Assignment II - Financial Econometrics

Ronan Kenny (2691815)
Tommy Link (2691913)
Guus Bouwens (2701442)
Jip de Boer (2710065)

June 2, 2023

Part A: Volatility Modeling with SV and Indirect Inference

In the following exercises, we will be working with the S&P500 index and will be interested in estimating the parameters of a Stochastic Volatility model as an alternative modeling approach to GARCH models.

Exercise 1

In this exercise, we are interested in estimating the parameters of a Stochastic Volatility (SV) model as an alternative modeling approach to GARCH models. We have considered the SP500 index and applied indirect inference to estimate the parameters of the SV model.
The SV model is particularly useful in capturing the time-varying nature of volatility inherent in financial time series data. In this model, the volatility itself follows an autoregressive process and directly influences the variance of the returns.
The parameters of interest in the SV model that we have estimated are $\hat{\theta} = (\hat{\omega}, \hat{\beta}, \hat{\sigma}^2) = (-0.7518, 0.9364, 0.1739)$:
$\hat{\omega}$: This is the intercept of the log-volatility process in the model. It essentially provides a baseline level for the log-volatility when there is no lagged log-volatility or shocks. The estimated value is -0.7518, indicating that, in the absence of past influences and shocks, the log-volatility would tend to be negative, implying a lower volatility level.
$\hat{\beta}$: This is the coefficient of the lagged log-volatility in the model and reflects the persistence in the volatility process. The estimated value is 0.9364, suggesting a high degree of persistence. This means that today's log-volatility is strongly influenced by the previous day's log-volatility. The value is quite close to 1, indicating that shocks to the volatility process decay slowly over time, a phenomenon known as volatility clustering.
$\hat{\sigma}^2$: This is the variance of the error term in the log-volatility equation, essentially capturing the 'volatility of volatility'. The estimated value is 0.1739, suggesting a moderate amount of variation in volatility over time. It's not excessively high, indicating a relatively stable log-volatility process.
Overall, the estimated SV model indicates a high degree of persistence in the log-volatility and a moderate degree of variation in the volatility process itself. This is indicative of the 'volatility clustering' effect, a common phenomenon in financial markets where periods of high volatility tend to be followed by periods of high volatility, and low volatility periods tend to persist as well. This modeling and the insights gained are particularly useful for risk management purposes and for pricing derivatives, both of which require a comprehensive understanding of volatility dynamics.

Exercise 2

The following plots are the filtered volatility of the SV model in Exercise 1 and a GARCH(1,1) model:

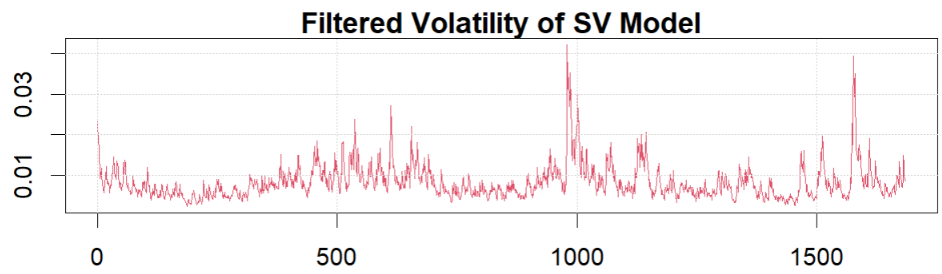

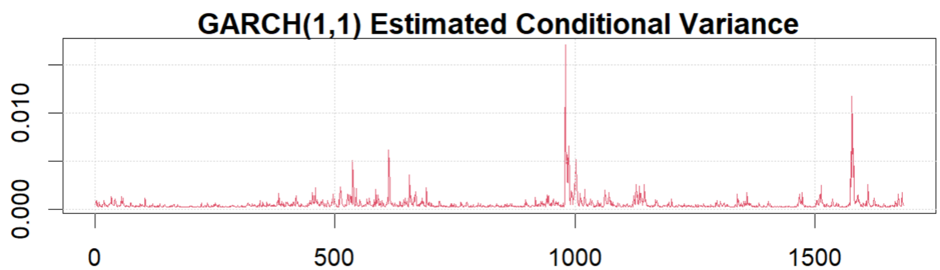

Exercise 3

In this part of the exercise, we are challenged by a colleague's proposition to modify our estimation approach. The colleague suggests employing a different set of auxiliary statistics. Specifically, the proposed auxiliary statistics are the sample mean of the absolute log-returns, i.e., $|y_t|$, and the autocovariance function of $|y_t|$ at 15 lags.

