Value at Risk: Assignment 2

Jake Mc Leroth (24913693)

24/03/2025

Contents

Introduction

In this assignment, we demonstrate the calculation of Historical Simulation Value at Risk (Historical VaR) in a South African context, based on the methodology used in chapter 3 of Alexander (2009). The methods of calculating VaR in chapter 2 of Alexander (2009) made assumptions about the form of the distribution. In historical VaR, we do not make any assumptions about distribution's form. This assignment will illustrate how to calculate historical VaR, how it differs from normal linear VaR, how it differs depending on the time frame used and how to apply it to an South African equity portfolio.

Question 1: Scale Exponents

Methodology

The square-root-of-time rule is a scaling law that allows use to accurately scale VaR to any time frame we want, assuming the returns are independently and identically distributed (i.i.d.) normal random variables (Alexander, 2009). The square-root-of-time can be written as follows:

$$x_{h,\alpha} = h^{1/2} x_{1,\alpha}$$

Where $x_{h,\alpha}$ is the h-day alpha quantile of the h-day discounted log returns. As can be seen, multiplying the 1-day alpha quantile by the square root of h (i.e. the square root of time) allows the 1-day quantile to be scaled to an h-day quantile. In practice, however, the normality assumption doesn't always hold up, with return distributions often being leptokurtic and skewed. In this section, we show how to estimate a scaling coefficient based on the empirical return distribution and implement it for the JSE All Share Index (ALSI).

More generally, we can write:

$$\begin{aligned}
x_{h,\alpha} &= h^{1/\xi} x_{1,\alpha} \\
\implies h^{1/\xi} &= \frac{x_{h,\alpha}}{x_{1,\alpha}} \\
\implies \frac{1}{\xi} &= \frac{ln(\frac{x_{h,\alpha}}{x_{1,\alpha}})}{ln(h)} \\
&= \frac{ln(x_{h,\alpha}) - ln(x_{1,\alpha})}{ln(h)} \\
\implies \xi &= \frac{ln(h)}{ln(x_{h,\alpha}) - ln(x_{1,\alpha})}
\end{aligned}$$

Where ξ^{-1} is the scale exponent. If $\xi = 2$, we get the square-root-of-time rule. To estimate ξ, $ln(x_{h,\alpha}) - ln(x_{1,\alpha})$ is plotted on the x-axis vs $ln(h)$ on the y-axis for different values of h. Then ξ is estimated as the slope using least squares regression, where $\hat{\xi} = \frac{\hat{\sigma}_{X,Y}}{\hat{\sigma}^2_X}$ (Alexander, 2008). The estimation is done in the R file "q1.R".

We estimate the scale exponents for the JSE ALSI at various time frames and alpha levels. If the return distribution is stable the time frame and alpha level should not affect $\hat{\xi}$ too much. The JSE ALSI data used runs from 30/06/1995 to 12/03/2025. The scale exponent is estimated from 1995, 2005 and 2015 onward.

The Log-log plot for 1995 onward versus a 5% quantile ratio is shown:

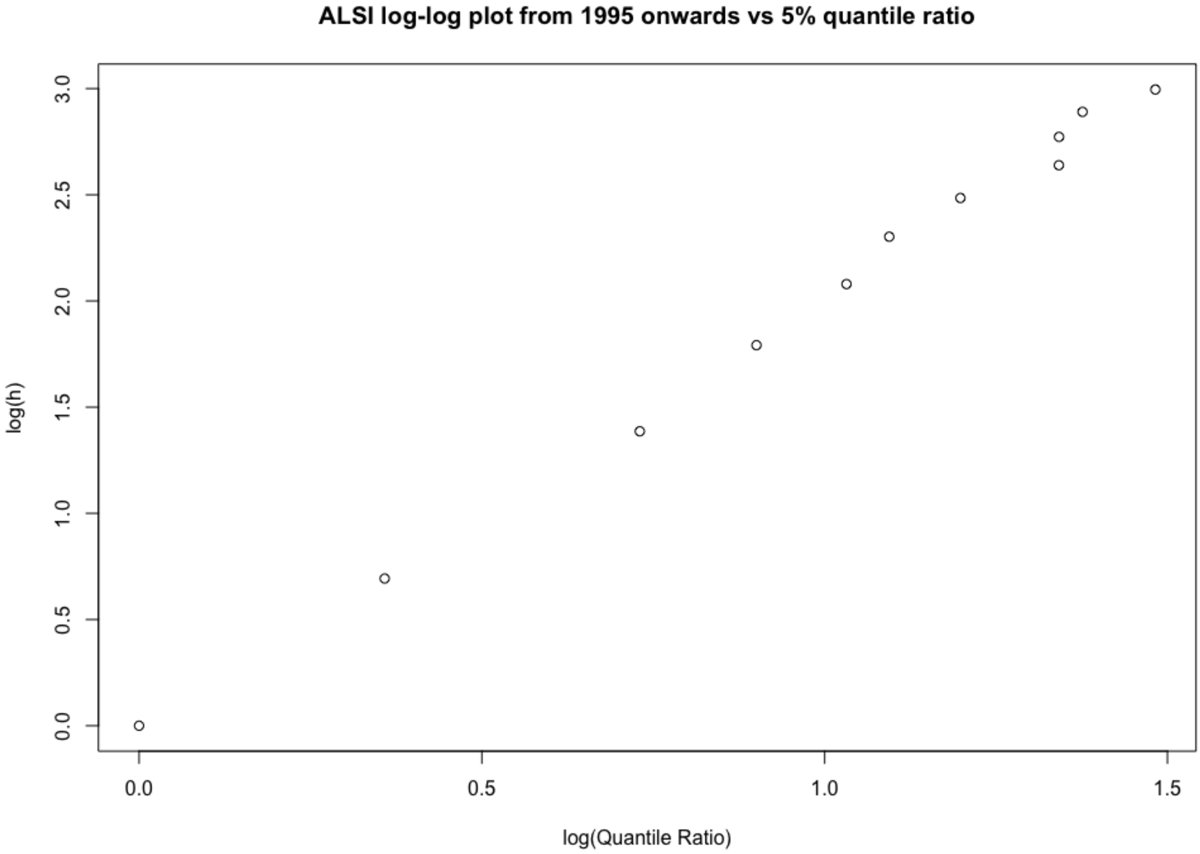

Results

The log quantile ratios are calculated for all periods and alpha levels 0.1%, 1%, 5% and 10%. The scale exponents are then estimated as $\hat{\xi}^{-1} = \frac{\hat{\sigma}_X^2}{\hat{\sigma}_{X,Y}}$, shown in the table below:

α	0.1%	1%	5%	10%
1995	0.5032	0.5500	0.4830	0.4582
2005	0.5132	0.4673	0.4377	0.4232
2015	0.4707	0.4161	0.4258	0.4383

Table 1: Estimated scale exponents for JSE ALSI

From the table, we observe that the exponents range from 0.42 to 0.55, with most of the values being below 0.5. We also see that they not only differ across time period but significance level too. This is evident of the JSE ALSI returns not being stable overtime, likely as a result of different market regimes with non-constant volatility. Had we used the square-root-of-time rule for any period and significance level, other than 1995 with $\alpha = 0.1\%$, we would have gotten an incorrect VaR value. An exponent lower the 0.5 means we would have overestimated VaR, and vice versa. For example, if we used data from 2015 onward, calculated 5% VaR and used the square-root of time rule to scale VaR, we would be massively overestimating risk. In practice, this leads to excessively conservative risk estimates and overstating capital requirements.

Question 2: Historical VaR v.s. Normal Linear VaR

Chapter 2 of Alexander (2009) calculates VaR under the assumption of normality. In chapter 3, no normality assumption is made. Instead, 1-day historical VaR is calculated as the negative alpha quantile of the historical daily return distribution. In this section, we illustrate two things: How normal linear VaR estimates differ to historical VaR estimates; and how the time period selected affects VaR estimates in general.

Methodology

The JSE ALSI price data is again used in this section with the same dates as before. We show daily the return series below:

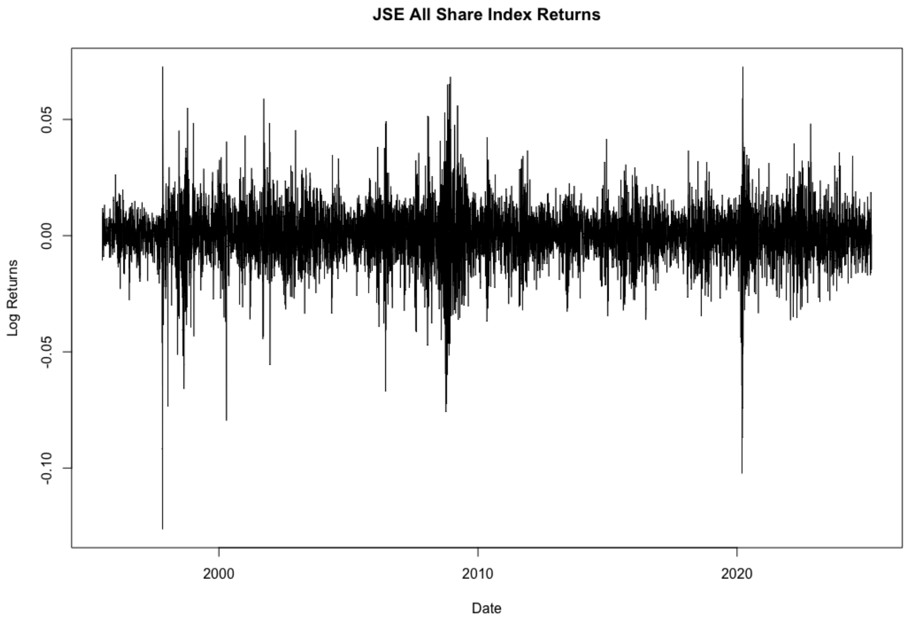

For both historical and normal linear VaR, we calculate the rolling VaR using the last 500 returns and the last 2000 returns. The 500 day rolling 1-day 99% historical VaR is calculated in the following steps:

1. Calculate the past 500 returns
2. Find the 1% quantile - the 5^{th} lowest return
3. The negative 1% quantile is the VaR for the current day
4. Move to the next day and start again at step 1

The same approach is used for the 2000 day rolling 1-day 99% historical VaR, except the past 2000 returns are used. This means that we can only start calculating VaR at the 501^{st} and 2001^{st} days of the price series.

