UPPSALA
UNIVERSITET

U.U.D.M. Project Report 2018:9

Stock Price Predictions using a Geometric Brownian Motion

Joel Lidén

Examensarbete i matematik, 30 hp
Handledare: Maciej Klimek
Examinator: Erik Ekström
Juni 2018

Department of Mathematics
Uppsala University

Stock Price Predictions using a Geometric Brownian Motion

Joel Lidén
Degree Project E in Financial Mathematics
Uppsala University
Supervisor: Maciej Klimek
Spring 2018

May 28, 2018

Contents

1 Acknowledgements

I would like to express my deepest gratitude to my supervisor Maciej Klimek for his support and guidance throughout this project. I am deeply greatful for his valuable inputs and also for the relevant research material that he has provided me with. I would also like to thank my good friend Yuqiong Wang for our interesting discussions and her valuable comments.

2 Abstract

In this study a Geometric Brownian Motion (GBM) has been used to predict the closing prices of the Apple stock price and also the S&P500 index. Additionally, closing prices have also been predicted by using mixed ARMA(p,q)+GARCH(r,s) time series models. Using 10 years of historical closing prices between 2008-2018, the predicted prices have also been compared to observed stock prices, in order to evaluate the validity of the prediction models. Predictions have been made using Monte Carlo methods in order to simulate price paths of a GBM with estimated drift and volatility, as well as by using fitted values based on an ARMA(p,q)+GARCH(r,s) time series model. The results of the predictions show an accuracy rate of slightly above 50% of predicting an up- or a down move in the price, by both using a GBM with estimated drift and volatility and also a mixed ARMA(p,q)+GARCH(r,s) model, which is also consistent with the results of K. Reddy and V. Clinton (2016) [1].

3 Introduction

In order to make financial investment decisions, simulated price paths of financial assets are often used to make predictions about the future price. The stochastic price movements of financial assets are often modelled by a GBM, using estimates of the drift and volatility. Currently there is an abundance of historical financial data available for download, which can be used to estimate parameters and compare simulations to actual historical prices. In this study, data containing ten years of historical closing prices of the Apple stock and the S&P500 index has been retrieved from NASDAQ's stock exchange [2], and predictions have been simulated and tested against historical data using standard statistical tests. The R software has been used in order to simulate price movements and to fit mixed time series models.

The predictions have been made by using a specified time frame of historical data which estimates the drift and volatility used in a GBM. Since assuming that the drift and volatility are constant throughout a long time frame is not realistic, the length of the time frame used has been varied in order to improve the simulated predictions. Different estimation methods, such as the standard bootstrap method, have also been implemented in order to improve the estimates of drift and volatility, and therefore also the model assumptions of the GBM.

The model assumptions of a GBM have also been investigated further. Especially the normality assumption of the logarithmic change in the price movement has been subject to debate in the past by researchers such as B. Mandelbrot [3] and G. Dhesi et al [4]. Empirical data shows signs of leptokurtosis when compared to the supposed normal distribution, and therefore using a modified distribution or bootstrap estimates might yield a better fit to the data.

4 Data Analysis

4.1 Apple Stock Price

In order to assess the validity of the prediction models, historical closing prices of the Apple stock has been compared to simulated prices by using basic statistical tests. A time series of the closing prices of the Apple stock during 2008-2018, as well as the log returns of the series can be seen in Figure 1. As can be seen by the time series of returns, the data shows signs of volatility clustering, with large volatility around the time of the financial crisis during the fall of 2008. Here, $S(t_i)$ is the price of the stock at time t_i, and $r_i = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right)$ is the log return at time t_i. Therefore the time series of stock prices can be expressed as $\{S(t_i)\}_{i=1}^{n}$ and the series of log returns as $\{r_i\}_{i=1}^{n}$.

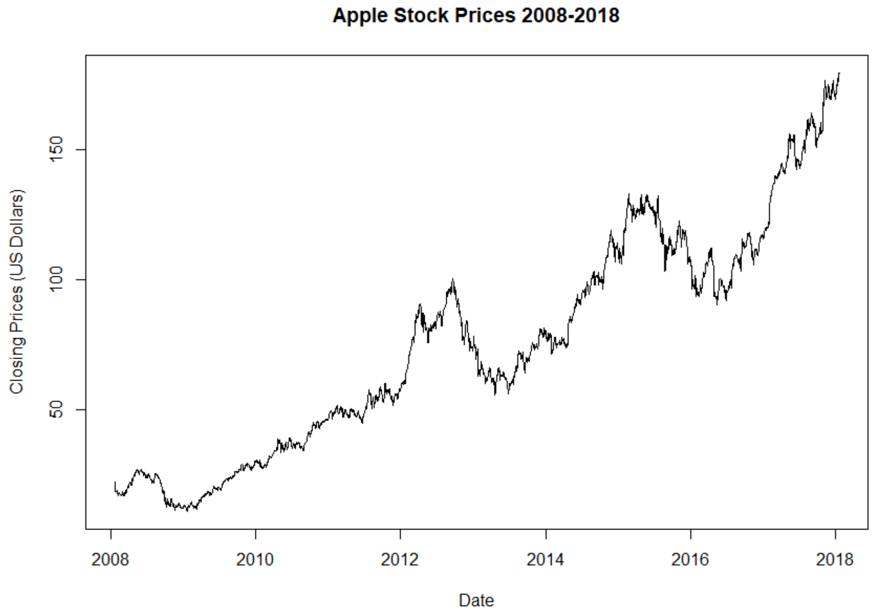

(a) Apple Stock Price

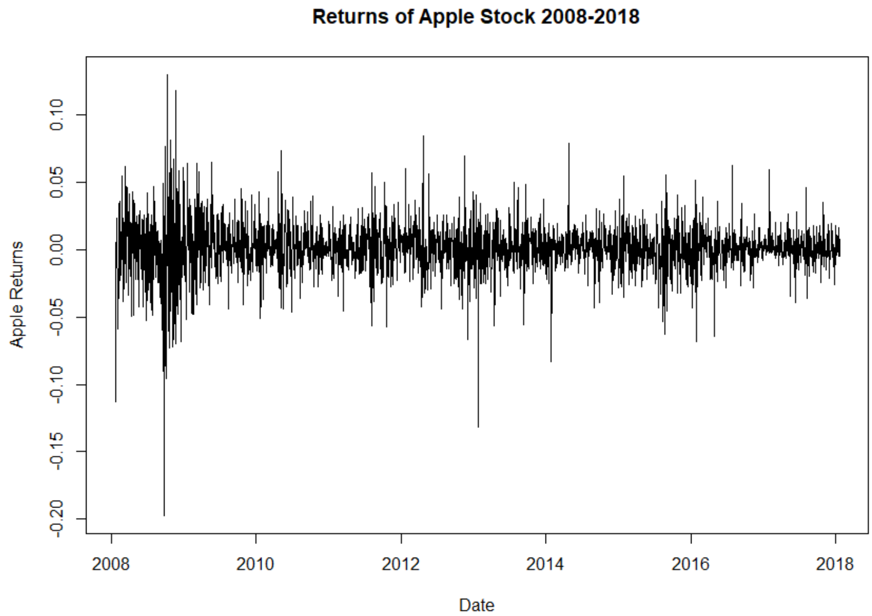

(b) Apple log returns

Figure 1: Apple stock prices and log returns during 2008-2018

In a standard time series model, stationarity of the series is usually assumed, with constant mean and variance of the error terms. Also, all error terms are assumed to be independent and normally distributed according to Equation 1.

$$r_i = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right) = \mu + \epsilon_i, \qquad \epsilon_i \sim Normal(0, \sigma^2) \tag{1}$$

Some basic statistics of the log returns of the Apple stock price can be found in Table 1. As can be seen, the distribution is negatively skewed, suggesting a left-skewed distribution. The normality assumption of the log returns is strongly rejected by the Jarque-Bera test.

Sample size	**Max**	**Min**	**Mean**	**Standard Deviation**	**Skewness**	**Jarque-Bera Test**	**p-value**
2517	0.1301903	-0.1974729	0.0008274646	0.01950108	-0.5083595	7878.3	2.2e-16

Table 1: Basic statistics of the Apple log returns

Figure 2 shows the autocorrelation function (ACF) and partial autocorrelation (PACF) of the Apple log returns series. As can be seen, the dependence structure between lags dies out slowly, since financial time series data displays evidence of "long memory" properties, and therefore the Apple stock price data can not be regarded as realizations of an i.i.d. process.

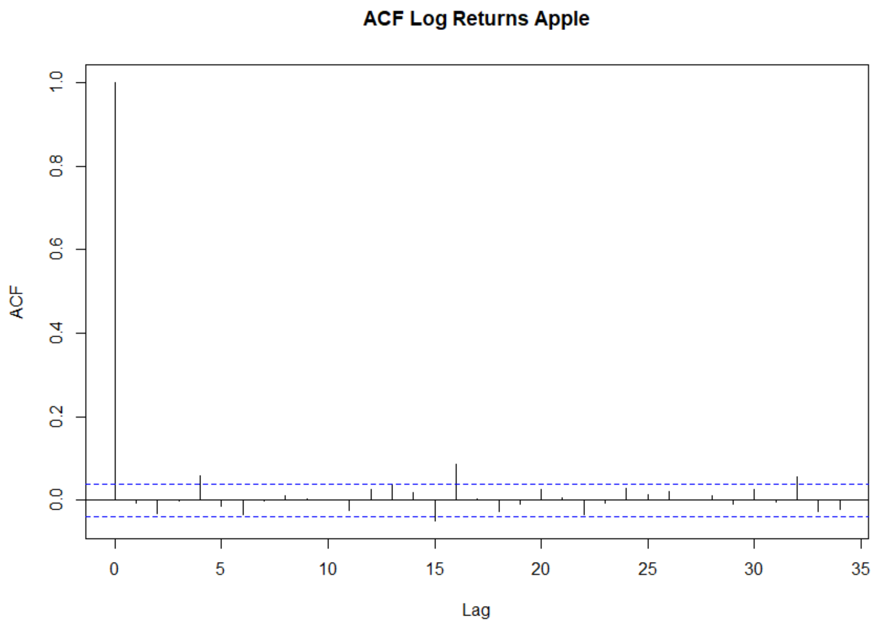

(a) ACF of the Apple log returns series

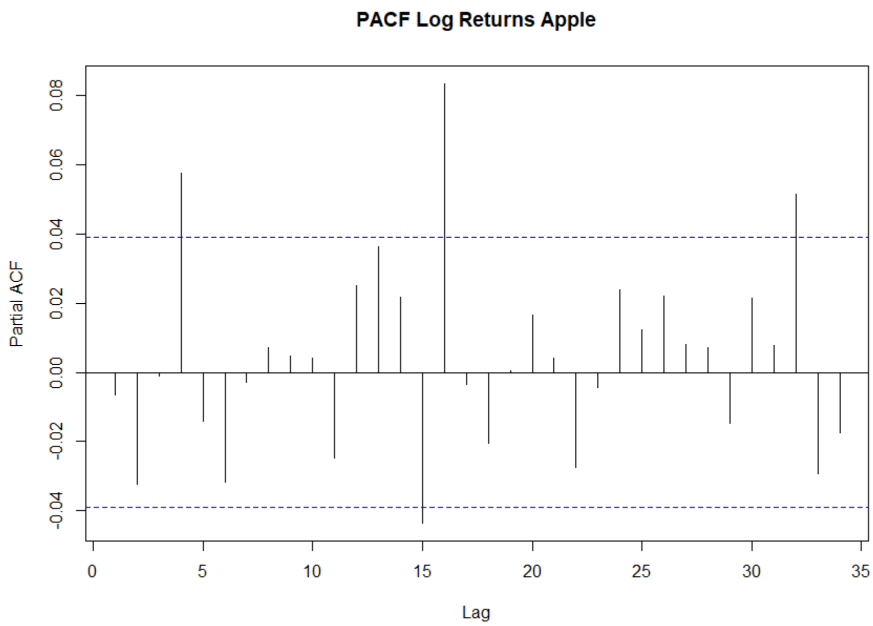

(b) PACF of the Apple log returns series

Figure 2: Autocorrelation and Partial autocorrelation of the Apple log returns

4.2 S&P500 Index

Historical S&P500 index prices has also been compared to simulated prices. A time series of the S&P500 Index during 2008-2018 can be seen in Figure 3. Analogous to the Apple stock prices, the log returns can be expressed according to Equation 1, which also assumes a stationary process with constant mean and variance, and also independent, normally distributed error terms.

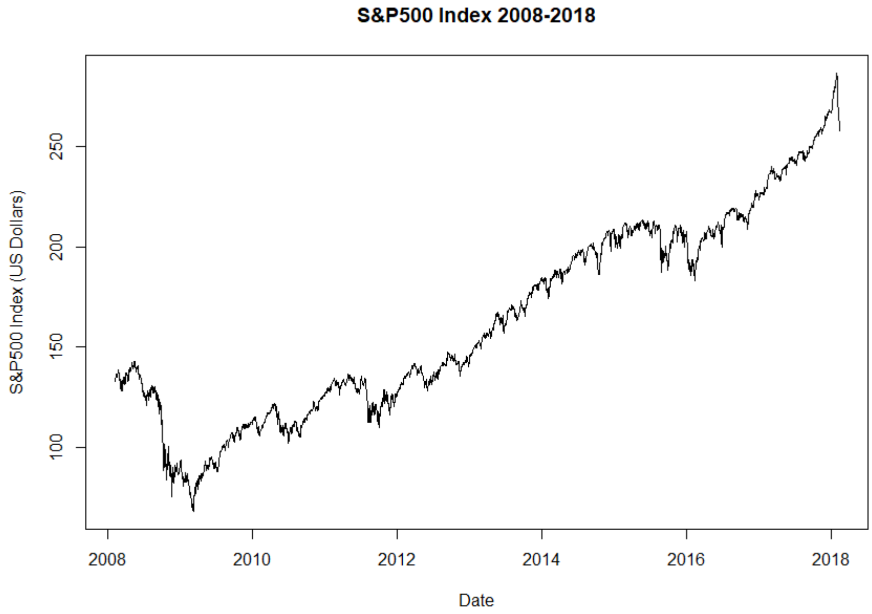

(a) S&P500 Index

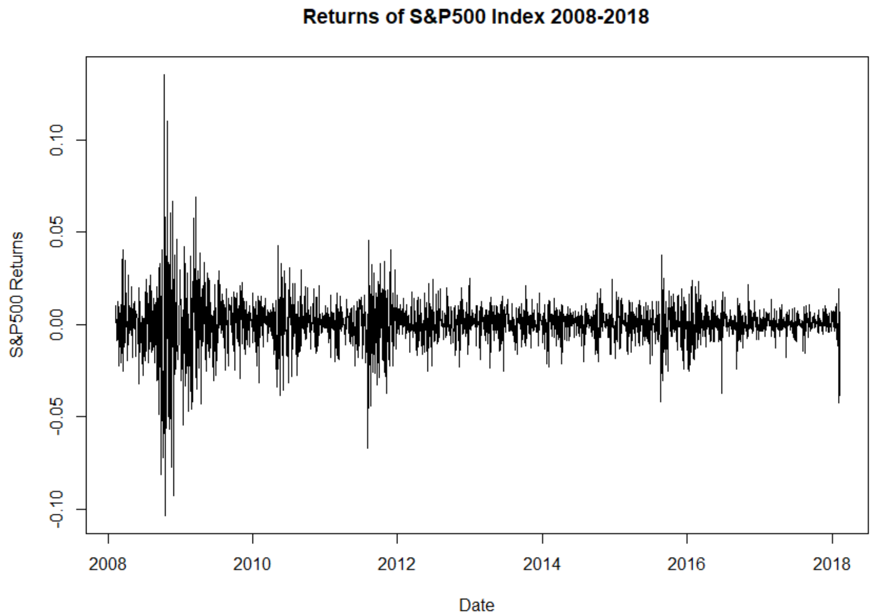

(b) S&P500 returns

Figure 3: S&P500 index prices and log returns during 2008-2018

In time series analysis, stationarity of the series is usually assumed, and all error terms are assumed to be independent. Basic statistics of the log returns of the S&P500 index can be found in Table 2. Just as the Apple stock, the log returns show signs of volatility clustering, especially around the time of the financial crisis in the fall of 2008. The normality assumption of the log returns is strongly rejected by the Jarque-Bera test, and the distribution has a slight left-skewness.

