Jollibee Foods Corporation Stock Price Forecasting with ARMA-GARCH

David Demitri Africa, Mark Kevin A. Ong Yiu, Carlo Gabriel M. Pastor

Abstract

In this paper, we model the log returns of the daily closing stock price of Jollibee Foods Corporation (JFC) from January 5, 2009 to February 7, 2023 using ARMA-GARCH. We test forecasts using the daily stock price for the next 100 trading days (February 8, 2023 to July 7, 2023), with the best model performing top-of-class on windowed forecasts and competitively on one-shot forecasts.

1 Introduction

Among the countless companies listed on the Philippine Stock Exchange (PSE), Jollibee Foods Corporation (JFC) stands out as a multinational food service corporation with a remarkable presence and market valuation. Having a multi-billion dollar market capitalization, it has historically been considered a "blue chip" stock with a track record of consistent growth and high liquidity. These characteristics have made it an attractive subject for stock price forecasting, with several studies [2, 3, 4] testing various model approaches on historical JFC stock price data.

In this section, we will briefly outline some introductory concepts and points of reference that will be needed to understand the rest of this study.

1.1 Statement of the problem

The Autoregressive Moving Average and Generalized Autoregressive Conditional Heteroskedasticity (ARMA-GARCH) is a powerful tool for modelling time series data with volatility clustering, which is a behavior typically observed in stock price log returns. However, a detailed theoretical discussion on the ARMA-GARCH modelling approach is both outside the scope of this study and redundant.

Rather, the primary focus of this study is to asses the suitability of ARMA-GARCH modelling approach for forecasting JFC stock prices. Using historical JFC stock price data, we aim to evaluate the forecasting accuracy of the ARMA-GARCH methodology by comparing it against several models reported in previous literature. Specifically, we seek to answer two central questions:

1. Is the ARMA-GARCH modelling approach appropriate for forecasting JFC stock price?
2. Does the ARMA-GARCH method outperform existing forecasting models documented in literature?

1.2 Previous literature

Several researchers have previously forecasted the stock price of JFC using various modeling approaches. In order to evaluate the effectiveness of our modelling approach on out-of-sample data, we compare our selection of parameter and model to the selected parameters and models of the following past forecasting attempts on JFC stock price data.

Submission for MATH 271.2 Advanced Statistical Methods, AY 2023-2024 Intersession

1.2.1 Forecasting Philippine stock market prices using ARIMA-GARCH models

Guiao [2] explored the viability of applying ARIMA, ARCH, and GARCH time series models to forecast the price of JFC stock. Daily closing prices of the stock were collected over trading days from January 2, 2014 to March 22, 2019, resulting in a total of 1,271 observations.

To build the time series models, the training data was transformed into log returns, and appropriate orders for an ARIMA model were determined from the analysis of the autocorrelation function (ACF) and the partial autocorrelation function (PACF). After evaluating various models and selecting the most appropriate based on Akaike Information Criteria (AIC) and Bayesian Information Criteria (BIC) scores, an ARMA$(1, 2)$-ARCH(1) model was chosen. This model was then used to perform one-hundred 1-step ahead forecast of JFC stock prices on the testing set, achieving a mean absolute percentage error (MAPE) of 0.77%.

1.2.2 Forecasting in business research using the ARIMA Box-Jenkins methodology

Malaya [3] applied an ARIMA modeling procedure over the daily closing prices of JFC stock from January 4, 1999 to September 30, 1999, resulting in a total of 190 observations. The first difference of the data was taken, and an AR(1) model was subsequently was fitted over the first 187 observations. The model was then evaluated on the last 3 observations, yielding a MAPE of 0.5%.

1.2.3 Bayesian estimation of A-PARCH model: an application to Jollibee Food Corporation stock market

Oliveros and Supe [4] attempted to model market and financial trends with a specific focus on modeling financial volatility. They addressed the limitations of the standard ARMA-GARCH models, which fail to account for asymmetric effects of positive and negative shocks, by employing the Asymmetric Power ARCH (A-PARCH) model (see section 1.3).