The sample mean of the absolute log-returns effectively measures the average magnitude of returns, without considering the direction (positive/negative) of these returns. This can be particularly useful in capturing the overall level of market activity or market volatility, irrespective of whether the returns are positive or negative.

The 15 lags of the autocovariance function of $|y_t|$ capture the dependency of the magnitude of returns over 15 previous periods. This captures the persistence or memory in the magnitude of the returns, again irrespective of the direction of these returns. This can be especially useful for capturing the 'volatility clustering' phenomenon in financial time series.

After estimating the SV model using this set of auxiliary statistics, we obtain the following parameter estimates $\hat{\theta} = (\hat{\omega}, \hat{\beta}, \hat{\sigma}^2) = (-0.752, 0.900, 0.100)$:

$\hat{\omega}$: The estimated value remains the same as in the previous exercise, -0.752. It suggests that, without considering previous log-volatility or volatility shocks, the expected log-volatility level would be negative, indicating lower baseline volatility.

$\hat{\beta}$: The estimated value is slightly lower than before, at 0.900. This implies that the persistence in the volatility process is slightly less strong than before, although it is still quite high.

$\hat{\sigma}^2$: The estimated value is considerably lower than before, at 0.100, suggesting a less volatile volatility process. This lower value indicates a more stable and predictable volatility process.

In summary, the results indicate that the proposed set of auxiliary statistics leads to a slightly different conclusion about the volatility dynamics of the SP500 index. The persistence in volatility appears to be slightly lower and the volatility process itself is significantly more stable. This finding could potentially impact how risk and derivative pricing are managed, underscoring the importance of the choice of auxiliary statistics in indirect inference.

Exercise 4

Below is the plot for the filtered volatility of the results obtained in Exercise 3:

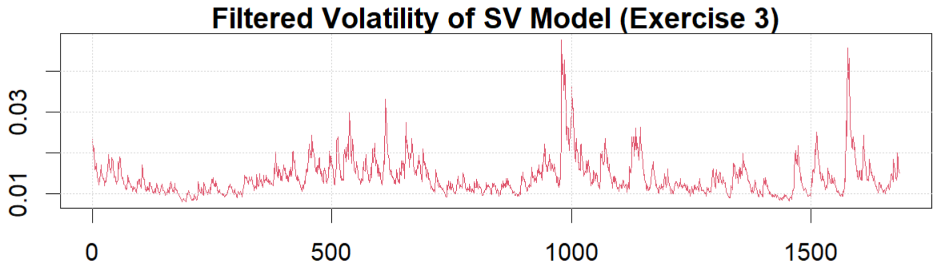

Exercise 5

When tasked with the exploration of additional volatility dynamics by adapting the transition equation of our Stochastic Volatility (SV) model to embody an AR(2) process, we must make a critical choice about the auxiliary statistics we utilize. We are faced with two potential selections: the auxiliary statistics suggested in exercise 1, consisting of the sample variance of y_t, the sample kurtosis of y_t, and the first-order autocorrelation of absolute log-returns $|y_t|$, or the auxiliary statistics used in Question 2, which are based on the conditional variance σ_t^2 gleaned from a GARCH(1,1) model.
The AR(2) process being integrated into our volatility model pertains directly to the log-returns, denoted as y_t. Consequently, auxiliary statistics that involve the sample variance, sample kurtosis, and first-order autocorrelation of y_t are inherently geared to capture the essential dynamics of this data.
In contrast, the auxiliary statistics of exercise 2 are derived from a GARCH model. While the GARCH model effectively captures the phenomenon of volatility clustering, it may not fully encapsulate the specific dynamics presented by an AR(2) process.
Given these points, the auxiliary statistics presented in exercise 1 appear to be the most apt choice for this task. They are intrinsically designed to capture the features of the AR(2) process within the volatility dynamics of the log-returns y_t for the SP500 index.