Sample size	Max	Min	Mean	Standard Deviation	Skewness	Jarque-Bera Test	p-value
2518	0.1355773	-0.1036372	0.0002623706	0.01281492	-0.1096676	21764	2.2e-16

Table 2: Basic statistics of the S&P500 index log returns

Figure 4 shows the autocorrelation function (ACF) and partial autocorrelation (PACF) of the S&P500 series. Just as the Apple stock price, the dependence structure between lags dies out slowly, and therefore the realizations of the S&P500 series cannot be regarded to be an i.i.d. process.

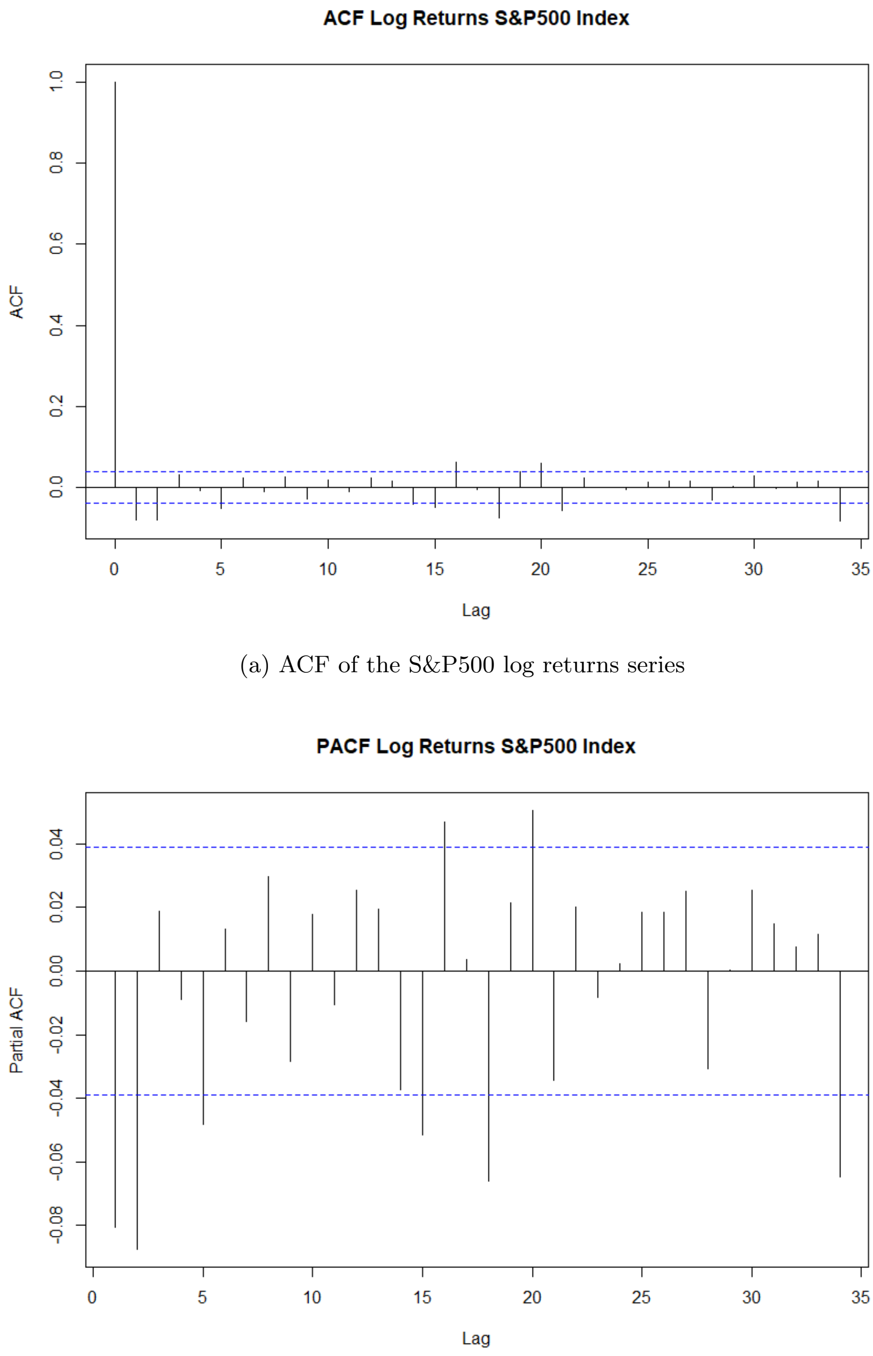

(a) ACF of the S&P500 log returns series

(b) PACF of the S&P500 log returns series

Figure 4: Autocorrelation and Partial autocorrelation of the S&P500 log returns

5 Theory

5.1 Expectation of a Geometric Brownian Motion

In order to find the expected asset price, a Geometric Brownian Motion has been used, which expresses the change in stock price using a constant drift μ and volatility σ as a stochastic differential equation (SDE) according to [5]:

$$\begin{cases} dS(t) = \mu S(t)dt + \sigma S(t)dW(t) \\ S(0) = s \end{cases} \tag{2}$$

By integrating both sides of the SDE and using the initial condition, the solution to this equation is given by:

$$S(t) = s + \mu \int_0^t S(u)du + \sigma \int_0^t S(u)dW(u)$$

Taking the expectation of both sides yields:

$$E[S(t)] = s + \mu \int_0^t E[S(u)]du + \sigma E\left[\int_0^t S(u)dW(u)\right]$$

The expectation of the stochastic integral is simply zero. Substituting $E[S(t)] = m(t)$ and using the initial condition $m(0) = s$, we can express the equation as an ordinary differential equation, according to:

$$\begin{cases} m'(t) = \mu m(t) \\ m(0) = s \end{cases}$$

Clearly, this simple ODE has the solution $m(t) = se^{\mu t}$. Therefore, the expectation of the stock price at time t is:

$$E[S(t)] = se^{\mu t} \tag{3}$$

To find the solution $S(t)$ to the SDE, we can use the substitution $Z(t) = logS(t)$, since the corresponding deterministic linear equation is an exponential function of time. Itô's formula yields:

$$dZ = \frac{1}{S}dS + \frac{1}{2}\left(-\frac{1}{S^2}\right)(dS)^2$$

$$= \frac{1}{S}(\mu Sdt + \sigma SdW) + \frac{1}{2}\left(-\frac{1}{S^2}\right)\sigma^2 S^2 dt$$

$$= (\mu dt + \sigma dW(t)) - \frac{1}{2}\sigma^2 dt$$

So we have the following equation for $dZ(t)$:

$$\begin{cases} dZ(t) = \left(\mu - \frac{1}{2}\sigma^2\right) dt + \sigma dW(t) \\ Z(0) = \log s \end{cases}$$

By integrating both sides and substituting back $S(t)$ yields the solution:

$$S(t) = s \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma W(t)\right) \tag{4}$$

Equivalently, we can express this equation as:

$$\log S(t) - \log s = \left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma W(t) \tag{5}$$

5.2 Distribution assumption

As mentioned in Section 4, a GBM assumes the logarithmic change of the stock price to be a normally distributed random variable according to:

$$r_i = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right) = \mu + \epsilon_i, \qquad \epsilon_i \sim Normal(0, \sigma^2)$$

This assumption can also be tested against historical data, as can be seen in Figure 5. The fitted normal distribution, which uses the entire sample period of 10 years of closing prices of the Apple stock and the S&P500 index to estimate the expected value and variance of the logarithmic change of the stock price, does not quite capture the actual distribution, which shows signs of leptokurtosis. Therefore, a modified distribution can be used to yield a better fit to the distribution of returns, as suggested by G. Dheesi et al [4]. As can be seen, the Cauchy distribution yields a better fit for both data sets, although the moments are not defined. Using a non-parametric kernel density estimation which is described in Section 8 yields an even better fit of the historical returns.

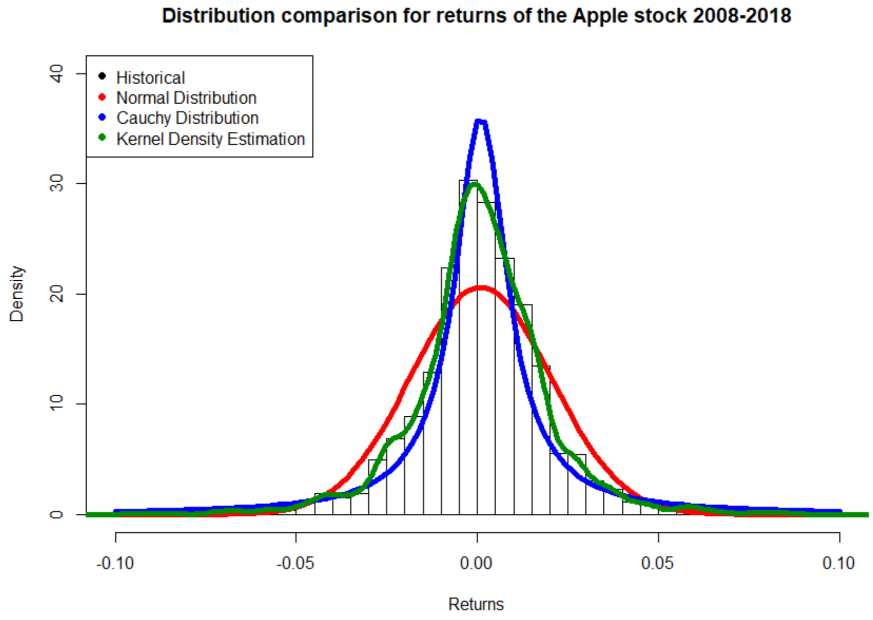

(a) Apple returns

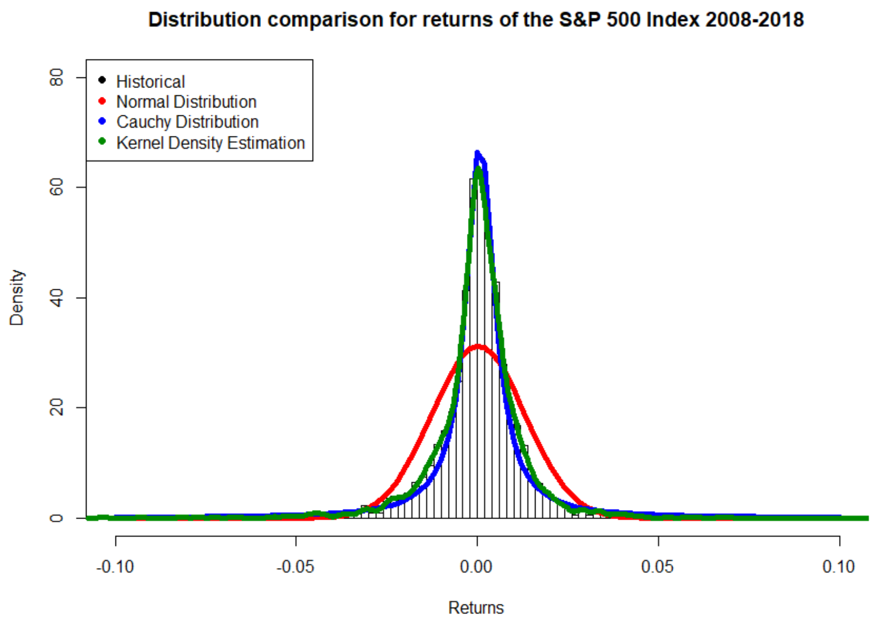

(b) S&P500 returns

Figure 5: Distribution comparisons for the returns of Apple stock and S&P500 index

In Figure 6, the skewness and kurtosis defined in Equation 6 and 7 of the returns in the historical Apple stock prices and S&P500 index between 2008-2018 can be seen. The distribution of Apple returns suggests a distribution that is slightly skewed and more heavy-tailed than the Gaussian, which has zero skewness and a kurtosis of 3. In the S&P500 case, the distribution does not show signs of skewness but also suggests heavy tails.