To test the applicability of the A-PARCH model on financial data, they similarly made use of daily closing prices of JFC stock from May 3, 2012 to April 23, 2019, resulting in a total of 1,692 observations. Similar to [2], this was transformed into log returns, and an ARMA$(1, 1)$-A-PARCH$(1, 1)$ model was selected as the best model based on ACF, PACF, and AIC. This was then used to compute a 1-step ahead forecast of the JFC stock market returns, yielding a mean squared error (MSE) of 1.7511%.

1.3 Asymmetric Power ARCH (A-PARCH)

The A-PARCH model is an extension of the GARCH model, which changes the conditional variance equation to accommodate asymmetric responses to positive and negative shocks. The equation is given by

$$\sigma_t^\delta = \omega + \sum_{i=1}^{m} \alpha_i (|y_{t-i}| - \gamma_i y_{t-i})^\delta + \sum_{j=1}^{s} \beta_j (\sigma_{t-j})^\delta$$

where γ_i determines the sign of the past shocks that have asymmetric effects on the volatility, and δ controls the power transformation and determines the nature of the asymmetric response. When $\delta > 0$, negative shocks have a larger impact on volatility than positive shocks. Conversely, when $\delta < 0$, positive shocks have a more substantial effect on volatility than negative shocks, leading to a positive leverage effect.

2 Methods

In this section, we will outline the methodology used to select and evaluate models, as well as the underlying definitions of important concepts.

2.1 Description of the data

We begin by pulling the daily close prices of JFC stock from January 5, 2009 to July 7, 2023 from the Wall Street Journal [1], resulting in a total of 3,539 observations. The data is assumed to have no trading activity on days with gaps, signifying nontrading days like holidays or weekends. Since the series being forecasted is the price, which is contiguous from the close of one trading day to the beginning of the next trading day, this assumption is reasonably sound to make. We reserve the last 100 observations (February 8, 2022 to July 7, 2022) as our testing set for model evaluation, and fit our models on the remaining data (see Fig. 1)

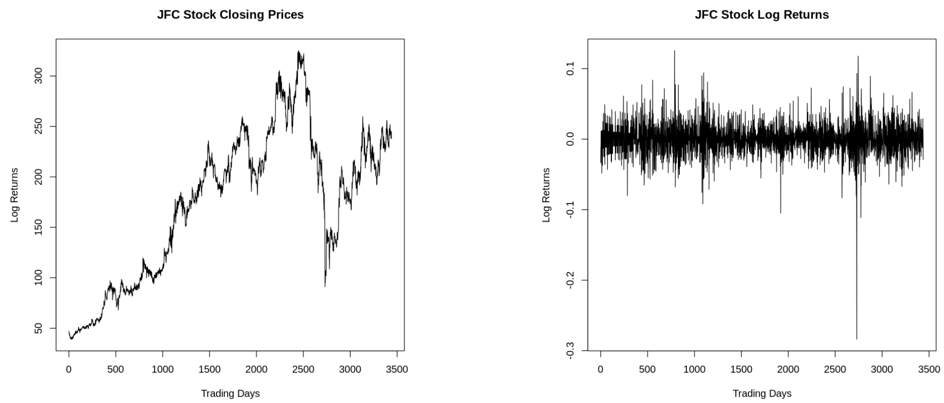

(a) Daily JFC close price

(b) Log returns of the daily JFC close price

Figure 1: Plots of daily JFC close price data before and after data transformation from January 5, 2009 to February 7, 2023

As is conventional with forecasting stock price data, the data was transformed to retrieve the log returns of the daily close price. Let r_t be the log returns of the daily close price of JFC stock (which we will call P_t). Then, r_t is defined as

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) = \ln(P_t) - \ln(P_{t-1}).$$

Following this, the forecasted price $\hat{P}_{t+1}$ given forecasted log return $\hat{r}_{t+1}$, is given by

$$\hat{P}_{t+1} = e^{\hat{r}_{t+1}} \cdot P_t.$$

To employ the ARMA-GARCH methodology, the data must be stationary and serially correlated. Moreover, nonseasonality must also be checked to guarantee the effectiveness of modeling. In this paper, we will use a 5% level of significance for our statistical tests.