The process for the 500 day rolling 1-day 99% normal linear VaR is as follows:

1. Calculate the past 500 returns
2. Calculate the mean and volatility of these returns
3. Calculate the normal linear VaR as $\sigma_t \Phi^{-1}(0.99) - \mu$
4. Move to next day and start agian at step 1

Again, the 2000 day rolling 1-day 99% normal linear VaR is calculated similarly and we can only start on the 501st and 2001st days. After all of these values are calculated through time, we compare them by taking the difference between the historical VaR estimates and the normal linear VaR estimates for both periods. This is done in the R file "q2.R"

Results

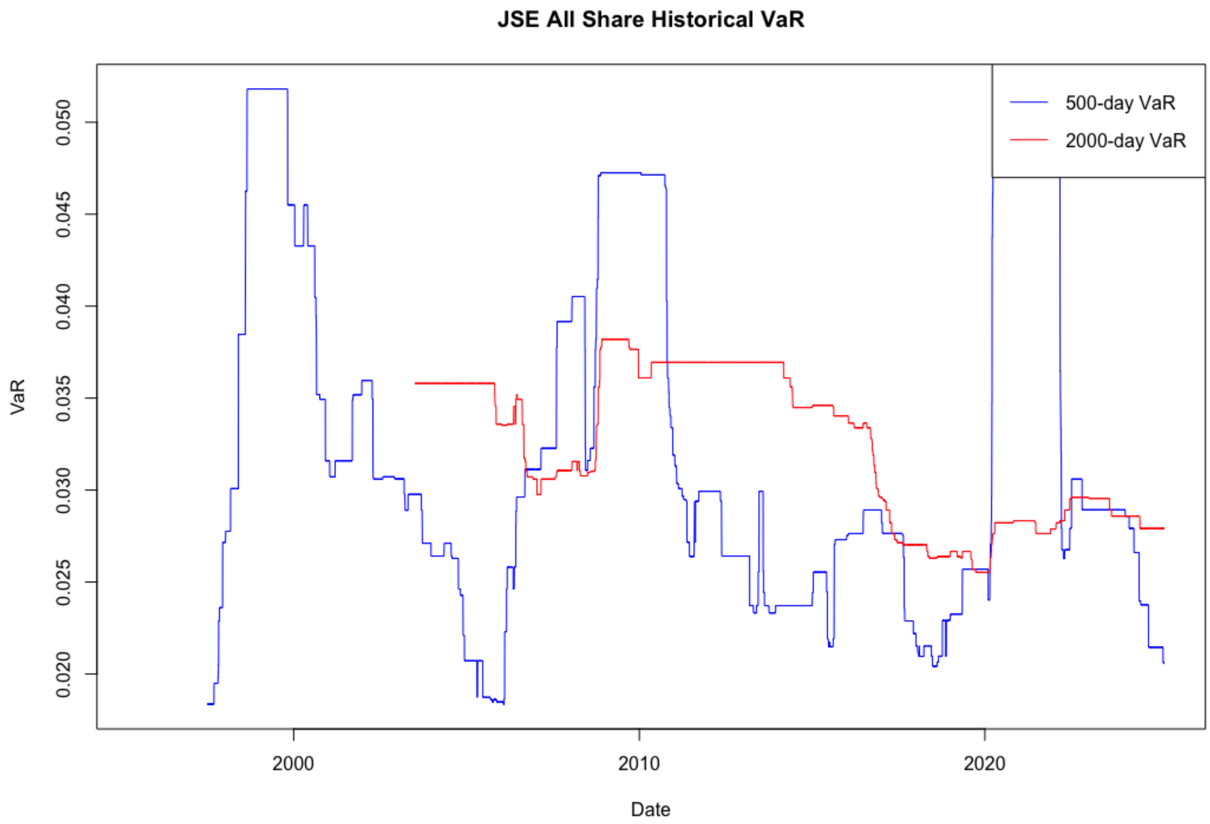

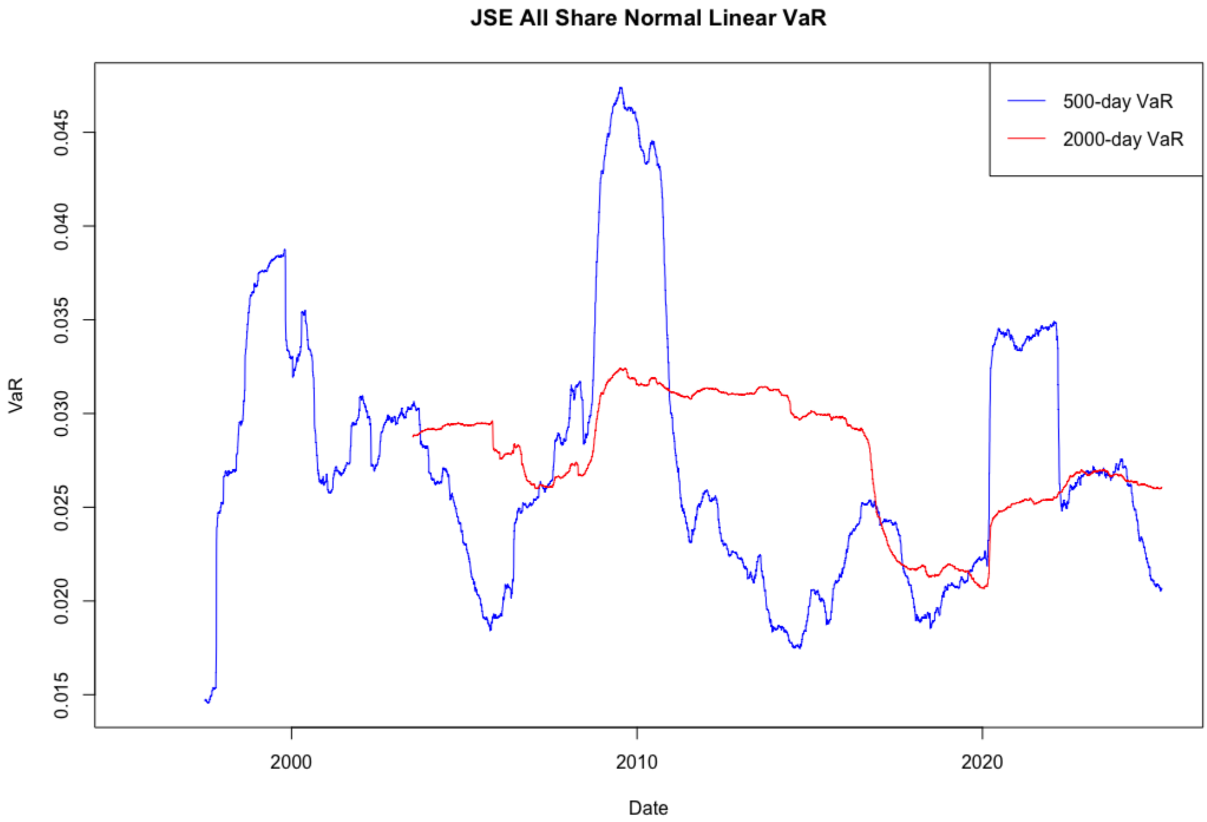

We observe that for both models, the VaR estimates vary quite substantially for the different window sizes. The smaller window size is much more reactive, while the larger window takes longer to change, exhibiting a ghost feature. This is due to the fact that all the returns are equally weighted in the calculation, so the longer window size means that the return from today will take another 2000 days before it is no longer used in the calculation, and throughout that time it will contribute the same amount as every other return - including the most recent.