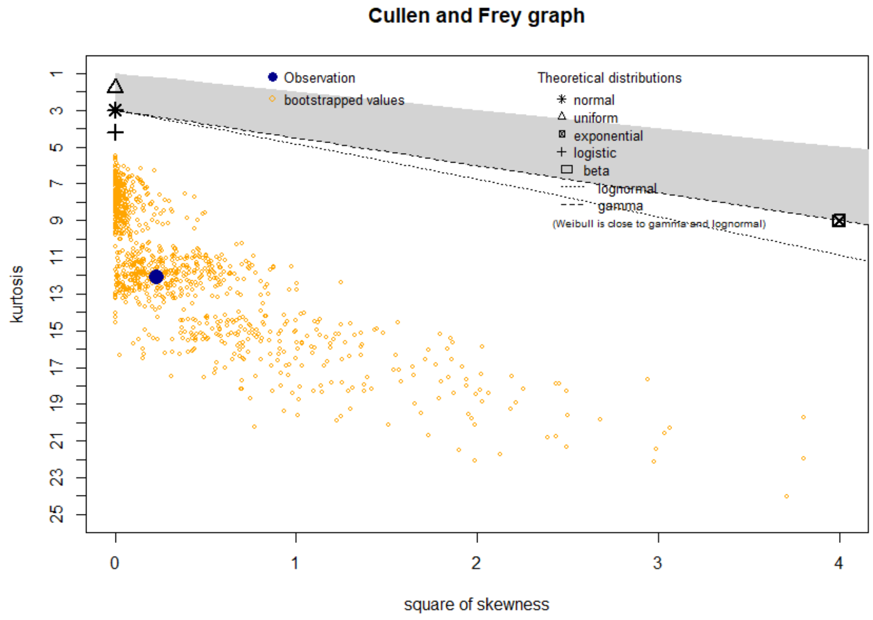

(a) Apple returns

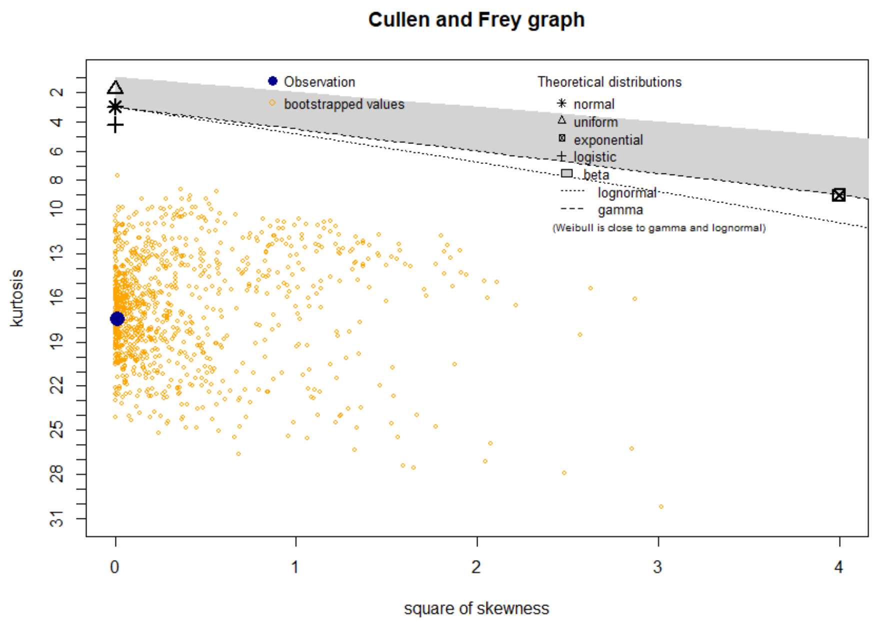

(b) S&P500 returns

Figure 6: Skew and kurtosis for both the Apple returns data and the S&P500 index data

$$Skew[X] = E\left[\left(\frac{X-\mu}{\sigma}\right)^3\right] = \frac{E[(X-\mu)^3]}{(E[(X-\mu)^2])^{3/2}} \tag{6}$$

$$Kurt[X] = E\left[\left(\frac{X-\mu}{\sigma}\right)^4\right] \tag{7}$$

Since the Cauchy distribution yields a better fit for the data, the logarithmic change in price is assumed to be a Cauchy distributed random variable according to:

$$r_i^* = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right) = \mu^* + \epsilon_i^*, \qquad \epsilon_i^* \sim Cauchy(0, \gamma)$$

5.3 Kernel density estimation

An alternative way to estimate a probability distribution to a data set is to use a non-parametric kernel density estimation (KDE). In this case a KDE has been used to estimate the distribution of both the Apple returns and the S&P500 returns, in order to find a better fit for the data. A KDE is defined as follows, where $(r_1, ..., r_n)$ is an independent and identically distributed sample drawn from a distribution with density function f:

$$\widehat{f(r)}_h = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{r - r_i}{h}\right) \tag{8}$$

Here $h > 0$ is a smoothing bandwidth parameter which essentially represents the width of each bin of the underlying histogram and K is a non-negative kernel function that integrates to one. Intuitively, in order to approximate the density f as well as possible, h should be made as small as possible. In this case, a Gaussian kernel has been used, which has the following properties:

$$K(u) = \frac{1}{\sqrt{2\pi}} exp(-\frac{u^2}{2})$$

$$K(u) = K(-u) \qquad \int_{-\infty}^{\infty} K(u)du = 1 \qquad \int_{-\infty}^{\infty} uK(u)du = 0$$

So the kernel is symmetric about zero, integrates to one and has expectation zero. The substitution $K_h(u) = \frac{1}{h}K\left(\frac{u}{h}\right)$ is also a kernel function and a density. The mean of the estimated density $\hat{f(r)}_h$ can therefore be estimated as follows, using the variable substitution $u = \frac{r - r_i}{h}$:

$$\begin{aligned}
\int_{-\infty}^{\infty} r\widehat{f(r)}dr &= \int_{-\infty}^{\infty} \frac{1}{n}\sum_{i=1}^{n} rK_h(r - r_i)dr \\
&= \frac{1}{n}\sum_{i=1}^{n}\int_{-\infty}^{\infty}(r_i + uh)K(u)du \\
&= \frac{1}{n}\sum_{i=1}^{n}\left(r_i\int_{-\infty}^{\infty}K(u)du + h\int_{-\infty}^{\infty}uK(u)du\right) \\
&= \frac{1}{n}\sum_{i=1}^{n} r_i
\end{aligned}$$

So the mean of the estimated density is simply the sample mean of the original observations $(r_1, ..., r_n)$.

5.4 Goodness-of-fit tests

In order to assess the fit of each proposed distribution, Chi-squared goodness-of-fit tests have been made. The test statistic is computed as follows:

$$X^2 = \sum_{i=1}^{n} \frac{(O_i - E_i)^2}{E_i} \sim \chi^2_{n-p-1}$$

Here $(n - p - 1)$ is the degrees of freedom used, where n is the number of customized bins used in the data set and p is the number of parameters used in the proposed distribution. The results of the tests can be found in Table 3.

Distribution assumption	**Data**	**Test statistic**	**Degrees of freedom**	**p-value**
Normal	Apple Stock	274.38	21	2.2e-16
Cauchy	Apple Stock	181.86	32	2.2e-16
KDE	Apple Stock	14.339	24	0.9386
Normal	S&P500	388.86	13	2.2e-16
Cauchy	S&P500	90.829	19	2.365e-11
KDE	S&P500	39.403	32	0.1725

Table 3: Chi-Squared Goodness-of-Fit tests for different distribution assumptions

Clearly, both the normal and the Cauchy distribution assumptions are rejected based on the Chi-square goodness-of-fit test, while the kernel density estimation (KDE) yields a better fit to the data for both the Apple- and the S&P500 index returns.

5.5 Estimating Drift and Volatility

Historical data has been used in order to find estimates of the drift $\hat{\mu}$ and volatility $\hat{\sigma}$. Since it's not reasonable to assume that drift and volatility are constant throughout a long time period, estimates have been made with a varying time frame of prior historical closing prices of the Apple stock and the S&P500 Index, generating prior observations $\{S(t_1), ..., S(t_n)\}$. To estimate σ, we can use the fact that S has a log-normal distribution. We can therefore define $\{r_1, ..., r_n\}$ as follows [5]:

$$r_i = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right)$$

The observations $\{r_1, ..., r_n\}$ are assumed to be independent, normally distributed random variables, which conflicts with the "long memory" property described in Section 4. However, since the dependence of the log returns is quite weak, the observations are assumed to be i.i.d. and therefore the expectation and variance can be expressed as:

$$E[r_i] = \left(\mu - \frac{1}{2}\sigma^2\right)\Delta t \qquad\qquad Var[r_i] = \sigma^2 \Delta t$$

The sample variance is given by:

$$S_\mu^2 = \frac{1}{n-1}\sum_{i=1}^{n}(r_i - \hat{\mu})^2$$

An estimate of μ and σ is therefore given by:

$$\hat{\mu} = \frac{1}{n}\sum_{i=1}^{n} r_i \qquad\qquad \hat{\sigma} = \frac{S_\mu}{\sqrt{\Delta t}}$$

Assuming a Cauchy distribution, estimating the location and scale parameter can be done by using the maximum likelihood estimate for a sample of size n according to:

$$\hat{l}(\mu^*, \gamma|\mu_1, ..., \mu_n) = -nlog(\gamma\pi) - \sum_{i=1}^{n} log\left(1 + \left(\frac{\mu_i - \mu^*}{\gamma}\right)^2\right)$$

Maximizing the log-likelihood with respect to μ^* and γ yields the following system of equations:

$$\begin{cases} \sum_{i=1}^{n} \frac{\mu_i - \mu^*}{\gamma^2 + (\mu_i - \mu^*)^2} = 0 \\ \sum_{i=1}^{n} \frac{\gamma^2}{\gamma^2 + (\mu_i - \mu^*)^2} - \frac{n}{2} = 0 \end{cases}$$

An estimate of the location parameter μ^* can then be found by an approximate numerical solution of the system. So therefore it's possible to find estimates of the drift and volatility by using historical data, and then find the expected stock price at time t_{n+1}, by using the expectation of a Geometric Brownian Motion expressed in Equation 3 and Euler's discretization process:

$$E[S(t_{n+1})] = S(t_n)e^{\mu\Delta t}, \qquad \Delta t = 1 \tag{9}$$

5.6 Nonparametric estimation using Standard Bootstrap Method

In order to retrieve a distribution of a statistic, bootstrap simulations can be used. Since the normality assumption of the log returns does not seem to be valid, estimating the drift by simply using the sample mean might not be accurate. Also, when using a small time frame in order to estimate drift and volatility, very few observations are used which can cause erroneous estimates. Therefore, using a bootstrap method can yield estimates based on a large number of resamples, which can reduce the estimation errors [7]. An arbitrary distribution can be assumed as follows:

$$r_i = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right) \sim F_i$$

In Section 4, the "long memory" property of both the S&P500 and the Apple series could be seen, since the ACF tails off slowly, showing dependence between historical lags. In standard bootstrap methodology, the observations are assumed to be i.i.d, which conflicts with the dependence structure of the data. However, since the dependence between lags is quite weak, the standard bootstrap method has still been used in order to find a distribution of the mean and variance of the drift. Therefore, assuming that all observations are an i.i.d. sample of n returns which can be expressed as $S = (r_1, ..., r_n)$, bootstrapping can be used to find a probability distribution of the estimator $\hat{\theta}(S_i)$. Formally, this is done by sampling the data randomly with replacement, drawing random resamples $S_i = (r^*_{i1}, ..., r^*_{in}), i = 1, ..., N$ from $S = (r_1, ..., r_n)$. By repeating this procedure N times, drawing N resamples and computing $\hat{\theta}(S_i)$ for each resample, a distribution of the parameter r_i can be found.

Also, one can estimate the parameter of interest $\theta(F_i)$, where F_i is the probability distribution of r_i. By resampling with replacement N times from the original sample S, and assuming that the observations are i.i.d., the distribution of means will approach normality by the Central

Limit Theorem. Using the Standard Bootstrap Method to estimate the drift will yield the same estimate as the sample mean. Let $S_i = (r_1^*, ..., r_n^*)$ be a sample from $S = (r_1, ..., r_n)$. Then let $\hat{\theta}_i$ be the arithmetic mean of S_i. Since the observations of the original sample are i.i.d., they're all assumed to have the same expectation $E[r_1] = ... = E[r_n] = \mu$. Now, the mean and standard error of the bootstrap estimator can be expressed as:

$$\bar{\theta} = \frac{1}{N}\sum_{i=1}^{N}\hat{\theta}_i \qquad\qquad SE(\hat{\theta}) = \frac{1}{N-1}\sum_{i=1}^{N}(\hat{\theta}_i - \bar{\theta})^2$$

5.7 Distribution of Drift using Standard Bootstrap Method

Since the normality assumption of the log returns does not seem to be valid, nonparametric estimation of drift and variance can be made using a standard bootstrap method. In Figure 7 the distribution of means of the Apple and S&P500 log returns can be seen using 10 000 bootstrapped resamples of the original sample data from 2008-2018. As can be seen, when using the standard bootstrap method, a normal distribution of the parameters seems to be valid when using an arbitrarily large number of resamples.

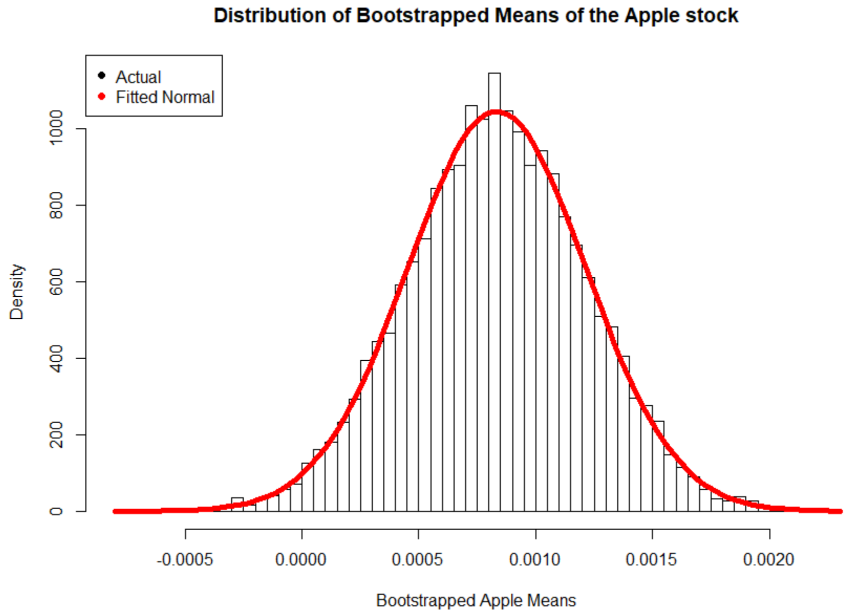

(a) Bootstrap Distribution of Apple Means

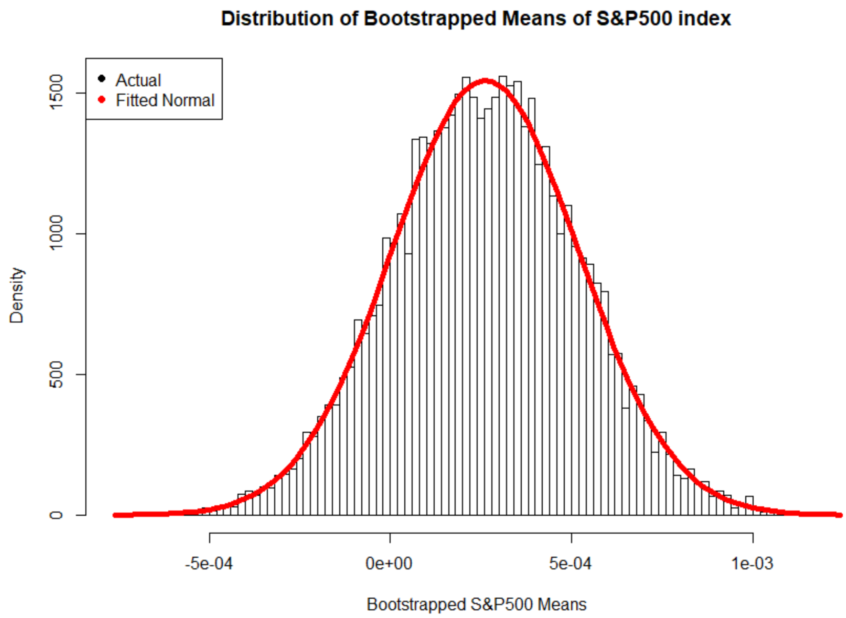

(b) Bootstrap Distribution of S&P500 Means

Figure 7: Distribution comparisons for the bootstrapped means of log returns for the Apple stock and the S&P500 index

In Table 4, comparisons between the actual sample estimates of the drift can be compared to the bootstrap estimates with a 95% confidence interval. For the sample estimates of the drift, a 95% confidence interval is simply given by $\hat{r} \pm t_p \frac{s}{\sqrt{n}}$, where t_p is the p^{th} percentile of the Student's t-distribution. For the bootstrap case, a 95% confidence interval is given by $[\delta_{0.025}, \delta_{0.975}]$, where δ_p is the p^{th} percentile of the resampled bootstrap distribution. As can be seen, the bootstrap estimates of the drift are approximately the same as the original sample estimates, although the bootstrap estimator reduces the error for a small sample size of 10, yielding narrower confidence intervals.