To check stationarity, the Augmented Dickey-Fuller (ADF) test was used, with the null hypothesis being that the time series contains a unit root and is nonstationary. Performing the test on r_t yields a p-value of 0.01, indicating that the null hypothesis can be rejected. Hence, r_t is stationary.

To check autocorrelation, the Ljung-Box test was used, with the null hypothesis being that the data exhibits no serial correlation. Performing the test on r_t for lags 1 to 30 shows significant p-values ($p < 0.05$), but only for lags greater than or equal to four. The significance of lag 4 can be observed in Fig. 2.

Finally, to check seasonality, several methods were employed to verify the data is nonseasonal. First, ACF and PACF of the data was examined, and a visual inspection of the plots show no seasonality (see Fig. 2). Second, using `decompose()` from the `forecast` package on the log return data shows no clear seasonal trend on the seasonal component of the plot (see Fig. 4). Finally, using `isSeasonal()` from the `seastests` package returns `FALSE` for lags 1 to 300 (except for 213, but we exercise our judgement that this is not significant nor interpretable).

2.2 Grid search

As mentioned above, autocorrelation only emerges in the log return data at lags 4 and higher. This means that the process of selecting an ARMA model by visual inspection is not so straightforward. So, grid search was employed on ARMA parameters p and q to exhaust all combinations $(p, q) \in [0, 5] \times [0, 5]$. Grid search is a heuristic that involves defining a set of parameters (in this case, p

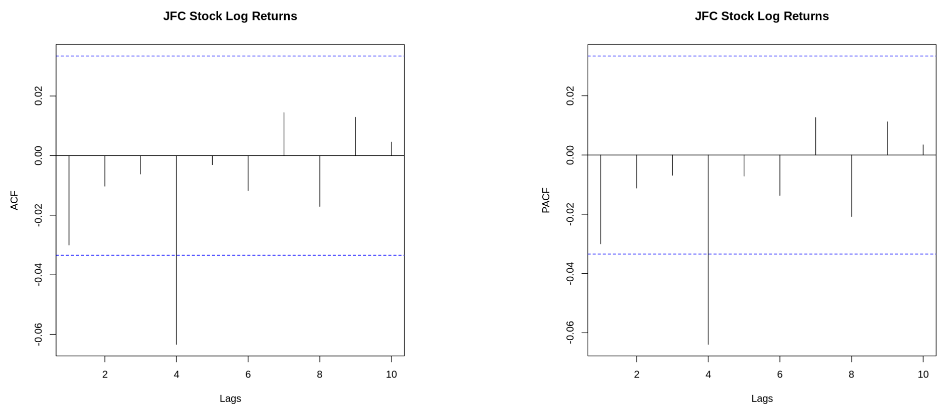

(a) ACF of daily JFC log returns

(b) PACF of daily JFC log returns

Figure 2: Diagnostic plots of autocorrelation in daily JFC log returns

and q) and their corresponding values to be considered in the search, after which ARMA models are generated based on the given parameters and data. The models were then compared based on their AIC, with the best model having the lowest AIC.

After performing grid-search, we found that the AIC scores were too close together to meaningfully differentiate (see Table 3). Thus, to select a best model, all ARMA models with insignificant parameters, i.e. coefficients that fail to be significantly different from 0, were discarded. This left three candidate models – AR(4), MA(4), and ARMA$(1,1)$.

The residuals and squared residuals of the remaining models were checked for autocorrelation (using the ACF, PACF, and the Ljung-Box test), with the requirement for GARCH modeling being that the residuals are uncorrelated, but the squared residuals are autocorrelated. After testing, all three remaining models fit the criteria for GARCH modeling.

Inspecting the ACF and PACF plots of the squared residuals for all models provides something more interpretable. Fig. 3 shows the plots for the squared residuals of the ARMA(1,1) model, though the plots for all the models (see Fig. 5 and Fig. 6) were visually similar. Here, the PACF plot cuts off after lag 3 while the ACF plot tails off, suggesting GARCH parameters $m \leq 3$, $s = 0$.

