Value at Risk: Assignment 4

Jake Mc Leroth (24913693)

06/05/2025

Contents

Introduction

The following assignment is split into two sections, Historical Value at Risk and Monte Carlo Value at Risk (VaR). In both cases we will look at calculating VaR for option portfolios. The first section on historical VaR is explained in a manner suitable for an audience with only a basic statistical knowledge. The section on Monte Carlo VaR is then presented to an audience with a deep understanding of statistics. In both sections, methods of calculating VaR for options portfolios will be looked at in a real-world context. For the South African data, the index used is the JSE Top 40. For volatility, the SAVI is used, and for the risk-free rate, the STEFI Composite Index is used. All calculations are done in **"hist_var.R"** and **"mc_var.R"**.

Option Portfolios: Historical Value at Risk

What is Value at Risk?

Imagine you have a portfolio of stocks. How do we measure risk in the portfolio? Value at risk is a way of doing this, by answering the question: what value are we sure our portfolio won't go below? More specifically, the 1-day, 1% VaR is the amount of our portfolio we are 99% sure we won't lose more than in 1 day((Alexander, 2009). To make this more clear, we look at an example. Lets say our portfolio is worth R1 million, and that our 1-day, 1% VaR is R100 000. This means that we are 99% sure we won't loose more than R100 000 in a day, or in other words, we are 99% sure our portfolio value won't go below R900 000 in 1-day. This value gives us an idea of the risk of a portfolio. If portfolio A has a higher VaR than portfolio B, then portfolio A is the riskier portfolio. In this section, we show how VaR is calculated for portfolios of options, using historical data. To do this, we adapt examples from Alexander (2009) to a South African context.

To calculate VaR, we need many values of hypothetical returns the portfolio might experience. This is called the returns distribution. These values are calculated from historical data. Our 1% VaR is then the value in our returns distribution that is smaller than 99% of the rest of the values, multiplied by -1 since we report it as a positive value. If we have 100 returns, then our 1% VaR is therefore the smallest value multiplied by -1. All options in this section are assumed to have a point value of R250.

Static and Dynamic Value at Risk

In this example, based on example 5.3 in Alexander (2009), we look at the static and dynamic VaR of a 30-day option on the JSE Top 40 index. This is calculated for both long and short positions on both call and put options. The data used is from the beginning of 2008 to the end of 2024. The 1% 10-day static VaR is essentially the VaR of a position that we hold for the full 10 days without making any changes. The 10-day dynamic VaR is the VaR if we were to hedge our position daily. We assume the option is at-the-money, meaning the strike price is equal to the price of the index. The index is currently at R75,381.31. For simplicity in this example, we will be using a risk-free rate of 0 and a volatility of 20%.

To calculate the static VaR, we first calculate the 10-day overlapping returns. Each of these returns represents a historical scenario that happened, and we assume that this behaviour could happen again. We use each of these returns to calculate a potential value for the index in 10 days' time. The same thing is done for the SAVI volatility index, giving us a potential volatility value in 10 days' time. We then use these values to get a potential price of our option in 10 days' time using the Black-Scholes-Merton formula, allowing us to calculate a hypothetical profit or loss. This is repeated for all the returns we have, giving us many potential values for the profit (or loss) we could experience in 10 days' time. From this distribution we get the 10-day static VaR.

To calculate dynamic VaR, we follow the same steps as with the static VaR, except we use daily returns instead of 10-day returns. This gives us potential profit or losses for 1 day in the future. We calculate the 1-day VaR from this distribution and then scale it to 10 days by multiplying by $\sqrt{h}$. This is called the square-root-of-time rule (Alexander, 2009).

The results are as follows:

Position	**Type**	**Call VaR**	**Put VaR**
Long	Dynamic	R768,956	R840,013
Short	Dynamic	R1,278,946	R1,249,679
Long	Static	R422,060	R425,742
Short	Static	R1,281,515	R1,564,591

Table 1: 10-day 1% VaR estimates for long and short call and put positions

We see that dynamic and static VaR are fairly similar for the short positions. The long positions, however, are very different, with the static being much lower. In this case, it seems that using the overlapping data has truncated the tail losses of the returns distribution.

Value at Risk and Expected Tail Loss of a Delta-Hedged Position

In this example we look at the effect of delta hedging on the VaR calculation. The Greeks are risk sensitivities that options are exposed to – delta is one of them. Delta measures the option's price sensitivity to a small change in the price of the underlying asset. For example, if the option's delta is 1, then the option value increases by R1 every time the underlying asset increases by R1. Delta-hedging means we take a position in the underlying asset to offset the effects of small changes in the underlying asset on the value of the option. We also calculate the Expected Tail Loss (ETL). Let's say we do exceed our 1% VaR estimate; then ETL is the amount we expect to lose now that we have exceeded the VaR value.