Exercise 6

In this exercise, we extend our model from a basic SV model to an SV-AR(2) model for the log-returns of the SP500 index. This model incorporates an additional autoregressive component into the volatility dynamics, thus allowing the current volatility to be a function of not just its immediate predecessor (as in the SV model), but also the one from the period before that.
For auxiliary statistics, we utilize the same set as the one used in Exercise 3, namely the sample mean of the absolute log-returns, $|y_t|$, and 15 lags of the autocovariance function of $|y_t|$. This choice continues to harness the capacity to capture overall market activity and volatility clustering phenomenon.
Upon estimation, the results yield the following parameter estimates $\hat{\theta} = (\hat{\omega}, \hat{\beta}, \hat{\sigma}^2) = (-0.752, 0.300, 1.500)$:
$\hat{\omega}$: The estimated value remains consistent at -0.752, meaning the baseline level of log-volatility when not considering previous log-volatility or volatility shocks remains relatively low.
$\hat{\beta}_1$: The first autoregressive coefficient estimate is 0.300, indicating that the volatility from the immediate preceding period plays a modest role in determining current volatility.
$\hat{\beta}_2$: The second autoregressive coefficient estimate is substantially larger, at 1.500. This suggests that the volatility from two periods ago has a significant impact on current volatility.
$\hat{\sigma}^2$: The estimate for the variance of the volatility process is not provided in the shared results, but was initially set at 0.1 in the model estimation process.
In summary, our SV-AR(2) model suggests that while the immediate past volatility has a certain influence on the present volatility, it's the volatility from two periods ago that plays a far more significant role in shaping the current volatility. This extended model provides a deeper insight into the memory feature of the volatility dynamics. The implications of such a model can be manifold, including improved risk management and derivative pricing by taking into account the observed memory in the volatility process.