We now plot the difference of the models:

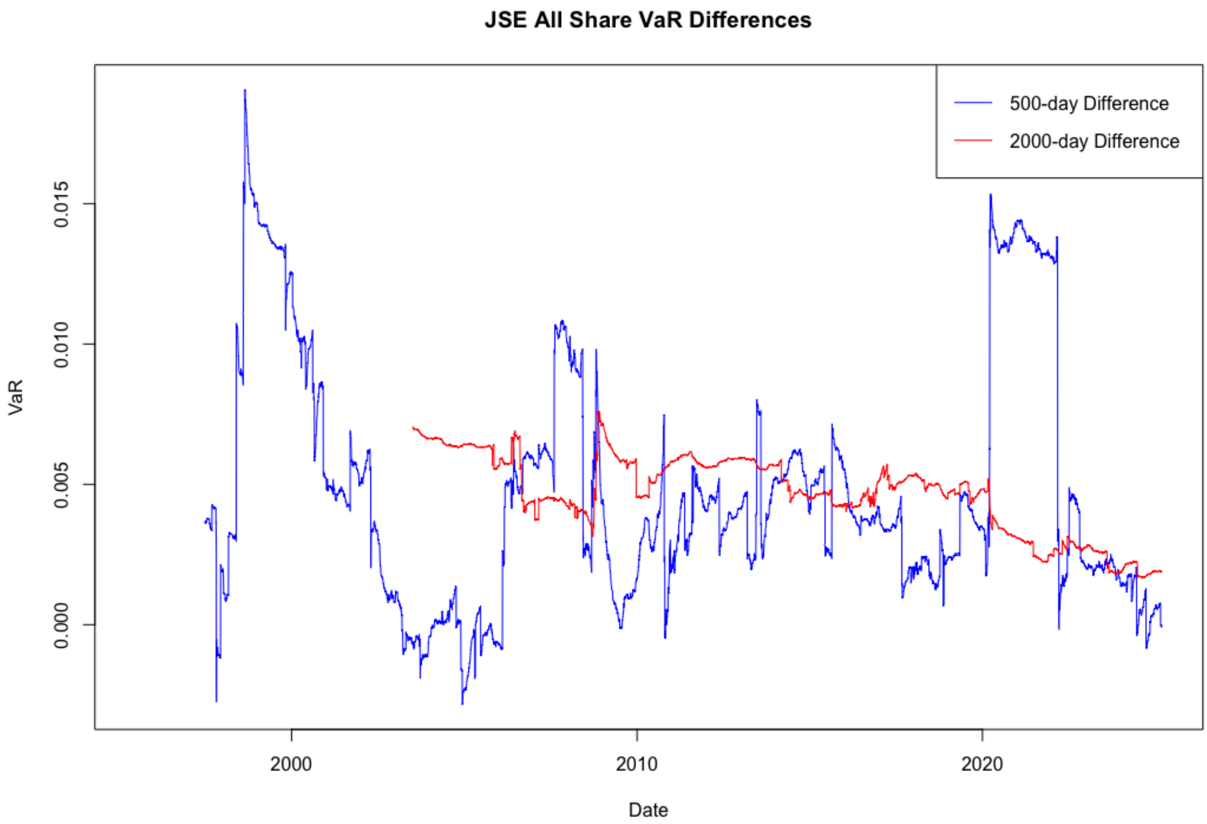

We observe that the historical VaR at most differs from normal linear VaR by 2%, with the longer rolling window being much more stable. This means that the variation between the two different window sizes within the historical VaR model and normal linear VaR model is greater than the variation between the two different models themselves. This tells us that the choice of sample size is more important in estimating VaR than the actual VaR model we choose. Moreover, equally weighting the returns means that our estimates suffer from ghost features and do not necessarily represent the present market conditions.

Question 3: GARCH Volatility Adjusted VaR

As shown in the previous section, using equally weighted returns is not ideal for calculating VaR. In this section, we make use of weighting returns before calculating the empirical distribution, as described in Alexander (2009).