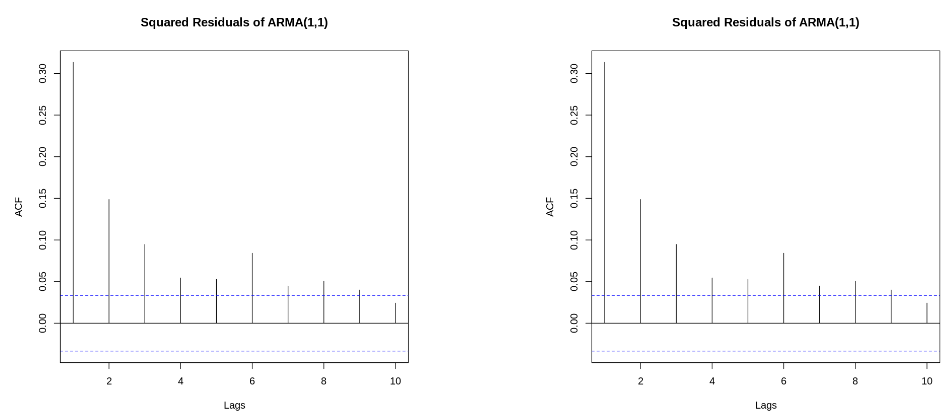

(a) ACF of squared residuals of ARMA$(1,1)$

(b) PACF of squared residuals of ARMA$(1,1)$

Figure 3: Diagnostic plots of autocorrelation in the squared residuals of ARMA$(1,1)$

Grid search was then performed again on the new parameter space, checking each set of parameters for ARMA $(p,q) \in \{(0,4),(4,0),(1,1)\}$ with each set of parameters for GARCH $(m,s) \in \{(1,0),(2,0),(3,0)\}$ using joint estimation and the `garchFit()` function. This gave us nine models to test, with their resulting AICs found in Table 4 in the Appendix. Again, the AIC scores were too similar to meaningfully suggest a best model, thus we have opted to evaluate our models using MSE and MAPE instead.

2.3 Evaluation

In order to assess the accuracy of a model's predictions, two metrics were used – MAPE and MSE. Given n datapoints Y_i and n corresponding predictions $\hat{Y}_i$, the mean absolute percentage error (MAPE) is given by

$$\text{MAPE} = \frac{100}{n} \sum_{i=1}^{n} \left| \frac{Y_i - \hat{Y}_i}{Y_i} \right|$$

and the mean squared error (MSE) is given by

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^{n} \left(Y_i - \hat{Y}_i \right)^2 .$$

To simulate real-world forecasting scenarios, we evaluate the model's performance on the excluded period from February 8, 2023 to July 7, 2023 using two methods – one-shot forecasting and windowed forecasting. For the former method, we calculate the MAPE and MSE for the model's predictions made 100 days into the future without shifting the dataset or updating the model. For the latter method, we evaluate the model's performance in a more dynamic setting, by iteratively shifting the dataset day by day and generating repeated one-day ahead forecasts. We compute the MAPE and MSE over all windows, enabling us to observe the model's adaptability and responsiveness to changing market conditions.

2.4 Checking validity of residuals

Afterwards, the standardized residuals $\frac{\hat{Y}_t^2}{\hat{\sigma}_t^2}$ of each resulting ARMA-GARCH model were checked. All standardized and squared standardized residuals were found to be uncorrelated (Ljung-Box test p-value > 0.05), stationary (ADF test p-value < 0.05), but not normal (Jarque-Bera test p-value < 0.05). This signifies that there are no obvious remaining trends in the mean or volatility of the data. All p-values can be found in Table 5 in the Appendix.

While the standardized residuals are not normally distributed, we compared it with a generated standardized t-distribution using the Kolmogorov-Smirnov test, whose null hypothesis is that the two random vectors in its input come from the same distribution. We found that the standardized residuals of all models have a standardized t-distribution with 4 degrees of freedom (p-value > 0.05). The implications of this may be the subject of future research.