Data	Sample Size	$\hat{\mu}_{\text{Sample}}$	$\hat{\mu}_{\text{Bootstrap}}$	95%CI$_{\text{Sample}}$	95%CI$_{\text{Bootstrap}}$
Apple Stock	10	0.00309	0.00309	[-0.00224, 0.00842]	[-0.00106, 0.00767]
Apple Stock	100	0.0011	0.0011	[-0.001, 0.0032]	[-0.00094, 0.00319]
Apple Stock	2517	0.00083	0.00083	[0,00007, 0.00159]	[0,00008, 0.00159]
S&P500	10	-0.0095	-0.00955	[-0.02373, 0.00474]	[-0.02151, 0.00197]
S&P500	100	0.00033	0.00033	[-0.00118, 0.00184]	[-0.00129, 0.0017]
S&P500	2518	0.00026	0.00027	[-0.00024, 0.00076]	[-0.00023, 0.00077]

Table 4: Sample estimates of drift compared to bootstrap estimates

5.8 Using a mixed GARCH model

When modelling stock returns, a Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model is commonly used if the time series data shows signs of time-varying volatility, i.e. periods of swings interspersed with periods of relative calm. Volatility clustering is present during financial instability, which was the case during the financial crisis of 2008. Since a GARCH model allows for a non-stationary volatility, which is the case in both the data of the Apple stock and the S&P500 index, it's appropriate for modelling returns. In general, a mixed ARMA(p,q)+GARCH(r,s) model can be expressed as follows [6]:

$$r_t = \alpha + \phi_1 r_{t-1} + ... + \phi_p r_{t-p} + \theta_1 w_{t-1} + ... + \theta_q w_{t-q} + w_t$$

$$w_t | w_{t-1} \sim N(0, \sigma_t^2)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 w_{t-1}^2 + ... + \alpha_r w_{t-r}^2 + \beta_1 \sigma_{t-1}^2 + ... + \beta_s \sigma_{t-s}^2$$

Therefore, the logged returns are modelled by a mixed ARMA(p,q)+GARCH(r,s) model with stationary mean but non-stationary volatility. Predicting the log return one time step forward and then transforming the log return to an actual stock price prediction yields the following, by substituting $r_t = \log\left(\frac{P_t}{P_{t-1}}\right)$:

$$\log\left(\frac{P_t}{P_{t-1}}\right) = \alpha + \phi_1 \log\left(\frac{P_{t-1}}{P_{t-2}}\right) + ... + \phi_p \log\left(\frac{P_{t-p}}{P_{t-p-1}}\right) + \theta_1 w_{t-1} + ... + \theta_q w_{t-q} + w_t$$

$$w_t | w_{t-1} \sim N(0, \sigma_t^2)$$

$$\Rightarrow P_t = P_{t-1} \exp\left\{\alpha + \phi_1 \log\left(\frac{P_{t-1}}{P_{t-2}}\right) + ... + \phi_p \log\left(\frac{P_{t-p}}{P_{t-p-1}}\right) + \theta_1 w_{t-1} + ... + \theta_q w_{t-q} + w_t\right\}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 w_{t-1}^2 + ... + \alpha_r w_{t-r}^2 + \beta_1 \sigma_{t-1}^2 + ... + \beta_s \sigma_{t-s}^2$$

6 Results

6.1 Predicting the Apple Stock Price using a Geometric Brownian Motion

6.1.1 Predicting One Time Step Forward

When predicting the Apple stock price by simply using the expectation of a GBM expressed in Equation 9, the validity of the historical data used to estimate drift and volatility can be questioned. In order to investigate if using a longer or shorter time frame improves the predictions, a varying number of historical closing days have been used in order to estimate drift and volatility. Distribution assumptions of the returns have also been varied, using both the normality and Cauchy assumptions. In order to estimate the drift, the sample mean has been used for the normality assumption, and a bootstrap estimate has also been used with 10 000 resamples. In the Cauchy case, the location parameter of a fitted Cauchy distribution to the data has been used.

Using the Mean Square Error (MSE), and the proportion of correct up or down movements in the price ($\hat{p}$), the results can be found in Table 5. Using 60 days of historical closing prices yields the lowest MSE for the normality assumption, while using 100 days yields the largest probability of predicting an accurate up- or down move in the price for the normal distribution assumption and by also using a bootstrap estimate of the drift. In the Cauchy case, 60 days also yielded the lowest MSE. Overall, the best predictions were only accurate slightly more than 50 % of the time. The Cauchy distribution provided lower MSE-values compared to the normality assumption.

Sample size	$\mathbf{MSE_{Normal}}$	$\hat{\mathbf{p}}_{\mathbf{Normal}}$	$\mathbf{MSE_{Cauchy}}$	$\hat{\mathbf{p}}_{\mathbf{Cauchy}}$	$\mathbf{MSE_{Bootstrap}}$	$\hat{\mathbf{p}}_{\mathbf{Bootstrap}}$
20	1.68844	0.5252202	1.696666	0.5152122	1.688678	0.5260208
30	1.676837	0.5188907	1.667149	0.5132637	1.676609	0.5204984
40	1.67217	0.5056497	1.666231	0.5092817	1.672332	0.5056497
50	1.665432	0.5202593	1.65844	0.5149919	1.665779	0.5198541
60	**1.66206**	**0.5174939**	**1.65229**	**0.5101709**	**1.662152**	**0.5183076**
70	1.665484	0.5216503	1.657948	0.503268	1.665281	0.5216503
80	1.66508	0.5307629	1.659285	0.5139459	1.665176	0.5307629
90	1.673651	0.5201812	1.669046	0.5004119	1.673862	0.5201812
100	1.676367	0.5326716	1.673352	0.5128205	1.676671	0.5326716
200	1.714528	0.5181191	1.708248	0.5163934	1.71448	0.5172563
300	1.782596	0.5081154	1.773526	0.5153291	1.782633	0.506312
400	1.859956	0.5080264	1.853068	0.5028329	1.859907	0.5042493
500	1.941562	0.5208127	1.93728	0.5178394	1.941525	0.5193261
600	2.018263	0.5218978	2.014919	0.5177268	2.018399	0.5218978
700	2.105357	0.5192519	2.104171	0.5181518	2.10529	0.5192519
800	2.202516	0.5157159	2.202301	0.5157159	2.202375	0.5157159
900	2.299643	0.5173053	2.296151	0.5166873	2.299727	0.5179234
1000	2.390633	0.5158103	2.386399	0.5158103	2.390786	0.5158103

Table 5: Different outcomes using a varied time frame of the Apple stock

In Figure 8 the actual stock prices have been compared to the predicted stock prices using a GBM with 60 days of historical data to estimate drift and volatility. As can be seen, volatile periods yield a larger difference between the actual and predicted price, which is expected since the expectation of a GBM does not depend on the volatility but simply the drift.

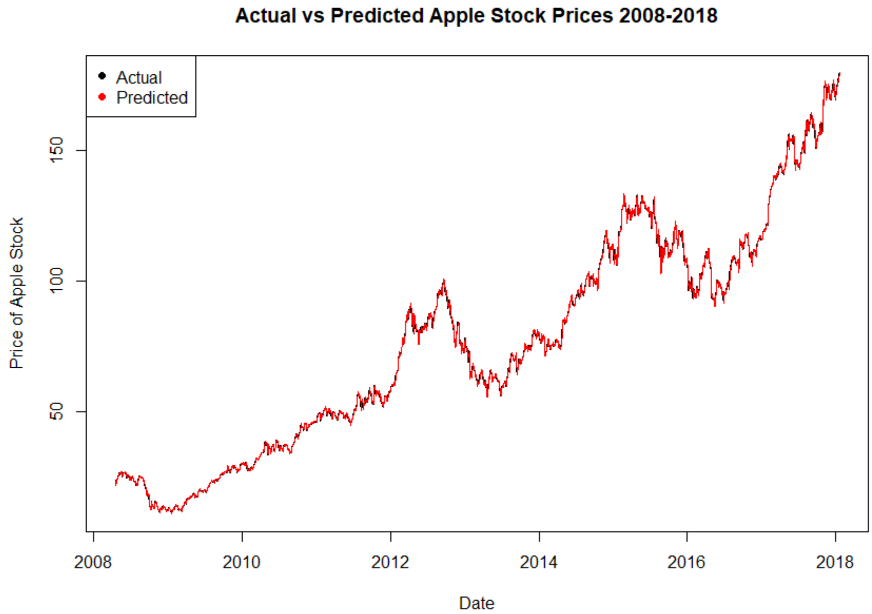

(a) Apple stock returns during 2008-2018

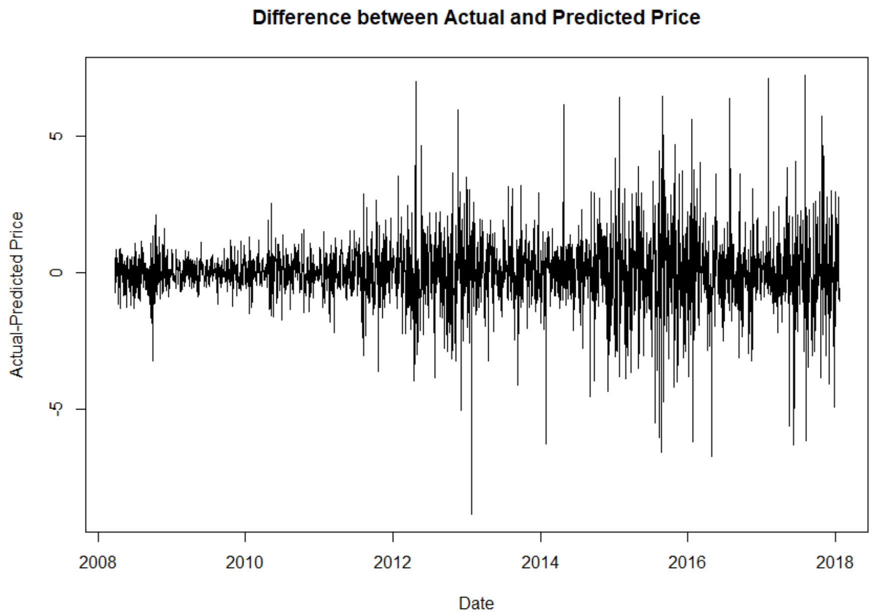

(b) Plot of differences

Figure 8: Actual vs Predicted Apple Stock Prices 2008-2018

6.1.2 Predicting a Longer Time Frame

In order to predict long time frames of Apple stock prices, Monte Carlo simulations have been made. Assuming constant drift and volatility throughout a longer time period is not realistic, which is why simulations have been made using a varied time frame to estimate drift and volatility. An example showing simulations of price trajectories using a GBM with different estimates of drift and volatility can be seen in Figure 9 a). For the normality assumption, sample mean and standard deviation have been used to estimate drift and volatility. For the Cauchy assumption, the location and scale parameter have instead been used to estimate drift and volaility, and for the bootstrap estimate 10 000 resamples have been used in order to simulate a distribution of drift and volatility, and thereafter both parameters have been estimated by taking the mean of the simulated distribution. The expected stock prices retrieved by 1000 Monte-Carlo simulations can be seen in Figure 9 b). In this case, the drift and volatility used in the simulations have been estimated by using the entire data set of closing prices during 2008-2018. In section 5.2 it was suggested that a Cauchy distribution might yield a better fit of the data. However, estimating the drift and volatility by using the location and scale parameter of a Cauchy distribution did not improve the prediction of the Apple stock, as can be seen in Figure 9 b). Bootstrap estimates of drift and volatility have also been used.

