



Yale SCHOOL OF MANAGEMENT
Program on Financial Stability

EliScholar – A Digital Platform for Scholarly Publishing at Yale

YPFS Resource Library

1-1-1900

Value at Risk: Parametric Method, Monte Carlo Simulation, Historical Simulation

RiskMetrics Group

<https://elischolar.library.yale.edu/ypfs-documents/557>

This resource is brought to you for free and open access by the Yale Program on Financial Stability and [EliScholar](#), a digital platform for scholarly publishing provided by Yale University Library. For more information, please contact ypfs@yale.edu.

[Open topic with navigation](#)

VaR: Parametric Method, Monte Carlo Simulation, Historical Simulation

Description: Worst-case loss over a specific time period at a specific confidence level.

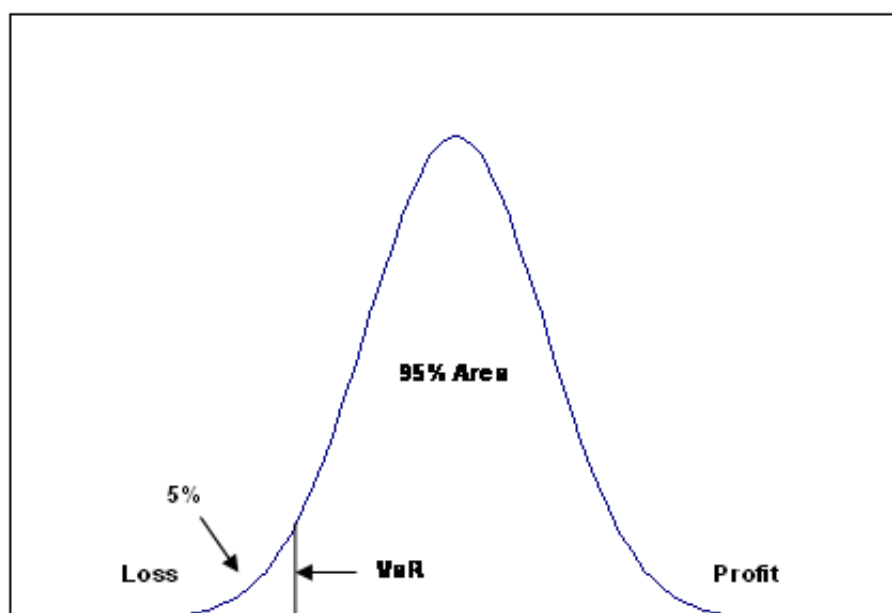
Key words: VaR; HVaR; Historical Simulation; Monte Carlo Simulation, Distribution Free; Fat Tails; N-Day, Risk Type.

Introduction

Value at Risk (VaR) is one of the most important and widely used statistics that measure the potential of economic losses. It has been adopted as the cornerstone and common language of risk management by virtually all major financial institutions and regulators. VaR measures the worst case loss over a specified time period. In mathematical terms, VaR corresponds to a percentile of portfolio P&L, and can be expressed as potential loss from the current value of the portfolio, or as the loss from the expected value at the horizon.

VaR answers the question, “how much can I lose with x% percent probability over a set horizon?” For instance, a VaR of 10 million dollars at a 95% confidence level and a 1 day horizon means that in 1 day in 20, you could expect to lose more than 10 million dollars due to market movement. The VaR amount coupled with confidence level and time period parameters, allow you to anticipate that losses exceeding 10 million dollars would occur 5% of the time; losses less than 10 million dollars would occur 95% of the time.

The figure below illustrates VaR within the tail end of a distribution curve using a 95% confidence level. The tail of the curve represents losses; the front of the curve represents profits.



VaR Methodology

VaR is typically calculated by first modeling the entire returns distribution for a position or a portfolio, then calculating the value at the percentile corresponding to the desired confidence level.