3 Results and discussion

In this section, we present the results of our methodology as well as discuss and compare them to the approach taken in previous studies. Due to some errors in the `fGarch` library, `garchFit()` was unable to do forecasts past the second day for AR(4)-ARCH(m) models.

The best model for each forecasting regime is highlighted in red, with ARMA$(0, 4)$-ARCH(1) performing best under the one-shot forecasting regime and ARMA$(1, 1)$-ARCH(2) performing best under the windowed forecasting regime. The full model equations for both models can be found in section A.1 of the Appendix. Summary statistics for both models can also be found in Fig. 6 and 7 in the Appendix.

3.1 Comparison to previous results

To compare our best models with those from past literature, we conduct a similar evaluation (one-shot and windowed) by training models using the specified parameters and architectures from literature on our dataset.

Table 1: ARMA-GARCH MAPE and MSE results on one-shot and windowed forecasting regimes

Model	One-Shot		Windowed	
	MAPE	MSE	MAPE	MSE
AR(4)-ARCH(1)	NA	NA	0.0215	44.0609
AR(4)-ARCH(2)	NA	NA	0.0215	43.9014
ARMA(4)-ARCH(3)	NA	NA	0.0215	43.8047
MA(4)-ARCH(1)	0.0665	289.1833	0.0214	43.7261
MA(4)-ARCH(2)	0.0694	311.1972	0.0213	43.4200
MA(4)-ARCH(3)	0.0682	302.1921	0.0214	43.4137
ARMA(1,1)-ARCH(1)	0.0671	293.9020	0.0214	43.6022
ARMA(1,1)-ARCH(2)	0.0695	312.0176	0.0213	43.3904
ARMA(1,1)-ARCH(3)	0.0680	300.5502	0.0213	43.4120

The table containing the results of the above can be seen below, with our two best models in the first two rows and the models of past studies in the succeeding rows.

Table 2: Comparison between literature and best model results on one-shot and windowed forecasting

Model	One-Shot		Windowed	
	MAPE	MSE	MAPE	MSE
MA(4)-ARCH(1)	0.0665	289.1833	0.0214	43.7261
ARMA(1,1)-ARCH(2)	0.0695	312.0176	0.0213	43.3904
ARMA(1,2)-ARCH(1)	0.0680	300.5502	0.0214	43.7729
AR(1)	0.0617	253.5806	0.0219	45.6498
ARMA(1,1)-A-PARCH(1,1)	0.0611	250.9817	0.0214	43.9286

Among the models tested, both the AR(1) and ARMA(1,1)-A-PARCH(1,1) models from the literature outperformed showed our best model in one-shot predictions. An interesting observation was the presence of an insignificant term in the MA(4)-ARCH(1) model, which might have contributed to its relatively lower performance. Nonetheless, the ARMA-GARCH models still provide somewhat competitive results to the literature.

However, when considering windowed forecasting, the ARMA(1,1)-ARCH(2) model outperformed the others. This is significant, since windowed forecasting more closely mimics the set-up of a day-trader, where the user gains information about trades as the day closes, allowing them to update their model and model forecasts.

In addition, it is worth noting that the training and forecasting processes for the models were time-consuming, with the A-PARCH model requiring the longest computational time due to having the most parameters.

3.2 Conclusion and recommendations

In this study, we found that the ARMA-GARCH modelling approach is appropriate for forecasting JFC stock price given its characteristics, namely, the stationarity and autocorrelation of the log returns and the autocorrelation of the squared residuals. After evaluating the performance of several models in forecasting JFC stock price, we found that the MA(4)-ARCH(1) and ARMA$(1,1)$-ARCH(2) models achieved the best results for one-shot and windowed forecasting respectively. Finally, although the AR(1) and ARMA$(1,1)$-A-PARCH$(1,1)$ models from previous literature outperformed our models in one-shot forecasting, our models generally demonstrated comparable results while also performing better in windowed forecasting.