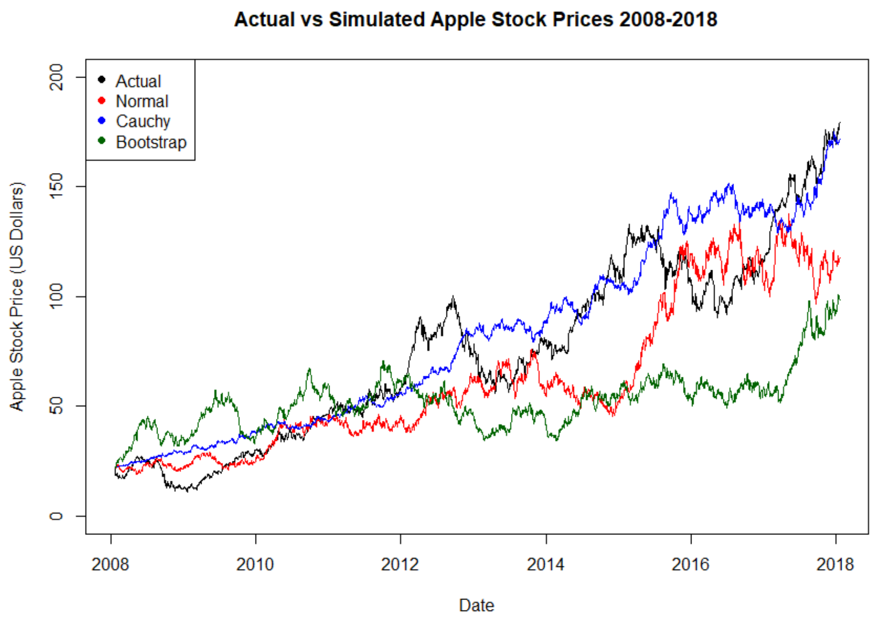

(a) Simulated Apple Stock prices during 2008-2018

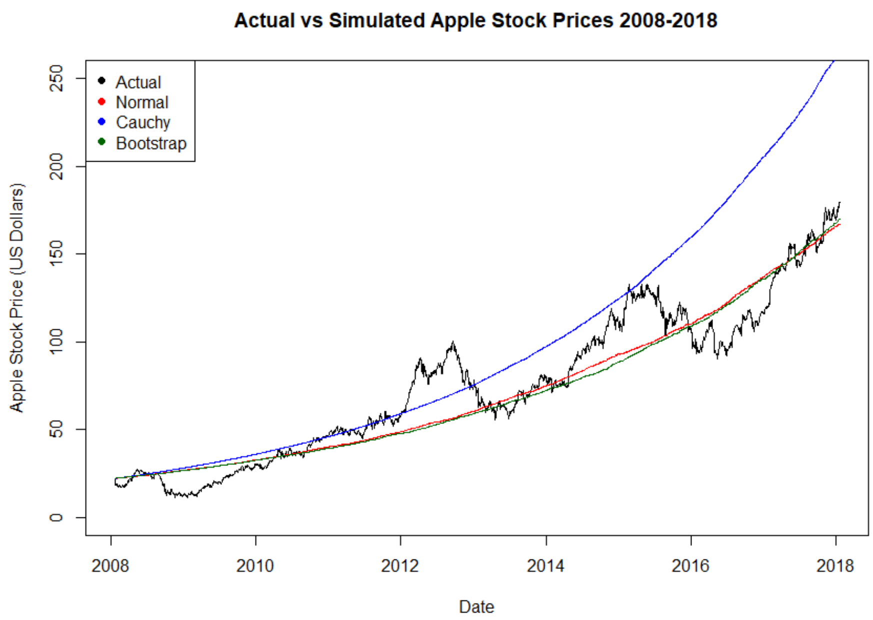

(b) Expected Apple Stock prices during 2008-2018

Figure 9: Simulated prices assuming constant drift of a standard GBM during 2008-2018

In Figure 9 a constant drift and volaility during 2008-2018 was assumed. Since this assumption is clearly not valid, blockwise intervals of equal length containing estimation data have been used in order to estimate drift and volatility in a standard GBM. The results of varying the block length for the estimates can be seen in Figure 10. Clearly, a smaller block length to estimate the drift and volatility yields simulations closer to the actual observed stock prices.

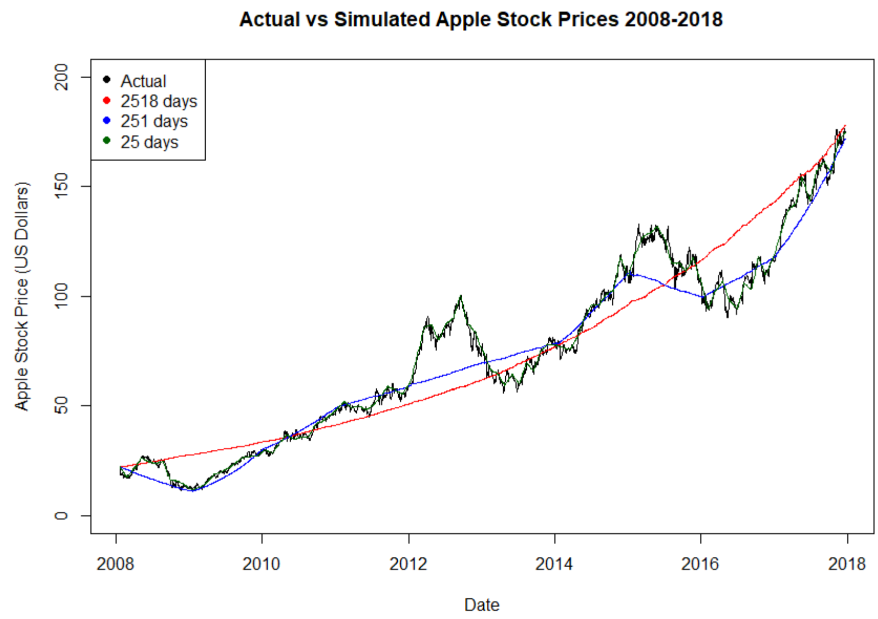

Figure 10: Apple stock returns during 2008-2018

In Table 6, the results of varying the block length to simulate the Apple Stock prices can be seen. In this case, both the normality and the bootstrap estimates for drift and volatility have been used. Since the drift and volatility varies greatly between different time intervals, using a smaller block length for estimates greatly reduces the MSE. Overall, using bootstrap estimates instead of sample estimates of the drift yielded about the same predictions.

Estimation period	**MSE Normal**	**MSE Bootstrap**
2518 days	224.4942	280.655
1258 days	232.3797	207.7554
503 days	190.9083	185.5588
251 days	91.95498	99.43726
125 days	30.46398	30.84619
50 days	13.76284	14.06368
25 days	6.149417	6.167841

Table 6: Different outcomes using a varied time frame of the Apple stock

6.2 Modelling the Apple Stock Price using a mixed GARCH model

The observed Apple stock price as well as the logged returns in the Apple stock price during 2008-2018 can be seen Figure 1. As can be seen, the log returns seem to have a constant mean of approximately zero throughout the series, although a larger volatility around the time of the financial crisis in September 2008. It can also be seen that there is no evident periodicity in the closing prices. Therefore a GARCH model can be used for predictions, which allows for non-stationary volatility.

By looking at the autocorrelation function (ACF) and partial autocorrelation function (PACF) in Figure 2, it's evident that the ACF tails off although lag 4, 16 and 32 show values outside of the bounds. So the data shows signs of long-term dependency, although the dependence is quite weak.

Fitting an ARMA(p,q)-model for the Apple stock prices between 2008-2018 which minimizes the AIC-value yields an ARMA(1,1)-model, suggesting a time series model with one autoregressive lag, one moving average lag and non-zero mean. The model can be expressed as follows, with parameter estimates according to Table 7, and where r_t is the predicted log return:

$$r_t = \alpha + \phi r_{t-1} + w_t + \theta w_{t-1}, \ \ w_t | w_{t-1} \sim N(0, \sigma_t^2) \tag{10}$$

α	ϕ	θ	**SE**(α)	**SE**(ϕ)	**SE**(θ)	$\boldsymbol{\sigma^2}$	**LL**	**AIC**
8e-04	0.3687	-0.3832	4e-04	0.5095	0.5094	0.0003805	6339.48	-12670.95

Table 7: Parameter estimates of the ARMA model

The squared residuals of the fitted ARMA(1,1)-model as well as the ACF and PACF plots can be seen in Figure 11. Volatility clustering can be seen during the time of the financial crisis in 2008. The ACF and PACF plot both show lags outside of the bounds, suggesting dependence between the squared residuals. Modelling the volatility by using a mixed GARCH-model, and choosing the model displaying the lowest AIC yielded a GARCH(1,1)-model.

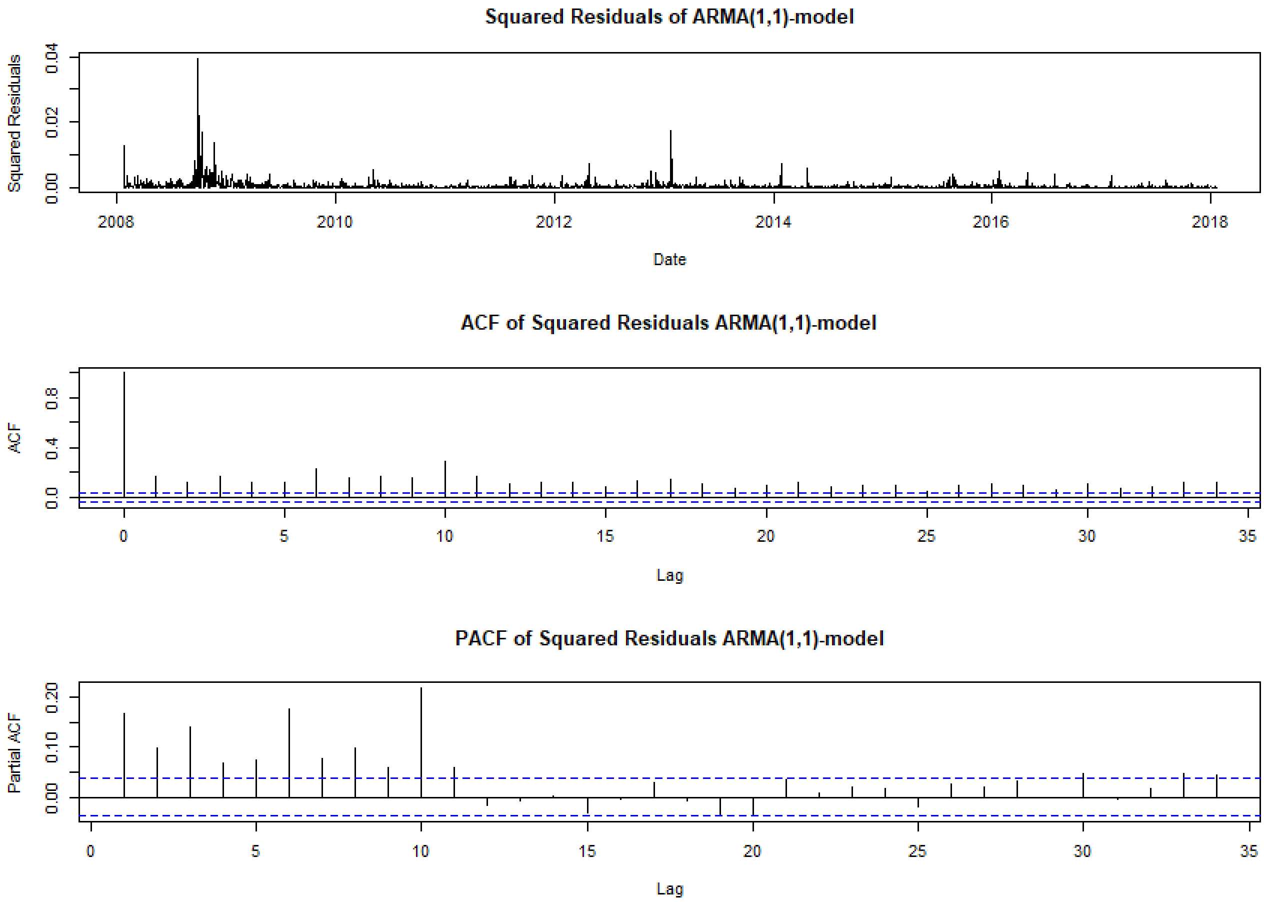

Figure 11: Squared Residuals of the ARMA(1,1)-model

The summary output of the chosen ARMA(1,1)+GARCH(1,1)-model can be seen in the Appendix. All parameter estimates of the GARCH(1,1)-model are significant. The p-value of the Ljung-Box test indicates that the residuals seem to be independently distributed, which is a model assumption. The Jarque-Bera test indicates that the residuals do not seem to have the skewness and kurtosis of a normal distribution, since this assumption is clearly rejected. The full ARMA(1,1)+GARCH(1,1)-model can therefore be expressed as follows:

$$r_t = \alpha + \phi r_{t-1} + w_t + \theta w_{t-1}, \quad w_t | w_{t-1} \sim N(0, \sigma_t^2)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 w_{t-1}^2 + \beta_1 \sigma_{t-1}^2$$

In Figure 12, the one-step predictions using an ARMA(1,1)+GARCH(1,1)-model to predict the Apple stock prices can be seen. A 95% Confidence Interval has been calculated by using the fitted conditional variance for each time step. The Mean Square Error of the predictions compared to actual prices turned out to be 1.606028, so the ARMA(1,1)+GARCH(1,1) model actually yielded better predictions of the Apple Stock price than the standard GBM model. However, the fitted ARMA(1,1)+GARCH(1,1) model uses the entire series of log returns between 2008-2018 to fit the parameters, and then the fitted log returns are then transformed to fitted stock prices as stated in Section 5.8. The one-step predictions of the GBM instead uses a varied time frame before each prediction to estimate the drift and volatility, and then the expected one-step prediction is calculated.