In this example we have a portfolio consisting of a 30-day call option on the JSE Top 40 Index futures, where the index futures price is R75,371.31 with a strike price of R75,381.31. The price of the option is R1700. The risk-free rate is 7.75%. To get the implied volatility, we plug in all the values into the Black-Scholes-Merton formula and solve for the volatility. We calculate the 1-day, 1% VaR and ETL of this option for both long and short positions. We then compare this to the 1-day, 1% Var and ETL of the delta-hedged portfolio. This is then repeated for a put with the same value and strike price. The VaR and ETL estimates are shown:

The effect of delta-hedging reduces VaR and ETL in all cases, sometimes drastically. This is clear evidence that delta-hedging substantially reduces the risk of the portfolio. We also note the VaR and ETL are less for long positions than they are for short.

VaR and ETL for Delta-Gamma-Vega Hedged Positions

We now look at the effect of hedging with respect to other sensitivities (Greeks). As said previously, delta is the option's sensitivity to small changes in the price of the underlying. Gamma is then the option's sensitivity to a change in delta. This is essentially a sensitivity of a sensitivity. Vega is the option's sensitivity to a small change in volatility. To hedge Gamma and Vega, we need to buy other options since the underlying asset has a Gamma and Vega of 0.

Position	Measure	Unhedged	Delta Hedged
Short Put	VaR 1%	392,348.3	115,180.9
	ETL 1%	579,458.2	214,254.1
Long Put	VaR 1%	265,724.4	40,778.6
	ETL 1%	316,208.1	59,581.5
Short Call	VaR 1%	403,357.5	119,546.0
	ETL 1%	617,590.3	217,740.7
Long Call	VaR 1%	242,096.5	37,901.3
	ETL 1%	291,557.6	64,816.9

Table 2: 1-day 1% VaR and ETL for Hedged and Unhedged Option Positions

To illustrate, we sell a 60-day call option on the JSE Top 40 index futures at a price of R1800 with a strike of R75,381.31. The index future is currently at R75,371.31. To hedge, we take positions in the underlying future, a long 30-day call with a price of R1700 and a strike of R75,366.31, and a long 90-day call with a price of R1850 and a strike of R75,396.31. We assume the risk-free rate is the same for all maturities at 7.75%. We again calculate the 1-day VaR and scale it up, assuming the positions are rehedged daily.

The 10-day, 1% VaR is R17,699.8 , and the 10-day, 1% ETL is R20,683.08. Compared to the VaR and ETL values we have seen before, these are very low for a 10-day time frame. We can see that hedging against more risk sensitivities reduces the risk of the portfolio. It is important to note, however, that the VaR is not 0. Hence, we were not immune to all risks, even when hedging. According to Alexander (2009), this is due to hedging only protecting against small changes of the risk factors, not large ones.

VaR with Greeks Approximation

We can approximate the profit or loss of an option by using the Greeks. Without getting into the mathematics of it, the more Greeks we account for, the more accurate the profit calculation. We show the effect of VaR calculations using just delta to estimate profit versus using many risk sensitivities.

The option in this example is a 30-day put on the JSE index futures, with a strike price of R75,381.31. The futures price is currently R75,371.31, and the risk-free rate is 7.75%. We look at 1-day, 1% VaR and 10-day, 1% static and dynamic VaR, in both cases for long and short positions. The results are as follows:

Table 3: Value at Risk Using Approximations of P&L - Short

PnL Approximation	Dynamic 10-Day VaR (R)	Static 10-Day VaR (R)
Delta	969,078	953,566.5
Delta-Gamma	1,201,671	1,667,096
Delta-Gamma-Vega	1,277,960	1,797,594
Delta-Gamma-Vega-Theta	1,255,922	1,728,037
All Greeks	1,251,586	1,727,272

Table 4: Value at Risk Using Approximations of P&L — Long

PnL Approximation	**Dynamic 10-Day VaR (R)**	**Static 10-Day VaR (R)**
Delta	1,010,180	994,010.3
Delta-Gamma	777,586.7	343,303
Delta-Gamma-Vega	728,417.8	314,297.6
Delta-Gamma-Vega-Theta	750,457.3	383,859.3
All Greeks	758,469	384,517

Table 5: 1-Day Value at Risk Using Approximations of P&L (ZAR)

PnL Approximation	**Long Position (R)**	**Short Position (R)**
Delta	319,446.9	306,449.4
Delta-Gamma	245,894.5	380,001.8
Delta-Gamma-Vega	230,345.9	404,126.4
Delta-Gamma-Vega-Theta	237,315.4	397,157.5
All Greeks	239,849	395,786.2