As a potential avenue for future research, incorporating multivariate forecasting may prove beneficial, as it would allow us to examine correlations with other financial instruments or stocks, providing a more comprehensive analysis of market dynamics. It may also be beneficial to take steps to account for the effect of the pandemic on JFC stock price, such as by fitting separate models that exclude trading dates during the pandemic.

References

[1] Factset. Jollibee foods corp. stock price & news. *Wall Street Journal*, 2023.

[2] J. E. Guiao. Forecasting philippine stock market prices using arima-garch models. 2019.

[3] M. Malaya. Forecasting in business research using the arima box-jenkins methodology. *DLSU Business and Economics Review*, 12, 2001.

[4] S. M. L. Oliveros and A. P. Supe. Bayesian estimation of a-parch model: An application to jollibee food corporation stock market. 2019.

A Appendix

A.1 Model equations

The model equation for ARMA(0,4)-ARCH(1) is

$$\begin{cases} r_t = 0.0005 + Y_t - 0.0817Y_{t-1} - 0.0400Y_{t-2} - \\ \qquad 0.0310Y_{t-3} - 0.0605Y_{t-4} \\ Y_t = \sigma_t \varepsilon_t \\ \sigma_t^2 = 0.0003 + 0.2899Y_{t-1}^2 \end{cases}$$

After removing non-significant terms (based on the difference between the coefficient estimate and the standard error), the model equation can be written a

$$\begin{cases} r_t = Y_t - 0.0817Y_{t-1} - 0.0400Y_{t-2} - 0.0605Y_{t-4} \\ Y_t = \sigma_t \varepsilon_t \\ \sigma_t^2 = 0.0003 + 0.2899Y_{t-1}^2 \end{cases}$$

On the other hand, the model equation for ARMA(1,1)-ARCH(2) is

$$\begin{cases} r_t = 0.0001 + 0.6789r_{t-1} + Y_t - 0.7634Y_{t-1} \\ Y_t = \sigma_t \varepsilon_t \\ \sigma_t^2 = 0.0003 + 0.2603Y_{t-1}^2 + 0.0901Y_{t-2}^2 \end{cases}$$

A.2 Figures and tables

A.2.1 Figures

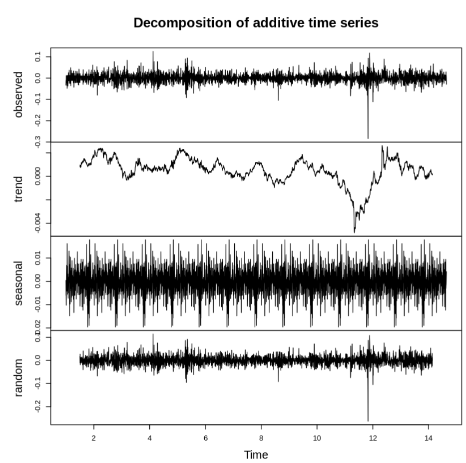

Figure 4: Seasonal decomposition of r_t using decompose()