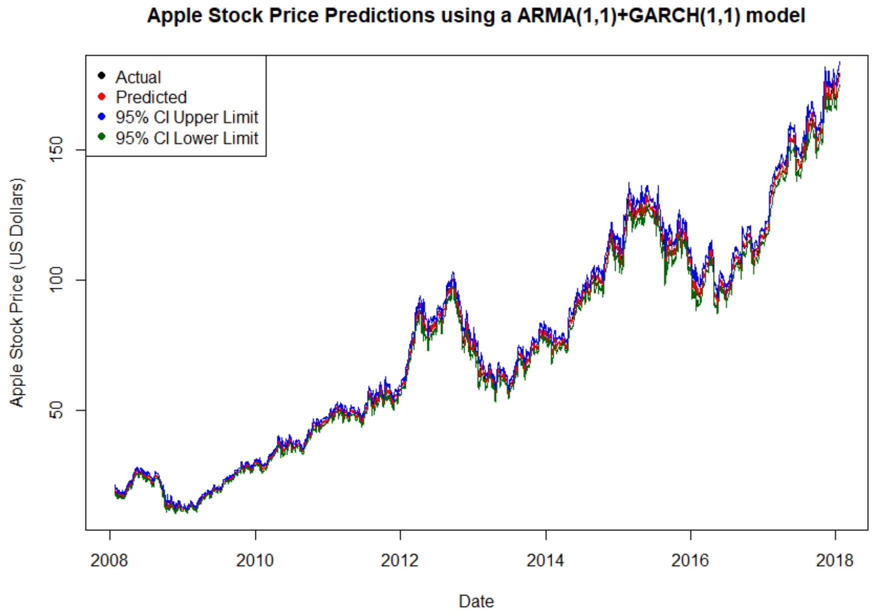

(a) Price Predictions of the Apple Stock

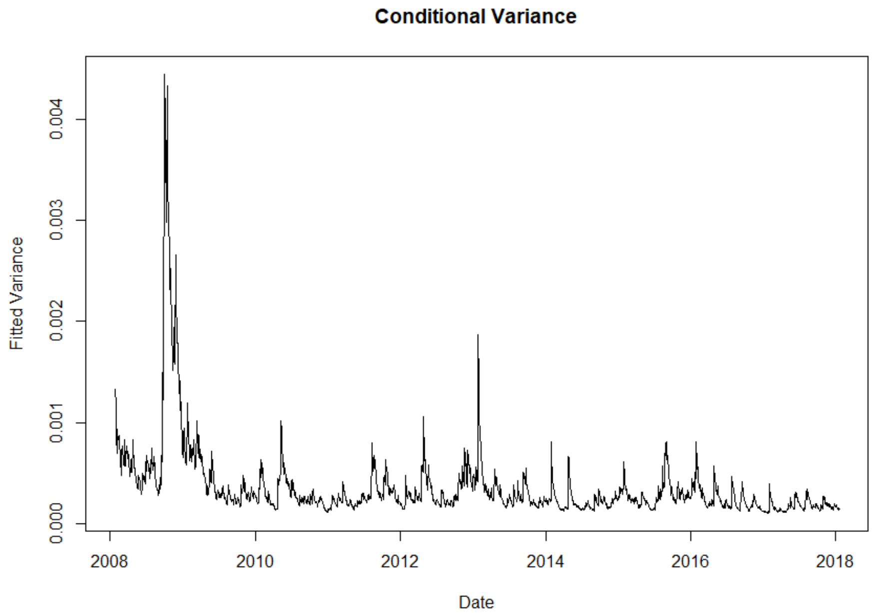

(b) Conditional Variances

Figure 12: Predictions and 95% Confidence Interval using a GARCH model

By looking at the residual plots in Figure 13, it's evident that the residuals of the ARMA(1,1)+GARCH(1,1) model provide a better fit of the normal distribution than the residuals of the ARMA(1,1) model, although both models yield a poor fit in the tails. Therefore using conditional variance modelled by a GARCH(1,1)-model yields a closer resemblance to the model assumptions. The ARMA model assumes a constant volatility during the entire time interval, which is not realistic in financial modelling. Instead, using a mixed model such as the ARMA(1,1)+GARCH(1,1) takes recent market fluctuations into account, and therefore provides a more accurate measure of volatility than the ARMA model.

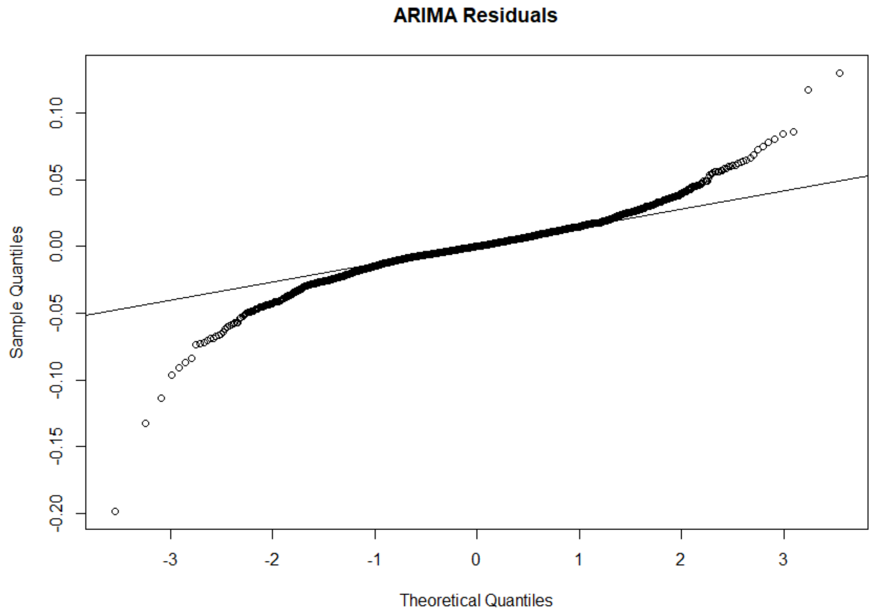

(a) QQ-plot of ARMA(1,1) model

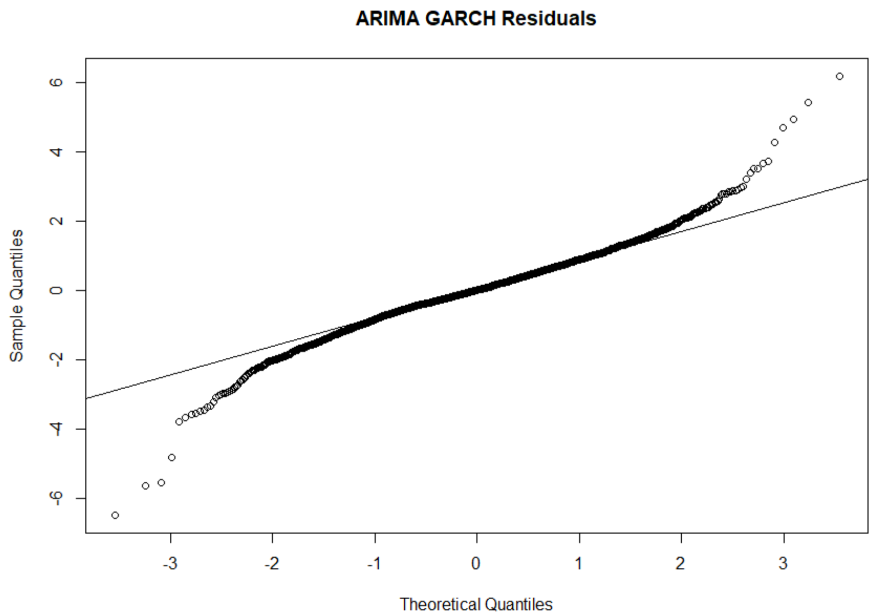

(b) QQ-plot of ARMA(1,1)+GARCH(1,1) model

Figure 13: Residual distribution assumption

6.3 Predicting the S&P500 Index using a Geometric Brownian Motion

6.3.1 Predicting one Time Step Forward

Predicting the S&P500-index by using a GBM with estimated drift assuming normality, a bootstrap estimate of the drift using 10 000 resamples and a Cauchy distribution assumption yielded the results seen in Table 8. The time frame used for estimates has also been varied in order to investigate the impact on the MSE. As can be seen, a time frame of 300 prediction days yielded the lowest MSE for both the normality, bootstrap and Cauchy case. In this case, the standard GBM assuming normally distributed returns performed better than the Cauchy assumption. Using a constant drift for the entire 10 years yielded an MSE of 2.715739 and a probability of accurately predicting an up- or a down move of 0.5466186 for the normality assumption and an MSE of 2.723693 and a probability of 0.5466186 for the Cauchy distribution assumption.

Sample size	$\mathbf{MSE_{Normal}}$	$\hat{\mathbf{p}}_{\mathbf{Normal}}$	$\mathbf{MSE_{Cauchy}}$	$\hat{\mathbf{p}}_{\mathbf{Cauchy}}$	$\mathbf{MSE_{Bootstrap}}$	$\hat{\mathbf{p}}_{\mathbf{Bootstrap}}$
50	2.739841	0.5188596	2.750076	0.5114035	2.740136	0.5188596
60	2.733141	0.5209159	2.744619	0.5178336	2.732843	0.5200352
70	2.695506	0.524838	2.702941	0.5239741	2.696275	0.5231102
80	2.689951	0.5195143	2.695731	0.5208153	2.690083	0.5195143
90	2.686302	0.5263846	2.690995	0.5180986	2.686209	0.5246402
100	2.680179	0.5192644	2.68783	0.5166375	2.680346	0.5188266
200	2.218271	0.5245317	2.21793	0.5396967	2.218228	0.5200714
300	**2.199803**	**0.5304191**	**2.200543**	**0.5443894**	**2.200125**	**0.5308698**
400	2.222911	0.5247758	2.229017	0.5497876	2.222954	0.5261916
500	2.275157	0.5433383	2.280613	0.5482912	2.275017	0.5433383
600	2.283222	0.5398645	2.285661	0.54716	2.282953	0.5398645
700	2.330804	0.5343595	2.336464	0.5492029	2.330855	0.5343595
800	2.40861	0.5386853	2.415768	2.415768	2.408445	0.5398487
900	2.33576	0.5398394	2.338786	0.5484867	2.335388	0.5410747
1000	2.248609	0.5497038	2.254401	0.5470704	2.248957	0.5483871

Table 8: Different outcomes using a varied time frame of the S&P500 index

In Figure 14, the actual S&P500 index has been compared to the predicted price index for data between 2008-2018. A standard GBM with a normality assumption and 300 prediction days has been used to predict the price movements.

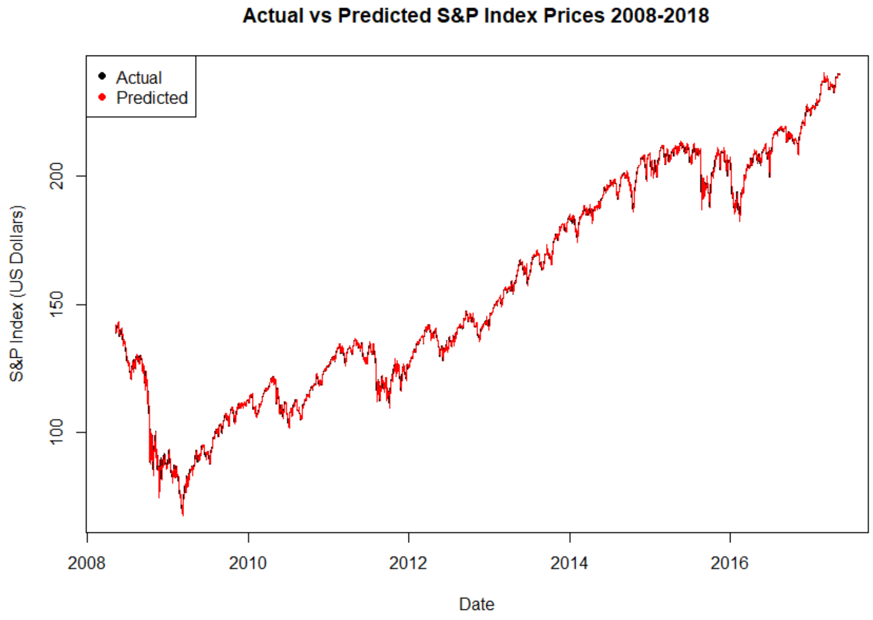

(a) SP500 index prices during 2008-2018

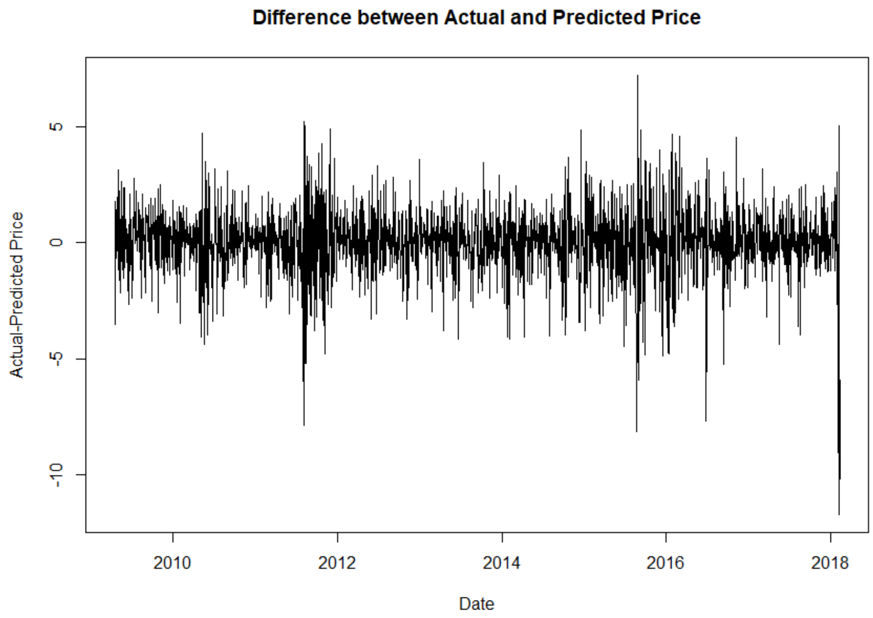

(b) Plot of differences

Figure 14: Comparison of the actual and predicted S&P500 index using a prediction time frame of 300 days

6.3.2 Predicting a Longer Time Frame

Just as in the case of predicting the Apple stock, Monte Carlo simulations have been made in order to predict the prices of the S&P500 index prices during 2008-2018. 1000 Monte Carlo simulations have been made in order to estimate the expectation of the future S&P500 index prices. In this case, a standard GBM using the normality assumption to estimate drift and volatility has been used, as well as a Cauchy distribution assumption. Predictions have also been made using bootstrap estimates of drift and volatility, by using 10 000 resamples of the original data. A Cauchy distribution assumption clearly did not improve the accuracy of the predictions, as can be seen in Figure 15, since the Cauchy location parameter clearly overestimates the drift. However, using bootstrap estimates of the drift and volatility slightly improved the predictions, even though the predictions turned out to be quite similar to the predictions assuming normality.