Methodology

Returns can be weighted by the following formula:

$$\tilde{r}_{t,T} = \left(\frac{\hat{\sigma}_T}{\hat{\sigma}_t}\right) r_t$$

Where $\hat{\sigma}_T$ is the estimated volatility today, $\hat{\sigma}_t$ is the estimated volatility at time t and r_t is the return at time t. Hence, $\tilde{r}_{t,T}$ is the return at time t adjusted for today's market conditions. With this equation it is clear that we are not assuming constant volatility, and, hence, we need to estimate it through time. In this section, we do this using the Normal GARCH and the Normal Asymmetric GARCH (A-GARCH) models. The GARCH equations from Alexander (2008) are as follows:

$$\begin{aligned} \sigma_t^2 &= \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \quad \varepsilon_t \mid I_{t-1} \sim N\left(0, \sigma_t^2\right) \\ r_t &= c + \varepsilon_t \end{aligned}$$

Where I_{t-1} represents all the information up to time $t-1$. We see that today's volatility estimate is dependent on yesterday's estimate as well as yesterday's returns. The A-GARCH equations are given in Alexander (2008) as:

$$\begin{aligned} \sigma_t^2 &= \omega + \alpha(\varepsilon_{t-1} - \lambda)^2 + \beta\sigma_{t-1}^2, \quad \varepsilon_t \mid I_{t-1} \sim N\left(0, \sigma_t^2\right) \\ r_t &= c + \varepsilon_t \end{aligned}$$

The lambda value here, captures the asymmetry of volatility. If lambda is positive, negative returns affect volatility more than positive ones. This is known as the leverage effect. We use both models on the JSE ALSI data, estimating the parameters using maximum likelihood in the R file "q3_q4.R". Since both models make use of recursive formulas, we need starting values for σ_0 and r_0. For both models, we take $r_0 = \frac{1}{n}\sum_{i=1}^{n} r_i$ and $\sigma_0^2 = \frac{1}{n-1}\sum_{i=1}^{n}(r_i - \overline{r})^2$ where n is the number of returns in the data set. Once we have estimated the parameters and calculated the volatility through time, we use the values to weight the returns. We then use the unadjusted, GARCH volatility adjusted and A-GARCH volatility adjusted returns to estimate the historical 1-day VaR at alpha levels 0.1%, 1%, 5% and 10% on 12/03/2025.

Results

The JSE ALSI GARCH and A-GARCH parameters estimated in R, the conditional volatility and the adjusted returns are all shown below:

Parameter	GARCH	A-GARCH
μ	0.0006526	0.0004230
ω	2.172841×10^{-6}	1.0×10^{-6}
α	0.0961	0.1396
β	0.8904	0.8395
λ	-	0.0048
$\alpha + \beta$	0.9865	0.9792
Long-term volatility	20.07%	22.62%

