Autoregressive Conditional Heteroskedasticity Models of Volatility in Asset Pricing

Simon Beyzerov[1,2]

svbeyzerov@wpi.edu

February 2023

Abstract

In this paper, we explore the development, identification, and motivations behind a fundamental class of time series models: autoregressive conditional heteroskedasticity (ARCH) models. First introduced by Robert Engle in his seminal work, ARCH models have given way to a rich class of econometric inquiry. Concerned with conditional-level fluctuations in variance, this unique class of models allows us to analyze dependence in residual series data. Finding particular applications within financial time series, ARCH modeling has shown continued success in risk analysis and volatility forecasting. They have especially seen applications within asset pricing, where risk is a central concern in making optimal financial decisions. We examine these considerations and more while discussing the motivations behind such models. By looking at the historical returns of one asset in particular, we provide a demonstrative analysis of the ARCH modeling process. We identify distinct effects in our data that suggest an ARCH process and go on to introduce a parsimonious model that fits our data well and forecasts future volatility.

[1] Worcester Polytechnic Institute
[2] Massachusetts Academy of Math & Science

1 Introduction

Risk is a central concern in portfolio optimization and a fundamental financial (and general) decision-making principle. At its core, it's the simple observation that an action may or may not result in an undesirable outcome—an inescapable reality that drives our unpredictable world and governs our decisions. Financial markets are no exception; in finance, we simplify the concept of risk to a binary outcome: the potential gain (reward) or potential loss in value of an asset or portfolio. This simplified, unambiguous view of a "desirable" and "undesirable" outcome has become the central paradigm of financial decision-making.

Optimal decision-making naturally encourages us to consider risk. In a set of possible actions, some provide larger reward while carrying the same level of potential loss, and thus are preferred—given the choice, we wish to maximize our gain and minimize our losses. But perfectly measuring *potential* movements is fundamentally difficult, if not impossible. To understand these systems, we need models that can estimate risk; we wish to use past observations to inform our future decisions. This simple challenge has been a central focus of econometricians and practitioners for the better part of a century.

Although many different ways of measuring risk have been considered, the variance of an underlying asset has been a consistent metric found to statistically quantify risk. First formally introduced in the Capital Asset Pricing Model (CAPM) by Sharpe (1964) and further explored by Black and Scholes (1973) in their pricing of options, the relationship between risk and variance would come together into the concept of *volatility*, computed from the square-root of the asset's sample variance. Typically, the *realized* volatility encompasses the notion of unconditional variance and is computed over a fixed period. But true volatility is theorized to be latent and not directly observable (Andersen et al., 2005). Many attempts have been made to use unconditional variance to approximate instantaneous volatility.

In time-series models, unconditional variance is assumed to be constant for stationarity. But while this realized volatility is constant over a time period, in the residuals of some time series, it was observed to fluctuate in clusters (B. Mandelbrot, 1963; B. B. Mandelbrot, 1997). It was clear that there were some fluctuations in *conditional* variance over the series, while the overall realized variance was fixed. This inherent heteroskedasticity at the conditional level was particularly common in financial systems and asset prices. A consistent model of dynamic risk forecasting was introduced in the seminal work of Engle (1982). Towards modeling inflation in the United Kingdom, Engle considered autoregressive conditional heteroskedasticity as opposed to traditional constant one-period volatility forecasts. The resulting ARCH model generalized sample variance into a time-varying process while allowing for straightforward parameter estimation from historical observations.

The ARCH model would spark a rich class of econometric research, particularly in finance, where risk is of particular interest. As evidenced by the many ARCH-type models that would arise, it was clear that Engle's approach was pivotal in assessing the volatility of an underlying asset and forecasting uncertainty in its price. In this work, we consider the application of ARCH-type modeling on one asset. Our goal is to demonstrate the structure of this powerful, fundamental model.

Paper organization In Section 2, we describe the general motivation behind time-varying volatility modeling in financial time series. We then go on to introduce the structure, strengths, and weaknesses of classic ARCH/GARCH models. Towards understanding the application of these two classic models in financial forecasting, we consider the daily returns of the NASDAQ Composite Index in Section 3. Following an initial data description and analysis, we outline the model-building process for both ARCH and GARCH models, discussing their differences along the way. In Section 4 we consider the ARCH/GARCH models proposed in Section 3 and provide relevant diagnostics. These models are then used to forecast the future volatility of the NASDAQ Composite index. Finally, in Section 5 we consider the forecasting performance of our proposed models and provide a discussion of our approaches. All references and appendices can be found at the end of our report—all code and links to outside repositories are included therein.