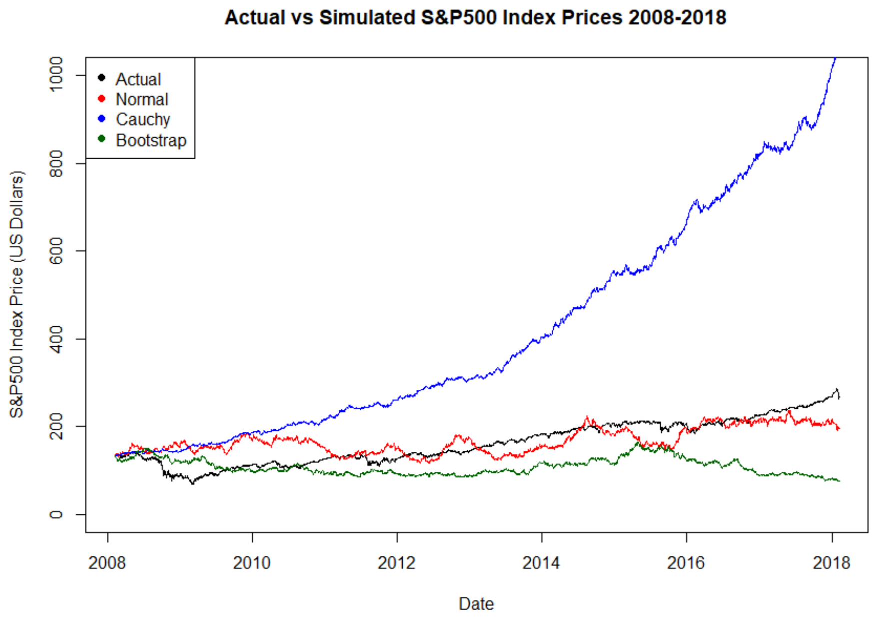

(a) Simulated S&P500 Index prices during 2008-2018

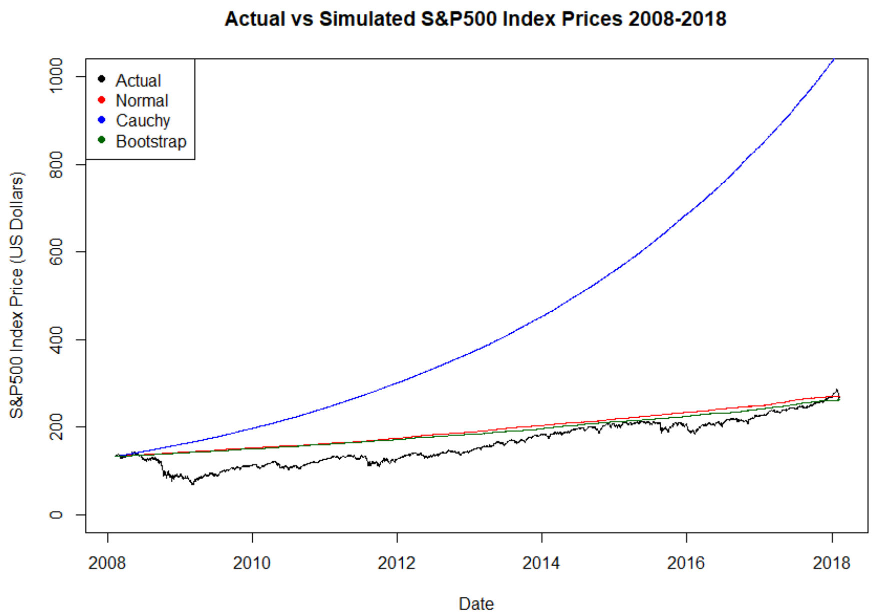

(b) Expected S&P500 Index prices during 2008-2018

Figure 15: Simulated prices assuming constant drift of a standard GBM during 2008-2018

Just as the case of the Apple stock, assuming a constant drift and volatility throughout a long time frame is not valid. In Figure 16, blockwise intervals of equal length have instead been used to estimate drift and volatility in a standard GBM, by estimating drift and volatility by using the sample mean and standard deviation, as well as bootstrap estimates. Decreasing the block length and thus the number of days for estimations yield more accurate results.

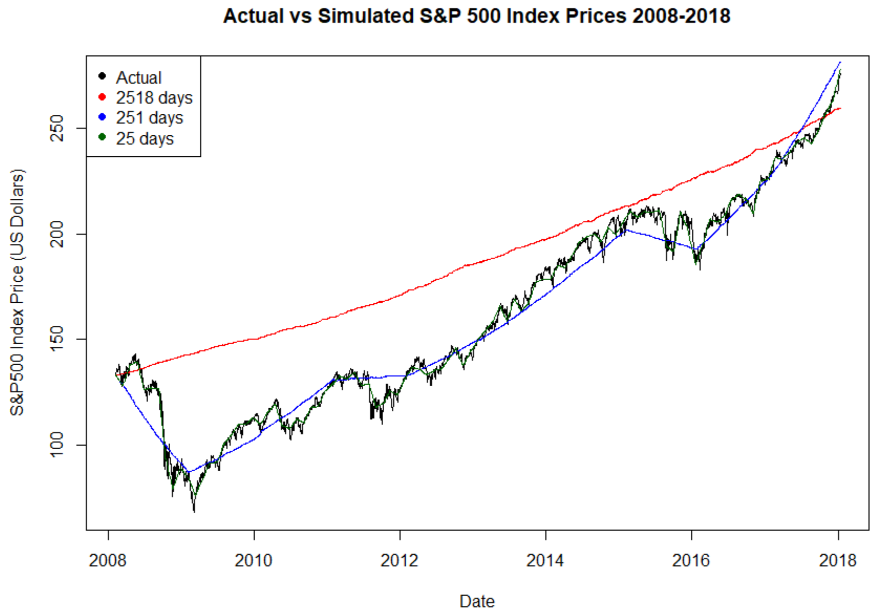

Figure 16: S&P500 index returns during 2008-2018

In Table 9, the results of varying the block length to simulate the S&P500 index prices can be seen. Since the drift and volatility varies greatly between different time intervals, using a smaller block length for estimates greatly reduces the MSE.

Estimation period	MSE GBM	MSE Bootstrap
2518 days	1064.954	1210.584
1259 days	458.3103	427.8717
504 days	194.6356	192.4257
252 days	70.69974	68.45065
126 days	32.4412	31.97207
50 days	15.53927	15.89237
25 days	8.90385	8.735084

Table 9: Different outcomes using a varied time frame of the SP500 index

6.4 Modelling the S&P500 Index using a mixed GARCH model

The historical prices and returns of the S&P500 Index during 2008-2018 can be seen in Figure 3. Similar to the Apple prices and returns, the data shows a larger volatility around the time of the financial crisis, while the mean of the series does not seem to be time dependent. So therefore it's appropriate to model the returns by using a mixed GARCH model, just like in the Apple case.

By looking at the autocorrelation function (ACF) and partial autocorrelation function (PACF) in Figure 4, it's evident that the ACF tails off although several lags show values outside of the bounds. So the S&P500 index data also shows signs of the "long memory" property, with long term dependence.

Fitting an ARMA(p,q)-model to the S&P500 data from 2008-2018 which minimizes the AIC-value yields an ARMA(0,2)-model, suggesting a time series model with two moving average lags and zero mean. The model can be expressed as follows, with parameters estimates according to Table 10, and where r_t is the predicted log return:

$$r_t = \theta_1 w_{t-1} + \theta_2 w_{t-2} + w_t \tag{11}$$

$$w_t | w_{t-1} \sim N(0, \sigma_t^2)$$

Coefficients	θ_1	θ_2	σ^2	**LL**	**AIC**
Estimates	-0.0838	-0.0782	0.0001621	7418.88	-14831.75
Standard Error	0.0199	0.0202	-	-	-

Table 10: Parameter estimates of the ARMA model

The squared residuals of the fitted ARMA(0,2)-model as well as the ACF and PACF plots can be seen in Figure 17. Volatility clustering is present around the time of the financial crisis in 2008, and both the ACF and PACF plot both show lags outside of the bounds. Modelling the volatility by using a GARCH-model, and choosing the model displaying the lowest AIC yielded a GARCH(1,1)-model.

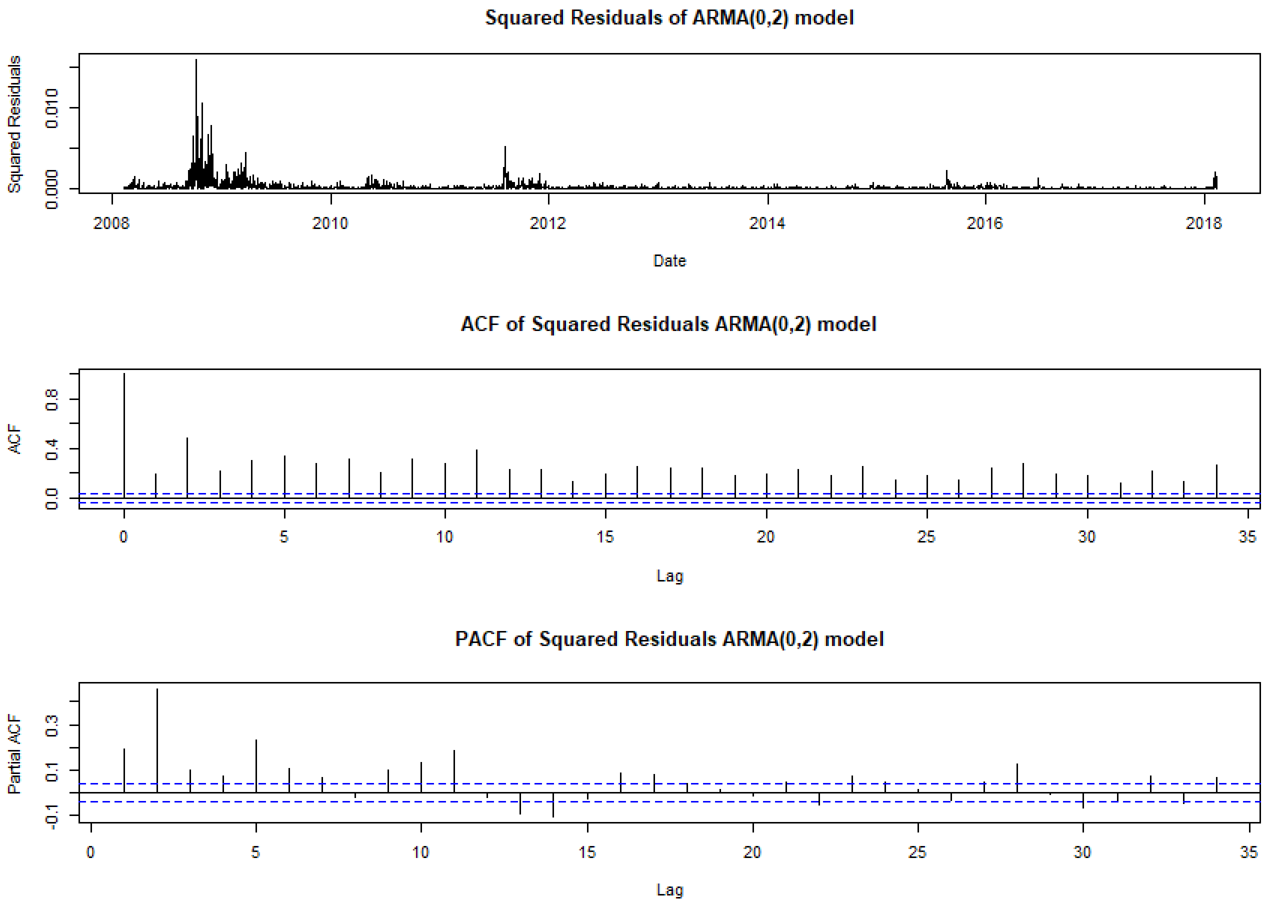

Figure 17: Squared Residuals of the ARMA(0,2)-model

The summary output of the chosen ARMA(0,2)+GARCH(1,2)-model can be seen in the Appendix. All parameter estimates of the GARCH(1,2)-model are significant. The p-value of the Ljung-Box test indicates that there seems to be no dependence in the residuals, which is a model assumption. At least the independence assumption cannot be rejected at any reasonable level of significance. The Jarque-Bera test indicates that the residuals do not seem to have the skewness and kurtosis of a normal distribution, since this assumption is clearly rejected. The full ARMA(0,2)+GARCH(1,2)-model can therefore be expressed as follows:

$$r_t = \theta_1 w_{t-1} + \theta_2 w_{t-2} + w_t, \;\; w_t | w_{t-1} \sim N(0, \sigma_t^2)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 w_{t-1}^2 + \alpha_2 w_{t-2}^2 + \beta_1 \sigma_{t-1}^2$$

In Figure 18, the one-step predictions using an ARMA(0,2)+GARCH(1,2)-model to predict the SP500 index prices can be seen. A 95% Confidence Interval has been calculated by using the predicted conditional variance for each time step. The Mean Square Error of the predictions compared to actual prices turned out to be 2.6576, so the ARMA(0,2)+GARCH(1,2) model did not yield better predictions of the SP500 index price than the standard GBM model. However, when using a mixed ARMA(0,2)+GARCH(1,2) model, the entire 10 years of historical data has been used to estimate the parameters, and the fitted log returns have later been transformed to fitted index prices according to Section 5.8.

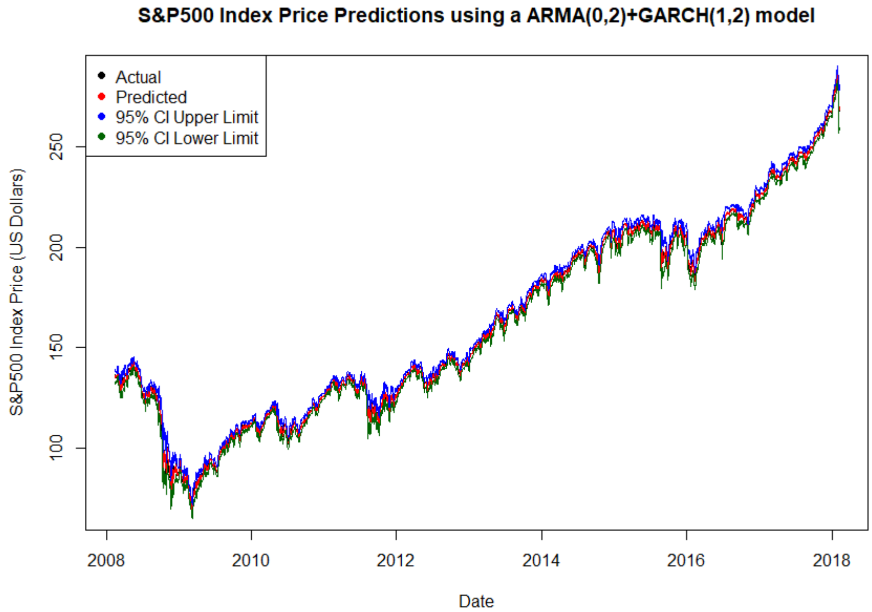

(a) Price Predictions of the SP500 Index Stock

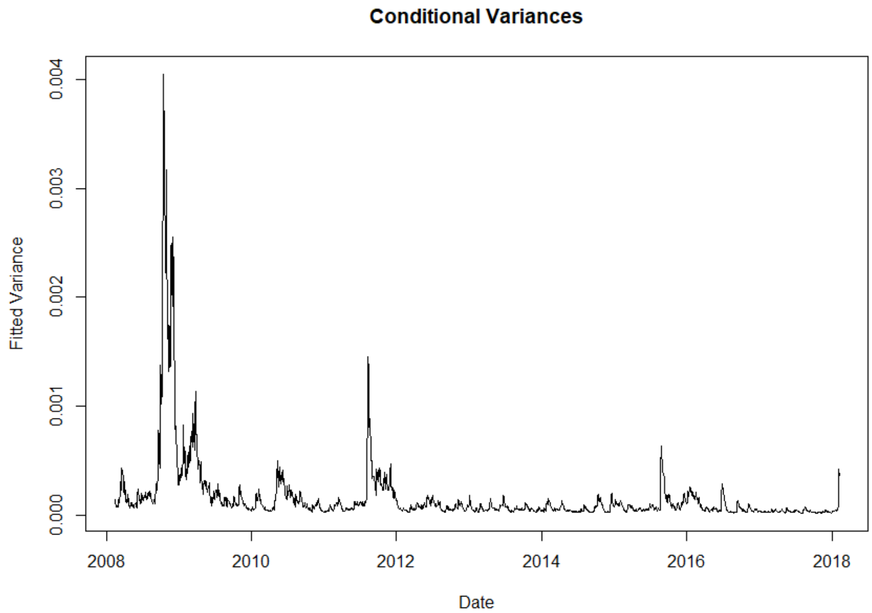

(b) Conditional Variances

Figure 18: Predictions and 95% Confidence Interval using an ARMA(0,2)+GARCH(1,2) model

By looking at the residual plots in Figure 19, it's evident that the residuals of the ARMA(0,2)+GARCH(1,2) model provide a better fit of the normal distribution than the residuals of the ARMA(0,2) model, although both models yield a poor fit in the tails. Therefore using conditional variance modelled by a GARCH(1,2)-model yields a closer resemblance to the model assumptions. The ARMA model assumes a constant volatility during the entire time interval, which is not realistic in financial modelling. Instead, using a mixed model such as the ARMA(0,2)+GARCH(1,2) takes recent market fluctuations into account, and therefore provides a more accurate measure of volatility than the ARMA model.