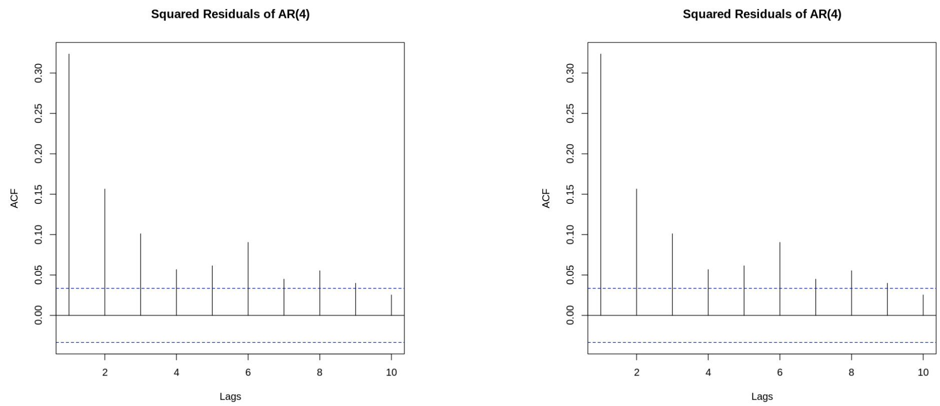

(a) ACF of squared residuals

(b) PACF of squared residuals

Figure 5: Diagnostic plots of autocorrelation in the squared residuals

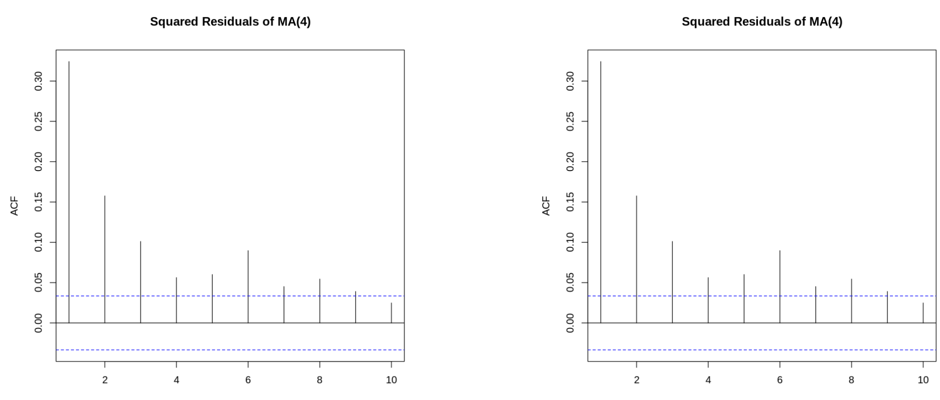

(a) ACF of squared residuals

(b) PACF of squared residuals

Figure 6: Diagnostic plots of autocorrelation in the squared residuals

A.2.2 Tables

Table 3: Grid search AIC values for ARMA(p,q) models

		p					
		0	1	2	3	4	5
q	0	-16852.15	-16854.87	-16852.76	-16852.76	-16864.65	-16861.95
	1	-16854.74	-16860.10	-16858.55	-16853.92	-16862.60	-16860.62
	2	-16852.53	-16858.57	-16856.70	-16854.08	-16860.85	-16860.63
	3	-16852.54	-16853.89	-16857.40	-16854.46	-16859.58	-16858.79
	4	-16863.87	-16862.08	-16860.59	-16859.09	-16858.18	-16857.04
	5	-16861.21	-16859.01	-16860.31	-16858.40	-16856.87	-16854.09

Table 4: Grid search AIC values for selected ARMA-GARCH models

	ARCH(1)	ARCH(2)	ARCH(3)
AR(4)	-5.054435	-5.064172	-5.069439
MA(4)	-5.055365	-5.065345	-5.070192
ARMA(1,1)	-5.054493	-5.063666	-5.068022

Table 5: P-values for diagnostic tests on standardized residuals of selected ARMA-GARCH models

	Ljung-Box SR	ADF Test SR	Jarque-Bera SR	Ljung-Box SR^2
AR(4)-ARCH(1)	0.4330284	0.01	0	0.4436046
AR(4)-ARCH(2)	0.2861988	0.01	0	0.7081430
AR(4)-ARCH(3)	0.4475912	0.01	0	0.6881576
MA(4)-ARCH(1)	0.3405286	0.01	0	0.4487458
MA(4)-ARCH(2)	0.1930888	0.01	0	0.7171979
MA(4)-ARCH(3)	0.3694904	0.01	0	0.7027299
ARMA(1,1)-ARCH(1)	0.5250369	0.01	0	0.4469825
ARMA(1,1)-ARCH(2)	0.3372798	0.01	0	0.7202630
ARMA(1,1)-ARCH(3)	0.3152959	0.01	0	0.7063508

Table 6: Summary statistics for AR(4)-ARCH(1)

AIC	BIC	SIC	HQIC
-5.055365	-5.042858	-5.055373	-5.050897

Table 7: Summary statistics for ARMA(1,1)-ARCH(2)

AIC	BIC	SIC	HQIC
-5.063666	-5.052946	-5.063672	-5.059837