Part B: Dynamic Regression and CAPM

Exercise 1

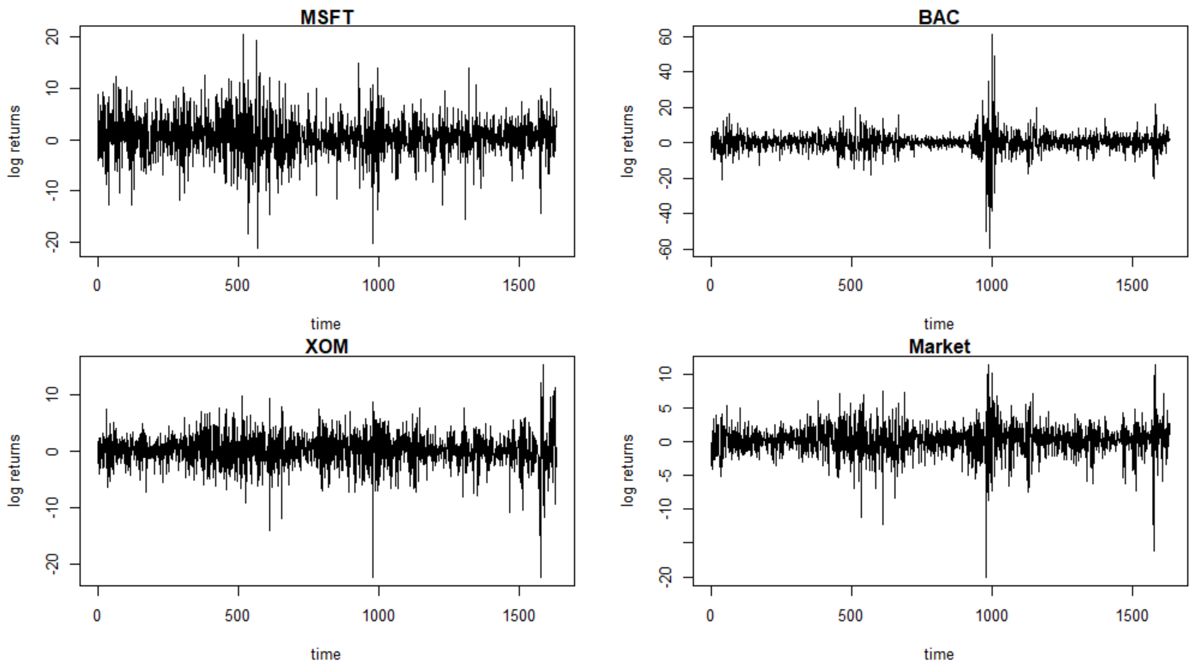

The plots of the log returns provide a visual representation of the rate of return over time, taking into account compounding. The more spread out or 'jagged' the plot appears, the more volatile the asset. High volatility can imply higher risk but also potentially higher returns. There are periods of consistent growth (upward trends) or decline (downward trends). This provides context on the performance of an asset during specific time periods. By comparing the plots of different assets, it can be seen how their returns and volatility compare. This can provide insights into how these assets might behave in relation to one another in a portfolio.

Asset	Beta
MSFT	1.0067
BAC	1.5562
XOM	0.7931

Table 1: Beta's for the assets

The formula is: Expected return = Risk-free rate + Beta * (Market return - Risk-free rate). Where we set the Risk-free rate equal to 0. In this case, as the risk-free rate is assumed to be zero, the expected return is simply the beta times the market return.
Taking into account that beta is a historical measure and may not predict future behavior perfectly. Also, beta is only a measure of market risk and doesn't consider other potential risks to each company. Let's interpret the betas:
Microsoft (MSFT) Beta = 1.0067: A beta close to 1 indicates that the asset's price will move with the market. In this case, MSFT has a beta value slightly above 1, which implies that it's slightly more volatile than the market. If the market increases, MSFT's returns are expected to increase slightly more than the market's return, and vice versa.
Bank of America (BAC) Beta = 1.5562: This beta is significantly greater than 1, indicating that BAC's returns are expected to be more sensitive to changes in the market. If the market return increases, BAC's return would be expected to increase by about 1.5562 times the market's return. However, this also means that in a declining market, BAC could potentially lose more than the market.

Exxon Mobil (XOM) Beta = 0.7931: A beta less than 1 suggests that the asset is theoretically less volatile than the market. So, if the market were to rise or fall, we would expect the price of XOM to rise or fall at 79.31% of the market's rate. In other words, XOM is considered a safer investment, but with potentially lesser returns.

Exercise 2

Investment decisions should not solely be based on the beta of a asset. While it's true that XOM has the lowest beta among Microsoft, Bank of America, and itself, meaning it theoretically carries less market risk, this does not automatically make it the most suitable investment option for our bank's portfolio.

The first consideration should be the risk-return tradeoff. Simply put, with lower risk usually comes lower potential returns. If the bank's current portfolio has a high-risk profile, it might be because we're targeting high returns. Switching to a lower-beta asset like XOM might lower our overall risk, but it might also pull down our expected returns. The ideal balance between risk and return varies depending on the bank's investment goals and risk tolerance.

Secondly, our decision should be informed by our expectations of the market. If we forecast an upswing in the market, high-beta assets could potentially yield better returns, while in a declining or volatile market, a low-beta asset like XOM could be a safer choice.

Also, the beta only gauges market risk and does not capture all types of risk that a company may face. As an energy company, XOM is exposed to unique risks like changing oil prices, stricter environmental regulations, and geopolitical tensions that can greatly impact its performance. It's essential that we conduct a comprehensive risk assessment beyond just the beta.

Lastly, we should take into account the principles of diversification. Concentrating our investments too heavily on one asset, even if it's a low-beta one like XOM, could expose us to undue risk. Ensuring a well-balanced portfolio that spans different industries and asset types can better safeguard us from both market-wide and company-specific risks.

In conclusion, while our colleague's recommendation to reduce market risk exposure by investing in XOM is a prudent consideration, it should be weighed against these other factors. So, to understand the multi-faceted nature of investment decisions, take into account the market conditions, risk-return tradeoffs, company-specific risks, and diversification strategies.