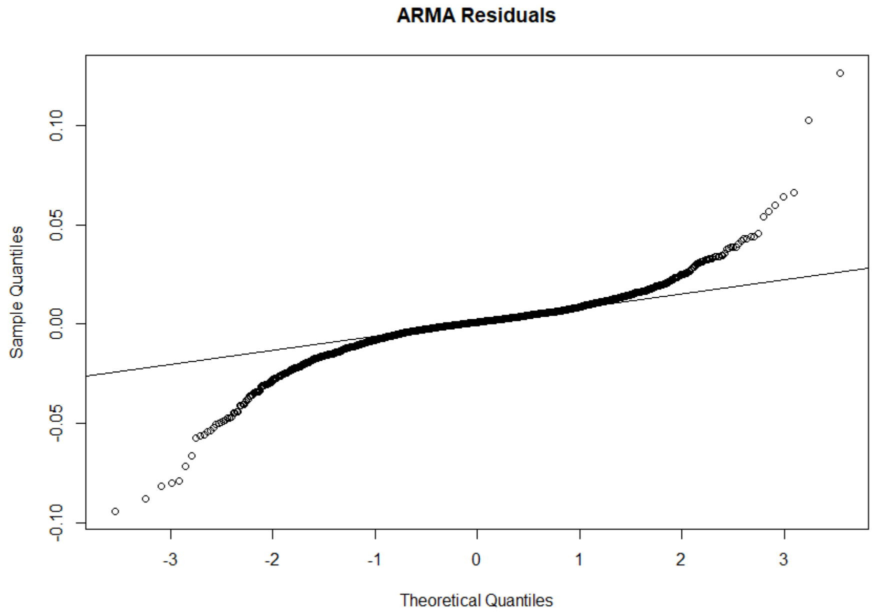

(a) QQ-plot of ARMA(0,2) model

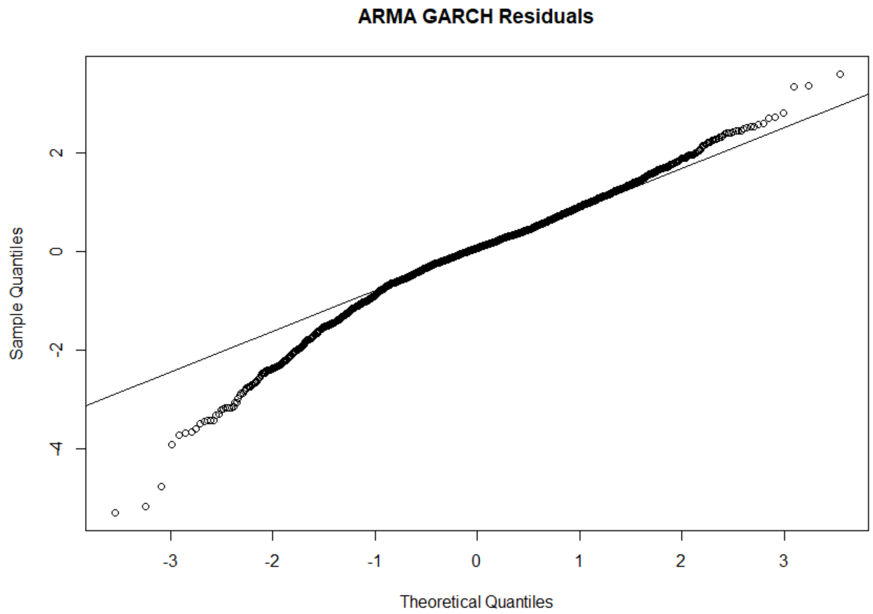

(b) QQ-plot of ARMA(0,2)+GARCH(1,2) model

Figure 19: Residual distribution assumption

6.5 Prediction Model Comparisons

Comparing the one step forward predictions of the Apple stock price, the Geometric Brownian Motion with estimated drift and volatility did not perform better than a mixed ARMA(1,1)+GARCH(1,1) model which yielded a lower Mean Square Error and a higher probability of accurately predicting an up- or a down move. However, in the case of the S&P500 index, a standard GBM model performed better than a mixed GARCH model, which yielded a lower MSE and a higher prediction probability. The results can be seen in Table 11. The one step predictions of the mixed ARMA(p,q)+GARCH(r,s) models are based on the entire history of log returns between 2008-

2018, using fitted log returns by the model to estimate the drift, and then one step predictions are found according to Section 5.8. Predictions by using a GBM have been made by estimating the drift and volatility with a varied time frame prior to the forward prediction, and then the estimated drift has been used to predict the forward price according to Section 6.3.1 and 6.1.1.

Prediction Model	**Data**	**MSE**	$\hat{\mathbf{p}}$
Standard GBM	Apple Stock	1.66206	0.5174939
ARMA(1,1)+GARCH(1,1)	Apple Stock	1.606028	0.5266296
Standard GBM	S&P500	2.199803	0.5304191
ARMA(0,2)+GARCH(1,2)	S&P500	2.6576	0.5031771

Table 11: Outcomes of prediction models

As can be seen in Table 11, both the standard GBM and the mixed ARMA(p,q)+GARCH(r,s) models yielded probabilities of accurately predicting an up- or down move by slightly more than 50%, suggesting some predictive power based on the 10-year history of the closing prices for both the Apple stock and the S&P500 index, consisting of about 2500 observations each. However, in order to investigate the accuracy of the models further, an even larger sample size would preferably be used since the edge shown by the models is quite small.

7 Conclusions

Both the financial time series data of the Apple stock price and the S&P500 index show evidence of the "long memory" property, meaning that there is a weak dependence between historical lags, and that present fluctuations in the price will have an impact on the future fluctuations of the price. The ACF and PACF both tail off slowly for both the Apple stock data and the S&P500 index, which can be seen in Figure 2 and Figure 4, showing evidence of long term dependence. The normality assumption of the log returns is clearly not valid, which is also consistent with the results of B. Mandelbrot (1963) [3] and G. Dhesi et al (2016) [4]. Therefore a Cauchy distribution was fitted as well as a kernel density estimator (KDE), using a Gaussian kernel function. The kernel density yielded the best fit based on the Chi-Square goodness-of-fit tests, which can be seen in Table 3. Both the normality and the Cauchy distribution assumption were both rejected, while the KDE was accepted.

Which can be seen in Section 5.2, a nonparametric standard bootstrap method has been used in order to generate a distribution of the means and variances for both the Apple stock and the S&P500 index. The bootstrap distributions of means can be seen in Figure 7. The standard bootstrap method relies on the i.i.d. assumption, and since the dependence between lags is quite weak, the standard bootstrap method has been used even though the i.i.d. assumption is not quite valid. Using 10 000 bootstrap resamples in order to generate a distribution of mean and variance compared to simply using the sample mean and variance did not improve the predictions generally, since the estiamate of the drift is the same as by simply using the sample mean, although the standard bootstrap method yielded narrower confidence intervals for small sample sizes, as can be seen in Table 4. The bootstrap estimates of the drift are approximately the same as by using the sample mean, and therefore using a standard bootstrap method also yielded approximately the same mean square errors, which can be seen in Table 6 and 9.

The one time step forward predictions of the Apple stock price and the S&P500 index yielded probabilities of predicting a correct up- or down move of a bit more than 50%, so the GBM with drift has some predictive power based on a sample size of more than 2500 price observations, which is also consistent with the results of K. Reddy and V. Clinton (2016) [1]. The one time step predictions yielded the best results when using 60 prediction days for the Apple stock as can be seen in Table 5, and 300 prediction days for the S&P500 index as can be seen in Table 8. Assuming a Cauchy distribution assumption and using the location and scale parameter to estimate drift and volatility clearly did not improve the predictions, which is evident when using Monte Carlo methods to find the expected stock price during 2008-2018 for both the Apple stock price as well as the S&P500 index which can be seen in Figure 9 and the S&P500 index which can be seen in Figure 15.

8 References

References

[1] Reddy, Krishna and Clinton, Vaughan, Simulating Stock Prices Using Geometric Brownian Motion: Evidence from Australian Companies (2016), https://pdfs.semanticscholar.org/6778/2faf27ffeb8acd203ad3e270ee8dfb831e9e.pdf. 2018-01-29.

[2] http://www.nasdaq.com/symbol/aapl/historical. 2018-01-18.

[3] Mandelbrot, Benoit (1963). The Variation of Certain Speculative Prices. The Journal of Business, Vol. 36, No. 4, pp. 394-419. The University of Chicago Press. https://web.williams.edu/Mathematics/sjmiller/public_html/341Fa09/econ/Mandelbroit_VariationCertainSpeculativePrices.pdf. 2018-01-29.

[4] Gurjeet Dhesi, Muhammad Bilal Shakeel, Ling Xiao (2016). Modified Brownian Motion Approach to Modeling Returns Distribution. School of Business, London South Bank University. https://arxiv.org/ftp/arxiv/papers/1507/1507.02203.pdf. 2018-01-29.

[5] Björk, Tomas (2009). Arbitrage Theory in Continous Time. Oxford, New York: Oxford University Press.

[6] Shumway, Robert H., Stoffer, David S (2011). Time Series Analysis and its Applications. New York: Springer.

[7] Cogneau, Philippe. Zakamouline, Valeri (2010). Bootstrap Methods for Finance: Review and Analysis. University of Liege. http://www.seminar.hec.ulg.ac.be/docs/Sem21.10.10_Cogneau.pdf. 2018-04-14.

[8] Cont, Rama. Long range dependence in financial markets, Ecole Polytechnique, France http://www.cmap.polytechnique.fr/~rama/papers/FE05.pdf. 2018-04-14.

[9] Zhuanxin Ding, Clive W.J. Granger and Robert F. Engle (1993). A long memory property of stock market returns and a new model, University of California, San Diego. https://pdfs.semanticscholar.org/945e/979486da21ca20b61339dd38546fff8bcb60.pdf. 2018-04-14.

9 Appendix

9.1 Apple Stock Price

The R summary output of the ARMA(1,1)-model of the Apple log returns can be seen below.

```
Call:
arima(x = HistoricalQuotes$returns, order = c(1, 0, 1))

Coefficients:
         ar1      ma1  intercept
      0.3687  -0.3832      8e-04
s.e.  0.5095   0.5094      4e-04

sigma^2 estimated as 0.0003801:  log likelihood = 6339.48,  aic = -12670.95

Training set error measures:
                        ME      RMSE         MAE MPE MAPE      MASE        ACF1
Training set -2.833982e-07 0.0194949 0.01339335 NaN  Inf 0.7019531 0.007475008
```

The R summary output of the GARCH(1,1)-model of the Apple log returns can be seen below.

```
Call:
garch(x = arima_model$residuals, order = c(1, 1), trace = F)

Model:
GARCH(1,1)

Residuals:
      Min        1Q    Median        3Q       Max
-6.487297 -0.523096  0.004813  0.588502  6.173866

Coefficient(s):
    Estimate  Std. Error  t value Pr(>|t|)
a0 1.129e-05   1.584e-06    7.129 1.01e-12 ***
a1 8.484e-02   9.667e-03    8.776  < 2e-16 ***
b1 8.815e-01   1.343e-02   65.634  < 2e-16 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Diagnostic Tests:
Jarque Bera Test

data:  Residuals
X-squared = 1318.4, df = 2, p-value < 2.2e-16


Box-Ljung test

data:  Squared.Residuals
X-squared = 0.00020317, df = 1, p-value = 0.9886
```

9.2 S&P500 Index Price

The R summary output of the ARMA(0,2)-model of the S&P500 log returns can be seen below.

```
Call:
arima(x = SP500$returns, order = c(0, 0, 2))

Coefficients:
          ma1      ma2  intercept
       -0.085  -0.0791      3e-04
s.e.    0.020   0.0202      2e-04

sigma^2 estimated as 0.000162:  log likelihood = 7412.93,  aic = -14817.86

Training set error measures:
                       ME       RMSE         MAE MPE MAPE      MASE         ACF1
Training set 2.583208e-06 0.01272627 0.008002666 NaN  Inf 0.6706361 -0.001882902
```

The R summary output of the GARCH(1,2)-model of the S&P500 log returns can be seen below.

```
Call:
garch(x = arma_SP500$residuals, order = c(1, 2))

Model:
GARCH(1,2)

Residuals:
    Min      1Q  Median      3Q     Max
-5.1817 -0.4730  0.1110  0.6341  3.6639

Coefficient(s):
    Estimate  Std. Error  t value Pr(>|t|)
a0 2.817e-06   3.725e-07    7.562 3.97e-14 ***
a1 8.728e-02   1.108e-02    7.878 3.33e-15 ***
a2 7.713e-02   1.890e-02    4.080 4.50e-05 ***
b1 8.176e-01   1.597e-02   51.187  < 2e-16 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Diagnostic Tests:
Jarque Bera Test

data:  Residuals
X-squared = 431.94, df = 2, p-value < 2.2e-16


Box-Ljung test

data:  Squared.Residuals
X-squared = 0.89521, df = 1, p-value = 0.3441
```