In modeling the returns distribution for a given security, RiskMetrics always begins by identifying the *risk factors* that affect the security's value. For example, the USD value of a cash position in GBP depends on one risk factor: the GBP/USD exchange rate. For a 10-year treasury bond, the risk factors would be points on the corresponding treasury curve up to 10 years out. For an equity option, the risk factors would be the underlying equity price, a discount curve, and (if specified) an implied volatility. The link between risk factors and securities allows us to *compute the value of a security in any scenario if we know the values of the risk factors in that scenario*. For example, Black-Scholes lets us price a European equity option (with a particular strike price and time to maturity) if we know the value of the underlying equity, the discount rate, and the equity volatility.

We always model the returns distribution for a security at a given horizon in terms of the risk factors that drive the security. There are three common methodologies for doing this: Historical Simulation, Parametric Modeling, and Monte Carlo Simulation. These methodologies, and their advantages and disadvantages, are discussed in more detail below.

Historical Simulation

In general, simulation methods model a distribution of a variable (in our case, the price of a security) by generating a set of (simulated) scenarios for the possible values of that variable. Historical simulation generates scenarios by applying risk factor moves from a historical period (defined by the user via start and end dates) to today's risk factor levels, then re-pricing each security from the new risk factor levels. There are *no* distribution assumptions (e.g., normality or lognormality) for either the risk factors or for the securities. Instead, distributions of underlying risk factors are taken exactly as found over the specified historical time period.

This method accurately prices all types of complex non-linear positions as well as simple linear instruments. It also provides a full distribution of potential portfolio gains and losses (which need not be symmetrical). If the underlying risk factors exhibit non-normal behavior such as fat-tails or mean-reversion, then the resulting VaR will include these effects. However, tail risk can only be examined if the historical data set includes tail events.

When computing Historical VaR, sampling history requires care in selection. Market conditions and currency devaluations may have occurred, dramatically shifting time series relationships. Also, high confidence levels (i.e., 99% and beyond) are coarse. For example, if 100 dates are used in the sample history, only one observation represents the 99% level.

In pricing models that have volatility as an explicit input used in valuation, **decay factor can change the resulting present value which in turn changes the calculated Historical VaR**. For example, in an equity option, if the user does not: a) supply an implied volatility point b) map to an implied volatility series or, c) explicitly choose to use historical volatility (calculated as historical volatility using unweighted returns from most recent 50 days), then RiskServer uses the default method for calculating historical volatility from the appropriate historical time series. The Risk Settings are taken into account when calculating this historical volatility. For example, if the user's Risk Settings specify a 1 year lookback and .94 decay and no explicit volatility input is specified by the user, RiskServer calculates historical volatility using 1 year of data and weights the returns using the .94 decay factor.



RiskServer 5.0 (9/2010) added a new *Loss Scenario* option to the VaR statistic. When *Loss Scenario* is specified, VaR returns the gain or loss (the PnL value multiplied by -1) corresponding to the *Loss Scenario* value entered (scenarios are sorted from the worst loss to the highest gain). Thus, if a user enters "5" as the *Loss Scenario*, RiskServer uses the 5th worst loss as the VaR statistic result. If both *Confidence Level* and *Loss Scenario* are specified, RiskServer uses *Confidence Level* only. Note that *Loss Scenario* is only valid for Historical and Monte Carlo methodologies, and when RiskServer Binning is not used. (RiskServer's Binning algorithm can be turned off in the Report Generation section of RiskManager's "Prefs" tab.



RiskServer 5.5 added a new *Loss Scenario Diagnostics* option to the VaR statistic, allowing users to obtain more details about the n-th worst P&L scenario. The new option provides two diagnostic functions:

2. **Historical Date:** Reports the historical date corresponding to the n-th worst P&L scenario. Note that this choice is only applicable to Historical risk methodology.
2. **Simulated Risk Factor Level:** Reports the Risk Factors' simulated level in the n-th worst P&L scenario. This function requires Risk Factor to be the last level of the drilldown, or the only one. This choice applies to either Historical or Monte Carlo risk methodology.