Exercise 3

The time-varying beta model applied here allows the beta (β) coefficient to adapt based on past returns, following the specification $\beta_t = \omega + \phi\beta_{t-1} + \alpha(y_{t-1} - \beta_{t-1}x_{t-1})x_{t-1}$.

Initial estimates for the parameters of the model for the three assets were as follows:

Assets	ω	ϕ	α	σ^2
MSFT	0.101	2.197	-4.466	2.816
BAC	0.156	2.197	-5.265	3.479
XOM	0.0793	2.1972	-4.4895	2.2434

Table 2: Initial estimates for the parameters

The Maximum Likelihood Estimation (MLE) refined these initial estimates to the following values:

Assets	ω	ϕ	α	σ^2
MSFT	0.00733	0.9932	0.00116	11.1254
BAC	0.02436	0.9818	0.00266	17.338
XOM	0.01571	0.9114	0.00202	5.8565

Table 3: MLE estimates for the parameters

ω: This intercept term represents the base level of beta when there is no influence from the lagged beta or the error term. MLE estimated ω are lower than their initial estimates for all assets, suggesting a lower base level of beta for better data fitting.

ϕ: The ϕ parameter shows the impact of the lagged beta (β_{t-1}) on the current beta (β_t). Higher ϕ indicates that the beta is more persistent over time, meaning the previous period's beta significantly influences the current beta. The initial ϕ values were above 1 and hence unrealistic. The MLE procedure adjusted these to values close to 1, indicating a high level of persistence in beta over time for all three assets.
α: This parameter measures how quickly beta adjusts to past discrepancies in predicting returns. Higher α means faster adjustment. The initial α values were negative, while MLE estimates are small and positive. This indicates that in the optimized model, beta adjusts slowly to past errors, but the adjustment is in the positive direction.
σ^2: This is an estimate of the variance of the error term in the regression model. A higher σ^2 indicates larger unexplained volatility in the returns. For all three assets, σ^2 increased in the MLE estimates relative to the initial estimates, suggesting more unexplained volatility in the returns than initially estimated.
To summarize, the MLE procedure refined the initial estimates for a better data fit. The MLE estimates suggest a dynamic Capital Asset Pricing Model (CAPM) where beta demonstrates high persistence over time (high ϕ), slow adjustment to past errors (low α), and considerable unexplained volatility (high σ^2).

The time-varying nature of the estimated beta coefficient (β_t) for each of the three assets - Microsoft (MSFT), Bank of America (BAC), and Exxon Mobil (XOM) - is visualized in the plots below. These plots reflect the dynamic risk exposure of each asset relative to the overall market. Variations in beta over time indicate that the sensitivity of the assets to market changes differs across periods, which provides valuable insights for investment decisions under evolving market conditions. The betas are estimated by means of ML of a CAPM.