For long positions, adding gamma decreases VaR, and vice versa for short positions. Then, adding vega to the long positions further reduced VaR due to the negative correlation between the changes in the index and volatility (Alexander, 2009). After adding Vega, the changes by adding more risk factors did not have a big effect on the VaR estimate. We can compare the values in table 5 with the values in table 2, since both examples are on the same option and time frame. We see when more risk factors are used for the approximation, VaR is always smaller than when calculated through pricing with the Black-Scholes-Merton formula. This is likely due to the risk sensitivities accounting for small changes, whereas the idea of VaR is about big changes.

Option Portfolios: Monte Carlo Value at Risk

Non-Linear, Non-Normal Monte Carlo Value at Risk

The following example, based on example 5.11 in Alexander (2009), shows how to capture a non-linear relationship between volatility and index returns and simulate non-normal VaR through Monte Carlo simulation.

Consider the following call option on the JSE Top 40 Index Futures:

- Call price: R1700
- Futures price: R75,371.31
- Strike price: R75,381.31
- Risk free rate: 7.75%

The implied volatility is calculated using the **uniroot()** function in R since there is no closed-form solution to volatility in the Black-Scholes-Merton formula. Firstly, we capture the relationship between the log returns of the index and the log returns of volatility using the following polynomial regression model from Alexander (2009):

$$Y_t = \beta_1 X_t + \beta_2 X_t^2 + \varepsilon_t$$

where Y_t is the log returns of volatility, and X_t is the log returns of the index. Note that this model is fitted without an intercept. Fitting this model using the SAVI for volatility and the JSE Top 40 index returns, we get the following:

$$Y_t = -1.189434017 X_t + 4.903345222 X_t^2 + \varepsilon_t$$

The standard error of the residuals is 0.03026223, which is used to simulate normally distributed error terms when using the model. Now that we have a model for volatility, we need to simulate returns. This is done using both the student-t distribution with 6 degrees of freedom and the normal distribution. In both cases the mean and standard deviation of returns were estimated from the JSE Top 40 Index returns. We then calculate the volatility using the quadratic regression formula by plugging in the simulated returns. From here, we can simulate future values of the JSE Top 40 as well as vol:

- $F_{sim} = F * exp(r_{sim})$
- $Vol_{sim} = Vol * exp(c_{sim})$

where F and Vol are the current futures price and volatility level, and r_{sim} and c_{sim} are the log returns of the index and volatility. This is done for both the student-t and normal simulated returns. For comparison, we drop the quadratic term in the regression model and calculate a linear version of volatility. The risk-free rate, in this example, is assumed to stay constant.

For each of the simulated values we get, we reprice the option using the Black-Scholes-Merton formula and then calculate the discounted profit and loss (DPNL) using the following:

$$DPNL = e^{-rh/365}V_h - V_0$$

where V is for the value of the option, since this applies to calls and puts (Alexander, 2009). This results in four long call DPNL distributions, the negatives of which are the short call distributions. We repeat what we've done with a put with the same parameters. This means we end up with 16 distributions. The 1-day, 1% VaR is calculated for each and shown in the following table: The results show the VaR estimates differ

Table 6: 1-Day Monte Carlo Value at Risk with Non-Linear Volatility (R)

Distribution	**Volatility**	**Long Call**	**Short Call**	**Long Put**	**Short Put**
Student-t	Quadratic	268,030	494,570	295,750	483,454
Student-t	Linear	270,698	492,002	298,278	480,686
Normal	Quadratic	286,477	507,712	283,189	465,341
Normal	Linear	286,820	507,421	283,458	465,165

more across distributions than across volatility models, meaning the choice of distribution has a much bigger impact than incorporating a non-linear relationship with volatility and the index.

One Step Versus Multi-Step Monte Carlo

In this section, we look at an example based on example 5.14 in Alexander (2009), to illustrate the effect of using Monte Carlo simulations to estimate VaR in one step compared to multiple steps. We have the following information for a 90-day long call option on the JSE Top 40 Index:

- Futures price: R75,381.31
- Strike price: R75,366.31
- Risk free rate: 7.75%
- Volatility: 20%

Using the Black-Scholes-Merton formula, we get the call price to be R3733.50. Now, we want to simulate the log returns for the index, volatility, and risk-free rate. To do this, we estimate the 1-day covariance matrix, which is estimated as:

$$\begin{bmatrix} 5.418090 \times 10^{-1} & 1.463890 \times 10^{-4} & 2.932839 \times 10^{-5} \\ 1.463890 \times 10^{-4} & 1.702610 \times 10^{-4} & -2.032202 \times 10^{-4} \\ 2.932839 \times 10^{-5} & -2.032202 \times 10^{-4} & 1.164891 \times 10^{-3} \end{bmatrix}$$

The order of the columns is the risk-free rate, index and then volatility. Multiplying the matrix by 10, we get the 10-day covariance matrix. For both of these we get the Cholesky decomposition, which we use to simulate the log returns.