Table 2: GARCH parameter estimates for JSE ALSI

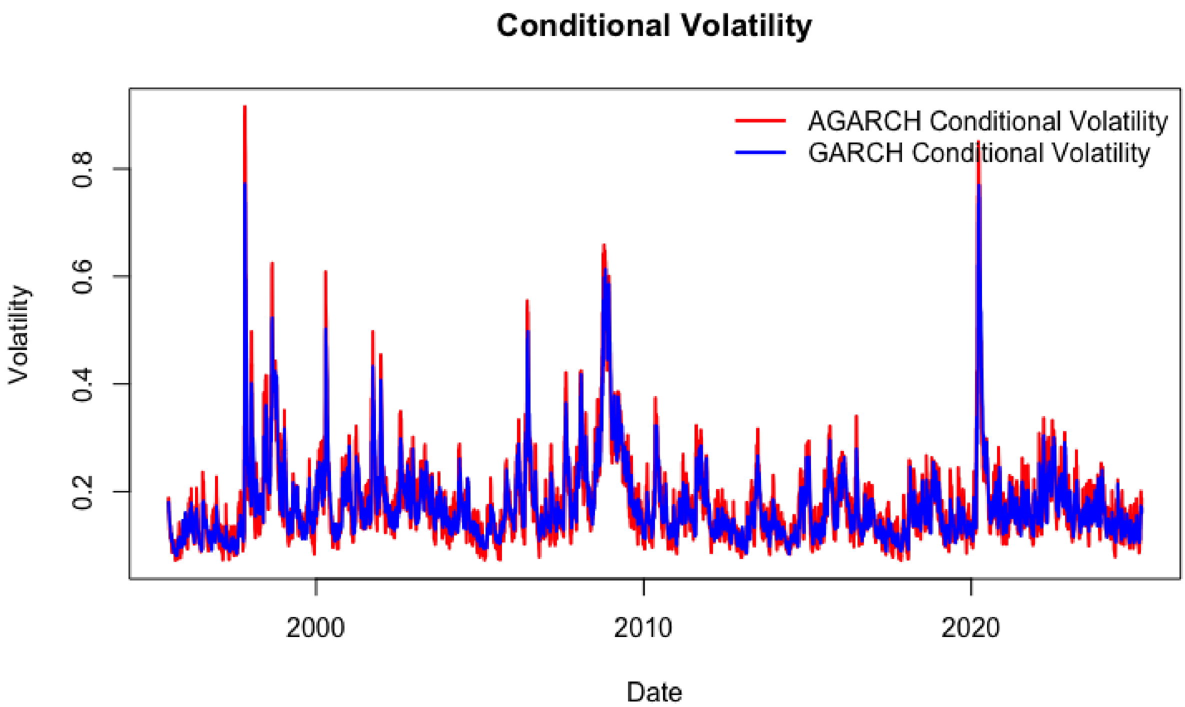

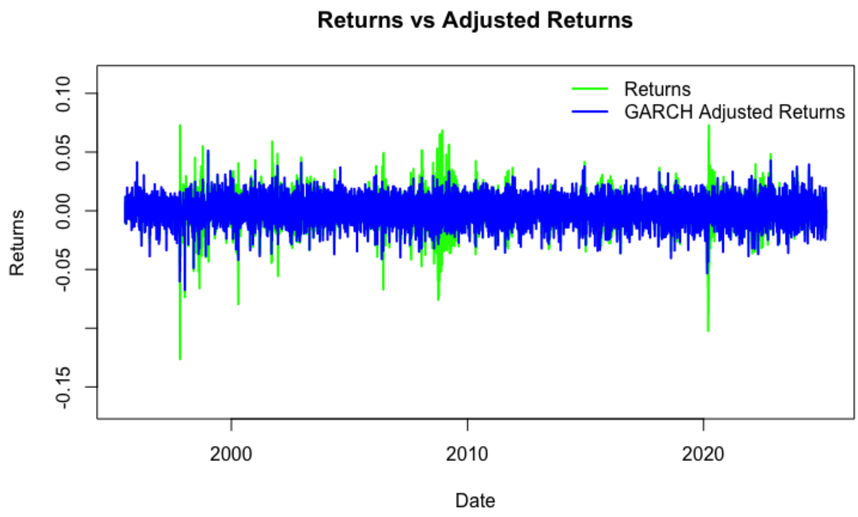
Returns vs Adjusted Returns
Returns
GARCH Adjusted Returns
Returns
0.10
0.05
0.00
-0.05
-0.15
2000
2010
2020
Date

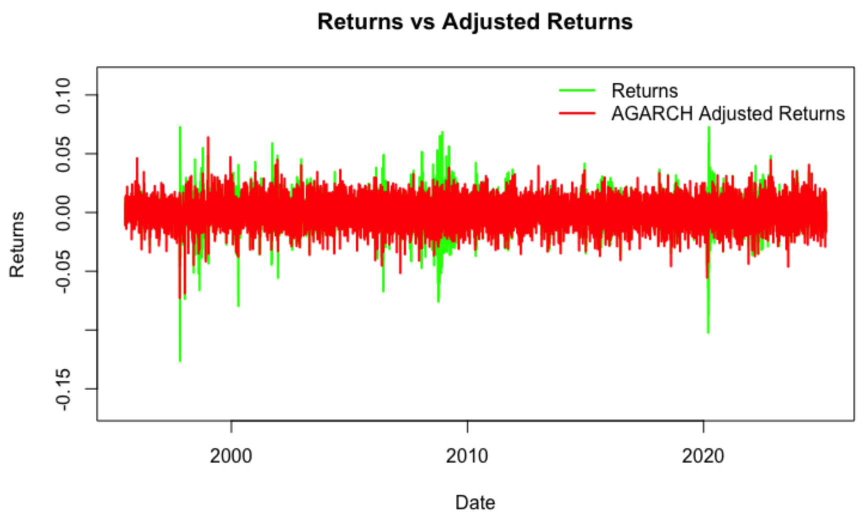
Returns vs Adjusted Returns
Returns
AGARCH Adjusted Returns
Returns
0.10
0.05
0.00
-0.05
-0.15
2000
2010
2020
Date

We can see the conditional volatility through time and the effect the λ value has on the conditional volatility, with the A-GARCH volatility values being more extreme. As a result of capturing the leverage effect, we see the unconditional volatility is higher in the A-GARCH model. Since the most recent value of volatility is fairly low, adjusting the returns to today's market conditions has the effect of making the returns less volatile through time. Using the unadjusted and adjusted returns, we calculate historical VaR and compare in the following table:

Quantile	Unadjusted	Volatility Adjusted	
		GARCH	A-GARCH
0.10%	7.01%	4.15%	4.49%
1%	3.21%	2.72%	2.90%
5%	1.81%	1.70%	1.83%
10%	1.28%	1.24%	1.35%

Table 3: Historical VaR estimates for JSE ALSI on 12/03/2025

Looking at the table, we see that the unadjusted VaR estimate is significantly larger than the GARCH and A-GARCH VaR in the tails. However, at α levels 5% and 10%, the estimates are similar. This shows the effect of adjusting the returns to today's conditions, meaning the unadjusted VaR is potentially overestimating VaR in the extreme quantiles. A-GARCH is also consistently higher than GARCH. This is expected as we are capturing the leverage effect with λ.

Question 4: Filtered Historical Simulation

Filtered historical simulation (FHS) is a method of simulating h-day returns by bootstrapping from the historical return distribution. In this section we illustrate calculating the 10-day VaR from FHS using A-GARCH volatility and compare it with the 10-day A-GARCH VaR, in which we scale up from the previous question

Methodology

Using A-GARCH parameters from the previous section, we now assume the following form according to Alexander (2009):

$$\hat{\sigma}_{t+1}^2 = \hat{\omega} + \hat{\alpha}(r_t - \lambda)^2 + \hat{\beta}\hat{\sigma}_t^2$$
$$\varepsilon_t = \frac{r_t}{\hat{\sigma}_t}$$

We set $\hat{\sigma}_0$ to the most recent A-GARCH estimate and r_0 to the most recent return, from which we can calculate $\hat{\sigma}_1$. We then calculate all ε_t using the equation above. Now, we simulate r_1 by bootstrapping values of ε_t and multiplying it by $\hat{\sigma}_1$. This is repeated many times, to get many simulated r_1 values. From this, we can then calculate many different values of $\hat{\sigma}_2$. We then bootstrap values of ε_t, this time multiplying by $\hat{\sigma}_2$ to get r_2, which we can then calculated $\hat{\sigma}_3$ from. We repeat this process until we h returns, each with many simulated, which we can then sum up to get many simulated h-day returns. We then calculate historical VaR using these simulate h-day returns. We calculate the FHS VaR for JSE ALSI using 1000 simulated 10-day returns, and compare it to the scaled volatility adjusted VaR. We also compare these values to VaR estimates assuming a most recent volatility of 10%. These simulations and calculations are done in the R file "q3_q4.R".