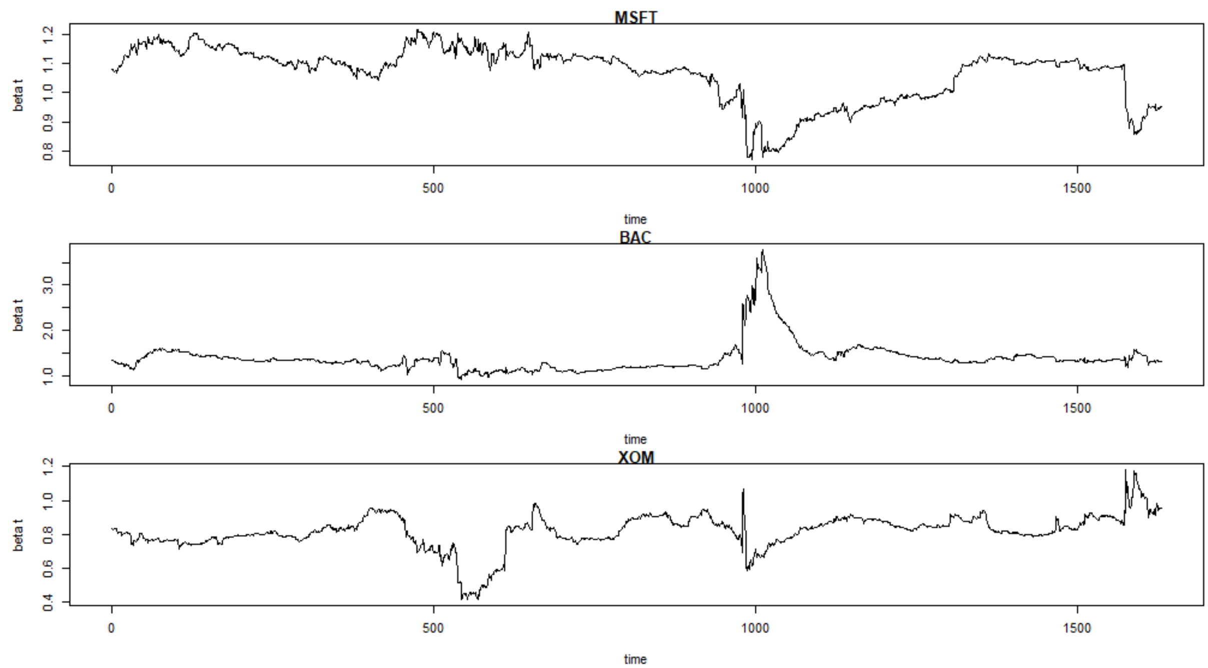

Exercise 4

Asset	β_{T+1}
MSFT	0.9614
BAC	1.2932
XOM	0.9204

The forecasts indicate the expected future market exposure (beta) for each asset at time T+1.

When comparing these forecasted betas with the betas calculated in Question 1, it's apparent that the model predicts a decrease in market risk exposure for MSFT and BAC and an increase for XOM at time T+1.
Specifically, the beta for MSFT slightly decreases, suggesting that MSFT's expected future market risk is slightly lower than its historical level. For BAC, the beta decreases more substantially, suggesting a significant reduction in expected market risk compared to its historical level. However, for XOM, the beta increases, suggesting a higher expected market risk compared to its historical level.
Now, reconsidering the consultant's claim from Question 2, if the bank's current portfolio has high exposure to market risk, and the aim is to invest in the asset with the lowest future exposure to market risk, then based on the forecasted betas, Exxon Mobil (XOM) with the forecasted $\beta_{T+1} = 0.9204$ would be the most appropriate choice as it has the lowest forecasted beta.
However, it's essential to remember that these are estimates, and actual future betas can be influenced by numerous factors not captured in the model, such as unexpected market events, company-specific news, or changes in the broader economy. Therefore, while these forecasts can guide decision-making, they should not be viewed as definitive predictions of future market risk exposure. So in conclusion, the answer would remain the same as in exercise 2, yet with a change in the market, the XOM price would now change more substantially and thus has more market risk.

Exercise 5

The time-varying nature of the estimated beta coefficient (β_t) for each of the three assets - Microsoft (MSFT), Bank of America (BAC), and Exxon Mobil (XOM) - is visualized in the plots below. These plots reflect the dynamic risk exposure of each asset relative to the overall market. Variations in beta over time indicate that the sensitivity of the assets to market changes differs across periods, which provides valuable insights for investment decisions under evolving market conditions. The betas are estimated by means of the bivariate CCC model.