This is done as follows (Alexander, 2009):

- Simulate a standard random normal vector $\underline{z}$
- Calculate $\underline{z}'Q + \underline{\mu}'$, this gives us a vector of simulated returns for the risk free rate, index, and volatility

Here, Q is the upper triangular Cholesky matrix, and $\underline{\mu}$ is the mean returns from the data. This is done for the 1-day and the 10-day Q matrix and $\underline{\mu}$ vector. We can then use these values to simulate values for the index, volatility, and the risk-free rate, as shown in the previous example. For the 1-day simulations, we reprice the option 1 day ahead to calculate the DPNL. We get the 1-day VaR, which we scale up using the square-root-of-time rule. For the 10-day, we reprice 10 days ahead and get the 10-day DPNL, from which we can calculate the 10-day VaR directly. For the one-step 1%, 10-day VaR, we get R710,379.40. For the multi-step, we get R709,839.70. We see the multi-step is slightly less. However, there is not much of a difference between the two estimates in this case.

Monte Carlo Value at Risk with Multivariate Delta-Gamma-Vega Mapping

In this example, based on example 5.20 in Alexander (2009), we look at a portfolio of options on the S&P 500, the FTSE 100, and the DAX 30 futures. We will illustrate the use of a multivariate risk factor mapping for the calculation of VaR. The current value of the Greeks for the portfolio is:

- FTSE Delta: -0.5
- S&P Delta: -0.2
- DAX Delta: 0.7
- FTSE Gamma: -0.005
- S&P Gamma: -0.001
- DAX Gamma: 0.004
- FTSE Vega: -150
- S&P Vega: -100
- DAX Vega: 200

The point values are FTSE 100: £10.00, S&P 500: \$250.00, and DAX: €5.00. The exchange rate is at £/\$ = 0.5 and e/\$ = 0.75. The 1-day covariance matrix between the log returns of the FTSE 100, the S&P 500, the DAX 30, the VFTSE, the VIX, and the VDAX is shown:

$$\begin{bmatrix} 0.000134745 & 0.000063415 & 0.000146388 & -0.000490705 & -0.000241770 & -0.000331826 \\ 0.000063415 & 0.000133116 & 0.000118051 & -0.000259339 & -0.000456837 & -0.000243348 \\ 0.000146388 & 0.000118051 & 0.000279472 & -0.000591391 & -0.000414085 & -0.000551510 \\ -0.000490705 & -0.000259339 & -0.000591391 & 0.003444646 & 0.001468086 & 0.002039951 \\ -0.000241770 & -0.000456837 & -0.000414085 & 0.001468086 & 0.002883070 & 0.001327802 \\ -0.000331826 & -0.000243348 & -0.000551510 & 0.002039951 & 0.001327802 & 0.002058965 \end{bmatrix}$$

We multiply this by 10 to get the 10-day covariance matrix. Again, we calculate the Cholesky matrix from this and use the method from the previous question to simulate values for the FTSE 100, the S&P 500, the DAX 30, the VFTSE, the VIX, and the VDAX. In this example, we assume the expected return is 0 and therefore don't add a $\underline{\mu}$. These simulated returns are then used in the following Taylor series expansion to estimate DPNL:

$$\Delta V \approx \Delta \cdot \Delta S + \frac{1}{2}\Gamma \cdot (\Delta S)^2 + \nu \cdot \Delta\sigma$$

We use the first term only to get the delta-only VaR, then add each term for comparison. The delta VaR is £3441.859, the delta-gamma VaR was £7774.639, and the delta-gamma-vega VaR was £6896.613. The equivalent dynamic historical VaRs are given as £4723, £7213, and £7606. For delta and delta-gamma-vega, the Monte Carlo estimate is smaller.

Conclusion

This assignment has illustrated the different methodology used for calculating VaR for option portfolios, calculated with Monte Carlo or historical simulation. Important issues that came up in the assignment was that the choice of distribution, as well as scaling, has drastic effects on the VaR estimate. Furthermore, using an approximation for DPNL can also lead to not estimating VaR correctly.

References

Alexander, C. (2009). *Market risk analysis, value at risk models.* John Wiley & Sons.