2 Background

Engle's starting point was a residual time series without significant serial correlation, but clear autocorrelation in its squares. This sort of behavior frequently appeared in residual series of financial models and suggested some underlying non-random behavior—some dependence in the series despite an absence of serial correlation. His approach was to model the variance of this residual series as a time-varying, autocorrelated process of lagged variance values (which he refers to as a type of "weighted variance" (Engle, 2004)).

Characteristics & Identification Prior to identifying any ARCH effects, we need to shift our time series into one that fits Engle's initial archetype: a lack of serial correlation but clear dependence. To achieve this, we typically remove any serial correlation with an ARMA model before analyzing the resulting residual series. A residual series exhibiting a strong ARCH effect will typically exhibit three main characteristics: *unpredictability*, *volatility clustering*, and *leptokurticity* (Engle, 1982, 2004). Loosely, these characteristics suggest a series where high and low volatility is sustained in periods, before reverting to a baseline level. When identifying an ARCH process, we search for evidence of time-varying volatility that is clustered and fat-tail.

Model Construction Once identification is successful, we must estimate the model parameters and orders. Although we will not go into depth regarding parameter estimation algorithms, we note that once an ARCH process is identified it is relatively difficult to estimate the order of an ARCH-type model (particularly for ARCH and GARCH models). In practice, it has been shown that low-order ARCH models perform just as well (or better) at forecasting than high-order models (Jafari et al., 2007). We elaborate on this further in Section 4.1.

Conditional Mean & Variance Taking a step back, suppose we have derived a time series r_t from a series of observations $\{P_t\}$. In our analysis, r_t will concern the log returns of an asset while P_t is the observed price of the asset at time t. We say $\mathcal{I}_t$ is the information set up to time t. Intuitively, we can consider $\mathcal{I}_t = \{P_1, \ldots, P_t\}$, or the set of all observations up to time t. In the literature, the information set is also seen as the set of all linear functions of past prices. Towards constructing an ARCH-type model, we typically decompose the base process into a mean process μ_t and innovation process ϵ_t,

$$r_t = \mu_t + \epsilon_t \tag{2.1}$$

or alternatively through conditional mean μ_t and variance σ_t^2 of r_t:

$$\mu_t = E(r_t \mid \mathcal{I}_t), \quad \sigma_t^2 = \mathrm{Var}(r_t \mid \mathcal{I}_t), \tag{2.2}$$

where ϵ_t is not serially correlated. As we've discussed previously, an ARCH model's primary focus is on describing the conditional variance term σ_t^2 as a time-varying series. Since Engle's base ARCH assumption is that ϵ_t is serially uncorrelated, the mean process μ_t should capture all serial correlation in the original series. Although the mean process μ_t will not be explicitly discussed in this report, in practice it is typically an ARMA(p, q) process. In our final models, we will explicitly state the ARMA process selected for the conditional mean.[1]

2.1 Autoregressive Conditional Heteroskedasticity Model (ARCH)

An Autoregressive Conditional Heteroskedasticity Model (ARCH) of order p ARCH(p) is one that expresses the conditional variance of ARCH process ϵ_t as an AR process:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \epsilon_{t-i}^2. \tag{2.3}$$

The ARCH model attempts to model the short-term instability of the underlying residual series through its conditional heteroskedasticity. In its base form, the ARCH model places equal impact on both positive and negative volatility periods—the impact in both directions is symmetric due to the square of the lags. Typically the residuals of the resulting model are assumed to follow a normal distribution, but different likelihood estimation algorithms may assume a different residual distribution.

2.2 Generalized-ARCH Model (GARCH)

A Generalized Autoregressive Conditional Heteroskedasticity Model (GARCH) of orders p and q GARCH(p, q) is one that expresses the conditional variance of ARCH process ϵ_t as an ARMA process:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^{q} \beta_i \sigma_{t-i}^2 \tag{2.4}$$

The GARCH model was presented by Bollerslev (1986) to further refine the modeling of conditional variance of an ARCH process. Although GARCH models are typically more parsimonious than ARCH models, they still fail to capture asymmetric shocks in volatility and suffer the same general weaknesses of the ARCH model.

3 Data and Model Specification

Data Description In order to demonstrate the use of ARCH modeling, we consider asset price data in our analysis. Specifically, we look at the monthly log returns of the NASDAQ composite index. Our data set was obtained from Yahoo Finance and contained the asset's monthly price from January 1985 to February 2023. We then fit our model on the log returns of the price data over one lag, a common transformation to achieve stationarity and to analyze volatility effects. In the rest of our report, we denote this log-return series as r_t and the original observed price data series as P_t.