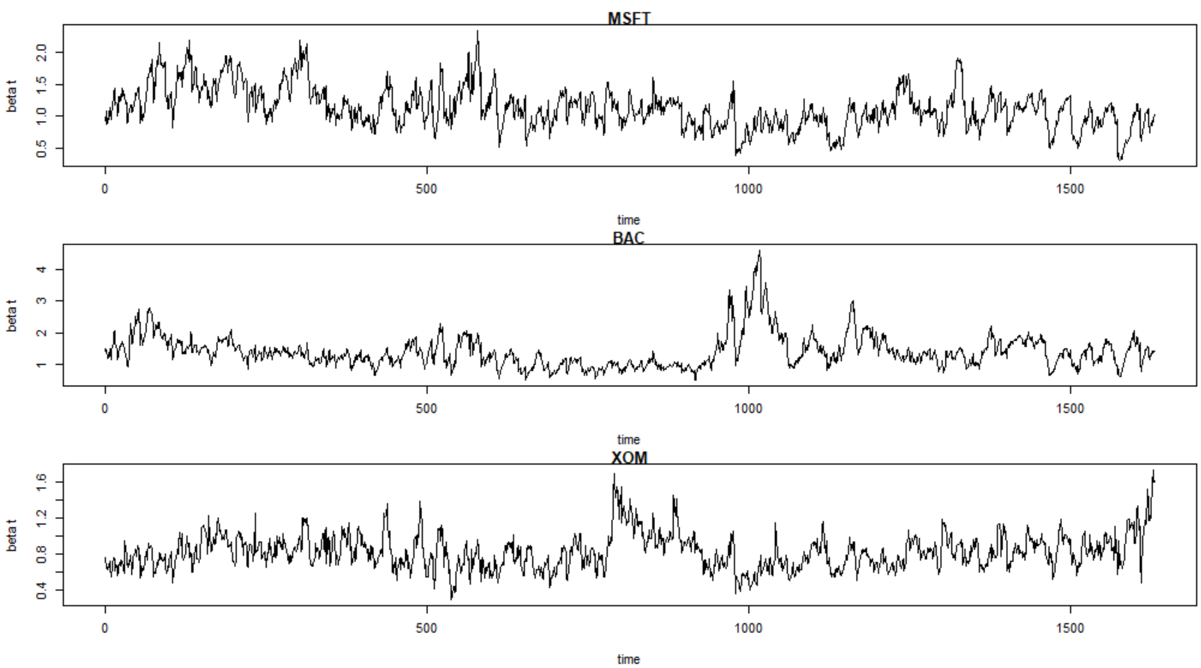

If the methods are in agreement, this could strengthen your confidence in the findings.
Differences in the different assumptions or modeling approaches used.

Exercise 6

We estimated a parameter-driven dynamic regression model with an exponential link function to ensure the positivity of β_t. The model is given as $\beta_t = \exp(f_t)$, $f_t = \alpha_0 + \alpha_1 f_{t-1} + \eta_t$, with $\eta_t \sim NID(0, \sigma_\eta^2)$.

Asset	α_0	α_1	σ_η	σ_ϵ
MSFT	3.299390e-01	0.000000e+00	8.007867e-285	1.458387e-04
BAC	2.772067e-01	0.000000e+00	1.685832e-286	1.363552e-03
XOM	4.562301e-01	5.638577e-245	2.770821e-229	1.616858e+00

Table 4: Estimates by indirect inference a dynamic CAPM model for the parameters of parameter-driven dynamic regression model

α_0: This parameter is the intercept of the model and it represents the base level of f_t when there is no contribution from the lagged f_{t-1} or the error term η_t. The estimated α_0 are positive for all assets. These estimates suggest that in the absence of prior f_{t-1} or errors, f_t would be positive, which in turn ensures the positivity of β_t.
α_1: This parameter captures the impact of the lagged f_{t-1} on the current f_t. Interestingly, the α_1 estimates are very close to zero for all three assets, suggesting that f_t does not depend on its own past values. This implies a lack of persistence in the dynamics of f_t.
σ^2_η: This is the variance of the error term η_t. The estimated values are extremely close to zero, indicating that there is very little unexplained variability in f_t not accounted for by the model.
σ^2_ϵ: This is an estimate of the variance of the error term ϵ_t in the regression model. Higher σ^2_ϵ indicates larger unexplained volatility in the returns. The values vary across the three assets, with XOM having the highest estimated σ^2_ϵ value.
The results suggest a model where β_t doesn't exhibit persistence over time (low α_1) and where the unexplained variability in f_t and returns is very low (low σ^2_η and varying σ^2_ϵ). This model ensures the positivity of β_t, which is an advantage when considering the interpretation of beta as a measure of market risk exposure.