The Parametric Model estimates VaR directly from the Standard Deviation of portfolio returns. It assumes that risk factor returns are normal (so risk factor levels are lognormal) and that portfolio returns are linear functions of the risk factors, and hence normal as well. The last assumption makes it easy to derive the formula for the portfolio distribution from the security distribution, without having to generate any distributions explicitly.

Volatilities and correlations for the risk factors are calculated directly from time series data over user-specified start and end dates. Users may also specify an optional decay factor (λ) to govern the weight that recent observations should be given relative to older ones.

The Parametric method is accurate for linear assets, but less accurate for options and other non-linear derivatives. It also becomes less accurate at longer horizons. Parametric calculations are faster than either simulation method and don't need extensive historical data (only the correlation and volatility matrices are needed). The method is not recommended for long horizons, for portfolios with many options, or for assets with skewed distributions.

Monte Carlo Simulation

Monte Carlo Simulation is similar to Historical Simulation in that it estimates VaR by simulating risk factor scenarios and revaluing all positions in a portfolio for each trial (i.e., full re-pricing). However, instead of generating risk factor scenarios from the historical distribution, it generates them from a lognormal distribution. Thus, the distribution assumption for risk factors is the same as in the parametric model, but the method for generating the price distribution of the security is different. The method accounts for all non-linearities in portfolio positions. As in the parametric method, volatilities and correlations for the risk factors are calculated directly from time series data over user-specified start and end dates. Users may also specify an optional decay factor, as well as the number of simulations to perform per analysis.

This method accurately prices all types of complex non-linear positions as well as simple linear instruments. It also provides a full distribution of potential portfolio gains and losses (which need not be symmetrical), but does not take into account any non-normality in the underlying factors (e.g. fat-tails, mean-reversion, although Monte Carlo Simulation quantifies fat-tailed risk only if scenarios are generated from appropriate conditions).

Simulations are computationally intensive and more time-consuming than either the parametric methodology (which is the fastest) or historical simulation.

The following sections provide a more detailed review of each methodology:

Parametric Example

As discussed earlier, The Parametric Model estimates VaR directly from the Standard Deviation of portfolio returns. Volatilities and correlations are calculated directly from user-specified start and end dates.

In general terms Parametric VaR can be calculated using the following formula:

$$\text{VaR} = \text{Market Price} * \text{Volatility.}$$

We generally use the term *volatility* to express a multiple of standard deviation, depending on a chosen confidence level for VaR (i.e., confidence level multiple * standard deviation). Therefore, if we are using a 95% confidence level, volatility refers to 1.65 * standard deviation. (The 1.65 confidence level scaling factor is used because the area under the standard normal curve is between -1.65 and 1.65. If we were to designate a 99% confidence level, volatility would refer to 2.33 * standard deviation.

As an example, assume we wish to calculate VaR for a portfolio consisting of only one asset: \$1,000,000 of XYZ stock. Assume that the volatility is 1% per day. At a 95% confidence level, with a normal distribution, we can calculate VaR using the following formula: \$1,000,000 * .01 * 1.65 = \$16,500. This means you have a 5% chance of losing \$16,500 over a one day period.

If a portfolio is composed of two assets, we must factor in the correlation percentage. For instance, assume a portfolio consists of two positions in Japanese Yen and Thai Bhat, and both currencies positions convert to USD 1,000,000.

The volatility for USD/JPY is 1.08%; the volatility for USD/BHT is 1.19%. To calculate VaR at a 95% confidence level, we first multiply 1,000,000 * 1.08% * 1.65 to obtain 17,800 VaR for the Japanese Yen position. We multiply 1,000,000 * 1.19% * 1.65 to obtain 19,600 VaR for the Thai Bhat position. The correlation coefficient between the two positions is .55.