Results

The estimated VaR values are shown:

Quantile	Current Volatility		10% Volatility	
	FHS	Scaled Volatility Adjusted VaR	FHS	Scaled Volatility Adjusted VaR
0.10%	11.43%	14.19%	7.92%	7.95%
1%	8.65%	9.18%	6.07%	5.14%
5%	5.11%	5.78%	3.54%	3.24%
10%	3.59%	4.27%	2.32%	2.39%

Table 4: Filtered Historical Simulation v.s. Scaling

The FHS values are consistently lower at the current volatility when compared to the scaled VaR. However, when we assume 10% volatility, the values are much closer.

Question 5: Portfolio Historical VaR

In this section, we turn to a portfolio of two stocks on the JSE, namely FirstRand and Sappi. We compare different ways of estimating the historical VaR of the portfolio and then decompose it into systematic and specific VaR.

Methodology

As seen in Alexander (2009), portfolio returns are calculated as:

$$r_t = \mathbf{w}'_{\mathrm{T}}\mathbf{x}_t, \quad t = 1, \ldots, \ \mathrm{T}$$

Where $\mathbf{x}_t$ is the vector of equity returns and $\mathbf{w}_{\mathrm{T}}$ is the vector of portfolio weights at time T. In our case, $\mathbf{w}'_{\mathrm{T}} = [0.3, 0.7]$, where we are invested 30% in FirstRand and 70% in Sappi. The price series of the stocks starts 02/01/1996 and end on 12/03/2025, from which we calculate the log returns. According to Alexander (2009), we adjust returns for portfolio historical VaR in three ways:

1. Adjust $\mathbf{x}_t$ using the following:

   $$\tilde{\mathbf{x}}_t = \mathbf{Q}_T\mathbf{Q}_t^{-1}\mathbf{x}_t \quad t = 1, \ldots, T$$

   where $\mathbf{Q}_t$ is the Cholesky matrix of the exponentially weighted moving average (EWMA) covariance matrix at t. Then use $\tilde{\mathbf{x}}_t$ to calculate the adjusted portfolio returns.

2. Calculate portfolio returns and adjust these using $\tilde{r}_{t,T} = \left(\frac{\hat{\sigma}_T}{\hat{\sigma}_t}\right) r_t$ where $\hat{\sigma}_t$ is the EWMA volatility estimate at t.

3. Adjust individual stock returns without accounting for correlation. Then use these to calculate the adjusted portfolio returns.

We compare estimates using all three methods. Note that we are choosing to use EWMA here, however other volatility estimates such as GARCH can be used. Again, the starting value is taken as the overall variance. Up till this point, we have just been calculating the total historical VaR. To disaggregate this into systematic and specific VaR for the equity portfolio, we need to calculate the the portfolio's sensitivity to the JSE ALSI. The market sensitivity of each equity, using EWMA volatility and covariance, is calculated as:

$$\beta_T = \frac{\sigma_{X,M}}{\sigma_X^2}$$

Where M is the JSE ALSI and X is the equity. We then get the portfolio β by multiplying each equity β by its portfolio weight and then adding these values together. From this, we can decompose our portfolio returns into systematic and specific returns. The systematic return is the portfolio β multiplied by the market return (i.e. the return explained by the market) and the specific return is the remaining unexplained return (i.e. total = systematic + specific). We then adjust these returns and calculate the historical VaR, comparing them to unadjusted VaR as well as normal linear VaR. Further, we do this when our data ends on 12/03/2025 as well as on 01/04/2020 during the Covid-19 pandemic.

Results

The EWMA volatility estimates are plotted: The VaR estimates are compared in the following table:

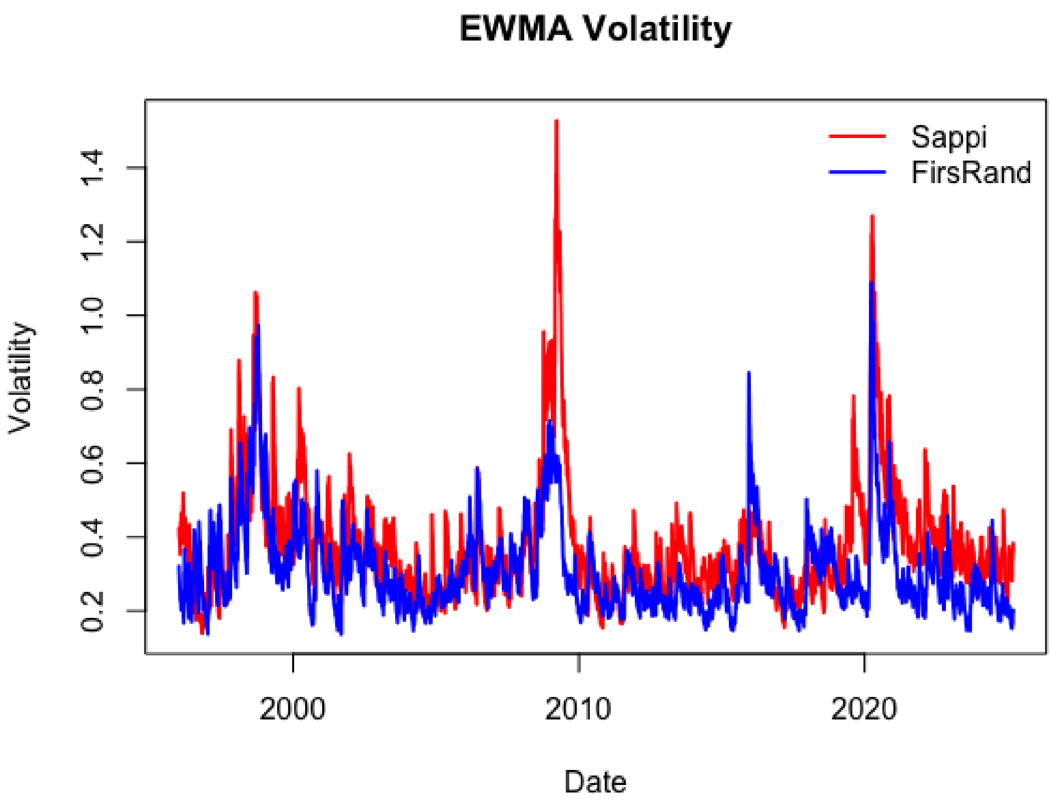

Significance Level	Unadjusted VaR	Method 1	Method 2	Method 3
0.1%	34.18%	17.71%	18.35%	18.22%
1%	18.07%	12.69%	12.97%	12.52%
5%	10.00%	8.80%	9.11%	8.60%
10%	7.12%	6.80%	6.93%	6.68%

Table 5: Comparison of Unadjusted VaR and Different Methods for Volatility Adjustment.

It is clear to see that based on current market conditions, unadjusted VaR significantly overestimates VaR. It is interesting to see that the three different values are all fairly close together. In this case, method 1 becomes redundant since the it is far more computationally expensive.

We show the adjusted portfolio returns based on the current market conditions:

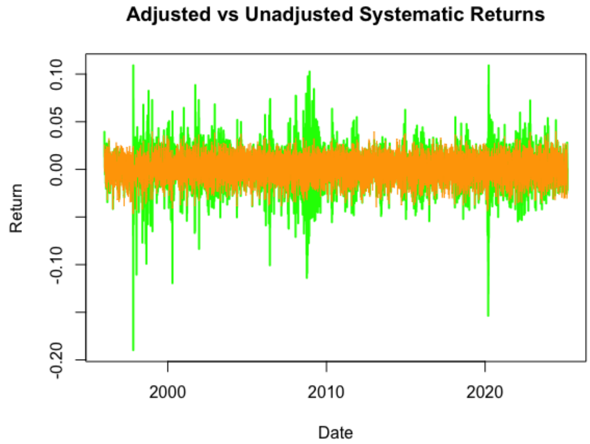

The returns have been shrunk, which is the opposite of what we saw during covid:

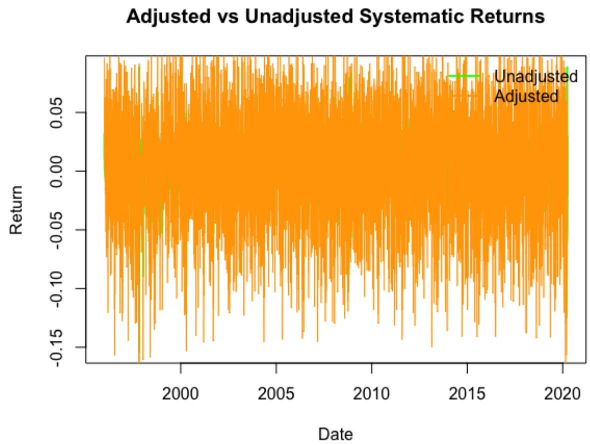

As we can see, the high volatility during the Covid-19 pandemic causes the adjusted returns to be completely inflated.

This is further evident in the VaR estimates:

	Historical VaR		Normal VaR	
Covid	**Unadjusted**	**Vol. Adjusted**	**Unadjusted**	**Vol. Adjusted**
Total VaR	17.83%	46.68%	15.16%	46.01%
Systematic VaR	13.18%	37.78%	11.02%	35.57%
Specific VaR	14.32%	30.63%	12.48%	29.35%
Post-Covid	**Unadjusted**	**Vol. Adjusted**	**Unadjusted**	**Vol. Adjusted**
Total VaR	18.07%	12.97%	15.31%	12.74%
Systematic VaR	15.32%	9.98%	13.34%	9.45%
Specific VaR	15.54%	8.87%	13.44%	8.60%

Table 6: Comparison of Covid and Post-Covid VaR under Historical and Normal Approaches.

The VaR estimates during Covid are massive once adjusted for volatility, for both normal and historical models. We see that using the current volatility estimates, the adjusted VaR is lower than the unadjusted. We also observe that the normal VaR is consistently lower, meaning that in this case the volatility adjustment actually increased the excess kurtosis. A final concerning observation, is how large the specific VaR estimates are, meaning there is almost as much risk that is not explained as there is risk that can be explained.

Conclusion

In this assignment, we have illustrated how to compute historical VaR estimates on South African data. Through this, we have explored how to adjust returns to better capture current market conditions, seen the affect sample choice has on the model and compared this methodology to normal linear VaR. The benefits of using historical VaR is that no assumption about the return distribution's form is made as well as the fact that the values used have been observed before. However, the downside is we need a lot of data to get accurate estimates.

References

Alexander, C. (2008). *Market risk analysis, quantitative methods in finance*. John Wiley & Sons.

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