3.1 Exploratory analysis

We illustrate both the monthly NASDAQ prices and log returns in Figure 1, where we note that the original price series is clearly non-stationary and has a visible trend. The log-return

[1] We refer to the mean process/innovation process and conditional mean/conditional variances interchangeably throughout this report, which we try to make clear from context.

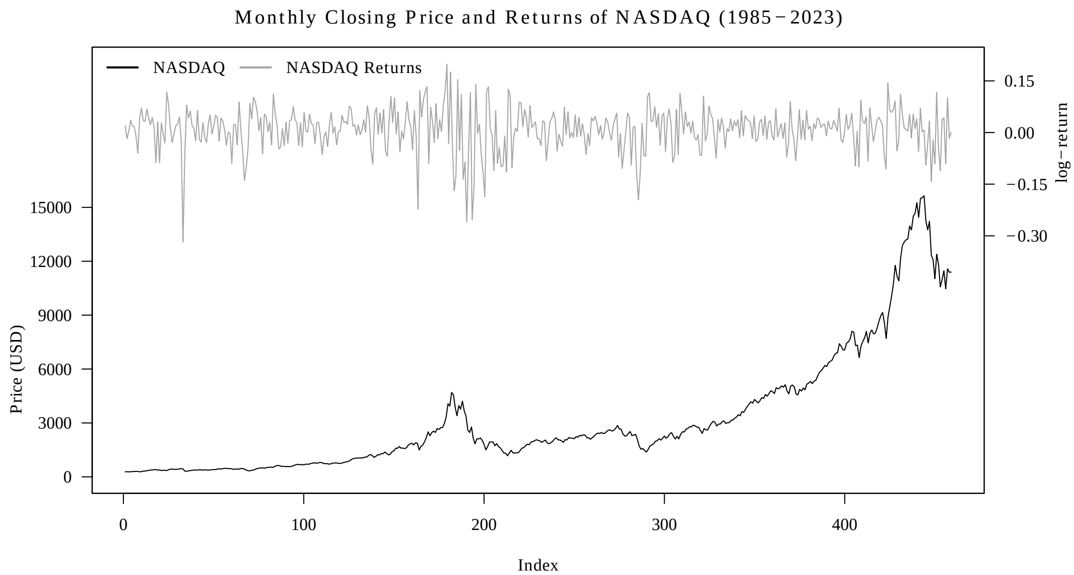

Figure 1: Daily asset price data of the NASDAQ Composite Index from 1985 to 2023: (bottom) daily price of NASDAQ, (top) daily log return of the asset.

series appears stationary; we perform an Augmented Dickey-Fuller test in Table 1 to verify stationarity. We see that the price series exhibits small jumps and falls in value over the course of the period. Oftentimes, these align with periods of high variance in the log-returns plot. This evidence of underlying volatility fluctuations urges us to use non-traditional time-series models.

Mean Process Isolation From our initial suspicion of an ARCH-type model, we attempt to isolate the log-return series into mean and innovation components. Observing the correlograms of the log-returns series in Figure 2, it appears there is no significant serial correlation in the series. This suggests we do *not* need to isolate a mean process, given that we confirmed the process to be stationary in Table 1. We go on to check if the log-returns series is an ARCH process, in absence of a mean series.

Table 1: Results of Augmented Dickey-Fuller Test (without trend) on log-returns of NASDAQ Composite Index, confirming that the series is stationary.

Test Statistic	Critical Value (1%)	Critical Value (5%)	p-value
-7.1274	-3.43	-2.86	< 0.01

Volatility Clustering The first property of an ARCH process we must identify is clustered volatility. In practice, clustering is typically identified empirically or visually by looking at the series. In our case, we identify clear clusters of high and low volatility in the log-returns series. The most identifiable clusters occurred around indices 40, 200, 300, and 450, as seen in Figure 1.