1.1

Portfolio Value	Volatility * 95% Confidence Level Scaling Factor	VaR
\$1,000,000	1.08%*1.65 (JPY/USD)	\$17,800.00
\$1,000,000	1.19%*1.65 (BHT/USD)	\$19,600.00

We can calculate the risk of the two linear positions using the following formula.

$$\text{VaR} = \sqrt{\text{VaR}_1^2 + \text{VaR}_2^2 + 2 * \text{correlation}_{1,2} * \text{VaR}_1 * \text{VaR}_2}$$

$$\text{VaR} = \text{square root} [(178,000)^2 + (19,600)^2 + (2 * 55\% * 178,000 * 19,600)] = \$32,935$$

Monte Carlo Simulation

The Monte Carlo methodology estimates VaR by simulating random scenarios and revaluing portfolio positions. Monte Carlo VaR is based on the concept of taking a current level and simulating returns of the underlying risk factors of a position instrument model. To the end user with significant non-linear exposure in their portfolio (such as derivatives with embedded options), this simulation approach with full position re-pricing is considered best practice, and is more accurate than the parametric method for estimating VaR.

Before calculating Monte Carlo VaR, one should review the underlying risk factors for an instrument and how an instrument's PV is calculated. Each underlying risk factor's input to an instrument model affects the PV of that instrument. For example, for an equity, a stock's time series of share prices is exactly the present value of one share of that equity for each daily closing price. For equities, the risk factor is the time series of closing prices transformed into daily returns.

In the case of an equity option, both the underlying equity time series, short term risk free interest rates and any implied volatility time series determines the present value of the position. With an equity option in a simple base case, two risk factors exist: equity time series and implied volatility time series (we hold short interest rates constant).

Examples on the CD-Rom

[Open accompanying Excel spreadsheet example.](#)

The Excel spreadsheet on the CD provides an example how Monte Carlo simulation can be used to calculate VaR. The spreadsheet is divided into four tabs: Time Series, Table of Random Numbers, Simulations and the PV calculation with the Monte Carlo VaR calculation.

Time Series tab

Column A consists of one hundred dates from March 21, 2005 to August 5, 2005. A table of closing prices for an equity, BHP Biliton, quoted in Australian Dollars, is listed in column B. Column C computes the daily log returns calculated by taking the closing price of the equity and taking the log of the current day's price divided by the previous day's price. Note that the choice of daily frequency is parameterized by the user in what is known as 'risksettings'. If we were to use a weekly frequency, the return horizon would be 5 business days; then we would take 5 day returns and skip over four business days. Extending this example for a fixed income positions, the returns of risk factors would be calculated for each key-rate point on the discount curve (used to model the instrument). We would then generate the log of daily returns calculated for each key rate point at the appropriate frequency (daily, weekly, etc.).

Random Numbers tab

The Monte Carlo Simulation method of VaR begins by using a random number generating function with a seed of number 1234567 to create a simulated independently distributed random number matrix. These simulated random scenarios are then used to simulate the risk factors (simulated returns tab) for each position. In the example case of a simple equity, the specified closing price from the time series is the only risk factor for that position.

We then generate a series of independently random normal variables. The number of these variables is determined by the level set by the user. The default number of simulations is usually 1000. In the case of the accompanying spreadsheet, the simulation is set to 10 (Shown in the Random Numbers tab in the Excel file).

Simulated Returns tab

Next we want to use the random variables to simulate the returns of risk factors. We do this by creating a matrix of random variables and multiplying it by the vector of returns of the risk factor (column D of the BHP timeseries tab).

Assuming an analysis date of 4th August, using the matrix of new simulated returns of the risk factor (row 101 of the BHP timeseries tab), we take the level of the risk factor (equity value) on the 4th August in this case, our analysis date, being equal to AUD\$19.80 and multiplying this by each simulated return. This effectively means we take the level (as of the analysis date) and the simulated returns to produce 10 NEW simulated levels of the August 4th 2005 BHP share value. If, you had 1000 simulations, you would effectively be creating 1000 new simulated levels for BHP on the 4th of August.

