLEX CONSIDERATIONS	6



Value-at-Risk is a measure that has become the industry standard for determining market risk across derivatives portfolios that include several asset classes. It is an estimate of the worst possible loss a portfolio might suffer due to changes in market rates and prices over specified time period and a specified degree of statistical confidence.


How is it used?

Value-at-Risk (VaR) is used for managing the market risk of derivatives portfolios. Typically, there will be limits set against the VaR number to ensure that traders or portfolio managers keep their market risk within agreed levels. It is also used to obtain a market risk measure across multiple asset classes as it provides a common statistical basis for this.

How is it measured?

Different methodologies have been developed for VaR, including historical simulation, Monte Carlo simulation and deterministic variance-covariance approaches. With better technology, historical simulation has now become the industry-standard methodology.

Historical simulation takes historic time series of market rates and prices and generates daily movements observed during that history. Each of those historically observed market data shifts is then applied to current market rates and prices to generate a range of possible scenarios for where markets might move in the immediate future. The portfolio of trades is then valued under each of these scenarios, sorted into order by size of change in mark-to-market, and then the worst-case loss at a particular confidence interval (typically 99%) is selected. This is the VaR.



VaR Histogram

Getting Technical with VaR

The first step in any historical simulation (HS) VaR calculation is to value the portfolio to give a base mark-to-market. This enables us to identify all of the market data required.

The next step is to gather historic time series for the identified market data. Users need to decide the size and period of the historic time series window to be used (e.g. 2 years counting back from today or a 1-year period in the past for stress reasons).

For single price series such as equity prices or foreign exchange spot rates, this is straightforward.

For term structure data such as interest rate curves, there is the possibility that the tenor points for the curves have changed over time with new points added or others removed. In the Vector Risk solution we take the current term structure bucket points and impose this structure on the historic curves, interpolating to get values not directly observable.

For surface data, for example for volatility surfaces, Vector Risk does the same thing with the ATM tenor points and then creates the surface based on those “spine” points.

The final outcome of this process is a set of historic market data with the same term structure appearance as the current market data.

Calculation Methods


“The preponderance of negative interest rates and values near zero in recent years can lead to inappropriate relative changes. You can set absolute shifts to mitigate this.”

The next step in the calculation process is to determine the daily changes that have been observed over time from these historic market data. Traditionally, this has been done using a relative shift methodology where the ratio of the rates on consecutive days is calculated and used as the daily change. This is the default behavior in Vector Risk.

However, with the preponderance of negative interest rates and values near zero in recent years, using this methodology can lead to relative changes that are inappropriate. As a result, the Vector Risk system allows particular market data curves and prices to be set to use absolute shifts, where the difference between rates on consecutive days is used rather than the ratio. This is part of the user configuration process.

For VaR time horizons other than one day, there are a number of ways to proceed. Suppose we are considering a 10-day VaR horizon. In a non-overlapping approach, the time series differences or ratios would be calculated on points 10 days apart. For example, day 1 to day 11, day 11 to day 21, etc. This significantly reduces the sample space, because for 2 years of historic data there would be only 50 such 10-day periods observed.

An alternative approach is to use overlapping 10-day periods, for example day 1 to day 11, day 2 to day 12, day 3 to day 13, etc. This introduces an element of autocorrelation into the calculation, but it does yield a larger number of scenarios to apply from our data. Finally, in the early days of regulatory use of VaR, it was proposed to use a “one-day scaled” value for longer VaR time horizons where a one-day VaR measure would be scaled by the square root of ten to get the 10-day VaR figure. All of these methods are supported in Vector Risk.




VaR Simulation Scenarios

“Trades can be tagged with different categories and the VaR calculated based on these.”

Now that the market data changes have been determined from the historic data, these changes are applied to the current market data to generate a series of scenarios for simulation. Under each of these market data scenarios, the portfolio is valued and a change in mark-to-market from the current portfolio mark-to-market is determined.


Once this is done for all of the scenarios, the mark-to-market changes are sorted into order and the worst-case outcome to within some statistical level is determined. Thus, for a 99% confidence interval, the 99th worst out of 100 scenarios would be chosen as the VaR number. Where this is not a round number of scenarios, the system allows users to choose how the VaR number is found. We recommend “round to tail”, meaning selecting the first scenario beyond the theoretical scenario number. Other methods are supported.

STEPS FOR DETERMINING VAR



Using Vector Risk, one can calculate the VaR across one or more business unit hierarchies as defined by the user so that it is straightforward to calculate VaR for sub-portfolios as well as the global portfolio.

Vector Risk also supports a concept of categories where trades can be tagged as belonging to free-form categories and the VaR calculated based on these. For example, you may wish to calculate VaR by product or trade type, currency pair, or customer account etc.



More complex considerations

Note that a separate attribution across risk factor asset class is done automatically so users can see contributions from interest rates, foreign exchange, volatility etc. These attributions are user-definable so the user can refine them if necessary.

Note also that full portfolio revaluation is undertaken at all times under all scenarios. It is possible to supply delta-gamma or other sensitivities to the system to be incorporated into a VaR calculation for exotic trade types not natively supported, but wherever possible, full valuation of trades will be done. Vector Risk uses advanced High Performance Computing techniques to ensure simulations run in a few minutes at most, even for complex portfolios.

For foreign exchange instruments, the system uses the concept of base currency so any trades not involving that base currency will be valued using the FX prices for each leg against the base. Suppose the base currency is USD. To value a spot EUR/GBP trade the system will use both the EUR/USD and GBP/USD spot rates. The only exception to this approach is for FX options, where a volatility surface for the specific currency pair will be required.

We hope you have found this guide on understanding how to calculate Value-at-Risk useful. For more information on how we can help with your risk management, please contact us at capitalmarkets@finastra.com or visit: finastra.com/risk-compliance

Last updated: July 2019



About Finastra

Finastra unlocks the potential of people and businesses in finance, creating a platform for open innovation. Formed in 2017 by the combination of Misys and D+H, we provide the broadest portfolio of financial services software in the world today—spanning retail banking, transaction banking, lending, and treasury and capital markets. Our solutions enable customers to deploy mission critical technology on premises or in the cloud. Our scale and geographical reach means that we can serve customers effectively, regardless of their size or geographic location—from global financial institutions, to community banks and credit unions. Through our open, secure and reliable solutions, customers are empowered to accelerate growth, optimize cost, mitigate risk and continually evolve to meet the changing needs of their customers. 90 of the world's top 100 banks use Finastra technology. Please visit finastra.com

Finastra and the Finastra 'ribbon' mark are trademarks of the Finastra group companies.
© 2019 Finastra. All rights reserved.

Corporate Headquarters

4 Kingdom Street
Paddington
London W2 6BD
United Kingdom

T: +44 20 3320 5000