BHP Timeseries MC tab

Once the simulated returns are prepared, the PV of the equity position can be calculated. This is done by taking the equation for the PV (cell C2 which is the August 4th value) and plugging in the new simulated market prices (C5 to C14). Mathematically this is accomplished by multiplying AUD\$19.80 by the exponential of the return.

I.e. $\$19.80 * \exp[\text{simulated return}^{[1]}]$. Repeating this calculation for each return produces 10 new “simulated” PV's (cells C5 to C14 of the ‘BHP timeseries MC VaR’ tab of the spreadsheet).

The final part of the process is to produce a P/L for VaR analysis. This is computed by subtracting each of the simulated new PVs from the PV on the analysis date. In the Excel example, this is shown in column F (cells F5 to F14) of the “BHP time series MC VaR” tab. In this way, the P/L is a simulated number from a constant “benchmark” value (from the 4th August 2005 price level C2).

As with all VaR numbers, the user must select a confidence level. For the purpose of simplicity we selected 90%. We are calculating the dollar amount that losses will exceed on one day out of 10 trading days for one share of BHP.

To calculate this number, we first rank, in descending order, the P/L levels of one share of BHP. Then, to obtain the 90% confidence level, we select the 9th out of 10 P/L number (of the sorted cells I5 to I14 and look at I13). In this case, the value is AUD\$2.23.

In other words, if one held 100 shares of BHP, the one-in-ten day loss would (90% MC VaR) be equal to a loss of AUD \$223.00.

Historical Simulation

Historical VaR is based purely on historical price movements and applied to current portfolio prices. An important distinction, compared to other methods, is that instead of generating a new distribution from historical statistics, we use the historical distribution directly. Historical VAR is often preferred as it is more likely to capture unlikely events or fat tails. Another reason for its popularity is its full transparency since losses can be associated with actual market movements with particular dates. This is especially popular with regulatory reporting where full explanation for excesses is required.

The measured VaR and volatility is influenced by the historical period chosen to calculate the VaR. If historical period price changes do not resemble the future period, then the forecasted VAR will not be a good indicator of the real risk.

The following sections are organized to first build on a simple equity model and explain historical VaR for 1 day. This is then extended to demonstrate the impact of VaR for greater than a day and of additional risk factors such as foreign exchange risks. Finally an equity option is explained. Accompanying this section is an Excel sheet that demonstrates the calculations.

Calculations: A Single Equity for 1 Day Holding Period

[Open accompanying Excel Spreadsheet example.](#)

Consider “Vodafone” trading in London Stock Exchange. VaR for 10Jun05 is displayed below with the following choices for calculations:

Analysis Date	10 Jun 2005
Pricing Date	10 Jun 2005
Analysis Horizon	1 Day
Sampling Frequency	1 Day
Historical Look-Back Period	1 Year
Decay Factor	1
Currency	GBP

Equity Price (GBP)	=	1.355
Present Value (1m Shares; GBP)	=	1,355,000
Historical VaR 95% (GBP)	=	21,843

This HVaR indicates that the portfolio can lose £21,843 or 1.6% of its value with 5% probability based on the volatility of the past 1 year historical price realization. To aid the more interested reader, the actual calculations are provided in the accompanying Excel worksheet on the CD. In the worksheet, the HVAR is returned using the Percentile function which pinpoints the HVaR level at the 95 percent level. The 1 day HVar in cell D11 is extracted from the 1 day P&L returns in the cells I24 through I286. The 10 day HVar in cell D12 is extracted from the 10 day P&L returns in the cells N33 through N286.

To arrive at the HVaR one of the important inputs is the daily simulated P&L from historical data, referred here and after as trials. Figure 1 displays this simulated P&L from historical volatility of Vodafone for the past.

[1] For example, sheet 'Simulated Returns cell C8 which is the sum of all simulated returns C10 to C108.