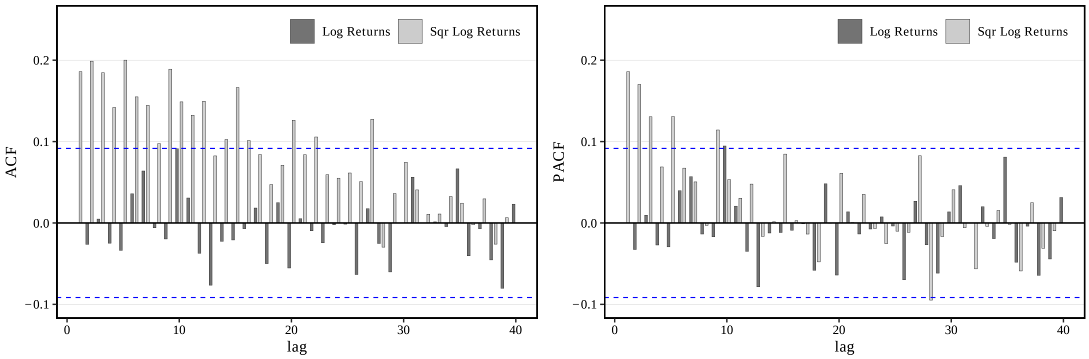

Figure 2: Correlelogram of the monthly daily returns of the NASDAQ Composite Index from 1985-2023: (left) ACF of monthly log return and squared monthly log return, (right) PACF of monthly log return and squared monthly log return.

Conditional Heteroskedasticity Next, we identify if the series is unpredictable and if there is evidence of conditional heteroskedasticity. Visually, it does appear there are fluctuations in volatility in Figure 1, but we take a look at the correlograms for the log-returns and squared-log return series in Figure 2. As before, we observe no serial correlation in the log-returns series. There appears to be a significant serial correlation in the series of square-log returns, as evidenced by the sample ACF and PACF. This autocorrelation in the squared series further supports an ARCH process.

Leptokurticity Finally, an ARCH process should be leptokurtic—peaks and troughs should frequently appear in the data. In our case, we look at the descriptive statistics in Figure 3 and conclude that the log-returns series is in fact leptokurtic. The Normal Q-Q plot (along with the large kurtosis) suggests the series distribution is fat tails. From the exploratory analysis above, we conclude that the log-returns series exhibits ARCH behavior. We proceed with our ARCH-type models. Recall that we did not need to isolate any ARMA mean process for this data sample and thus proceeded to verify the log returns themselves behave as an ARCH process.

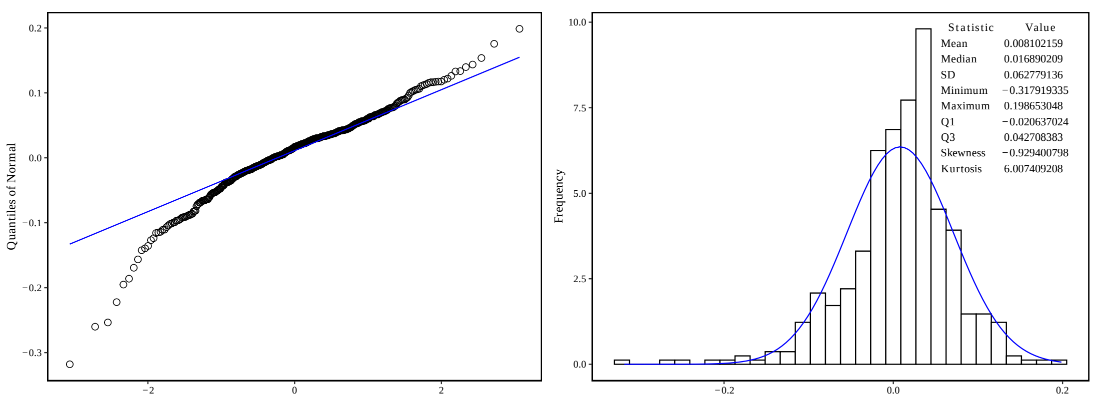

Figure 3: Summary statistics of NASDAQ daily returns from 1985-2023: (left) Normal Q–Q plot of daily log returns series, (right) descriptive statistics of daily log returns series.

Table 2: Model comparison on NASDAQ historic monthly log returns, fitted under normal and Student's t-distribution residual assumptions.

Mean Process	Innovation Process	Normal distribution AIC	Normal distribution BIC	Student's t-distribution AIC	Student's t-distribution BIC
ARMA(0,0)-	ARCH(1)	-2.795142	-2.768110	-2.879634	-2.843592
	ARCH(2)	-2.826293	-2.790251	-2.902406	-2.857353
	ARCH(3)	-2.850076	-2.805023	-2.929181	-2.875117
	GARCH(1,1)	-2.864962	-2.828919	-2.950039	-2.904986
	GARCH(1,2)	-2.859456	-2.814403	-2.944416	-2.890353
	GARCH(2,1)	-2.859487	-2.814434	-2.945114	-2.891050

4 Empirical Analysis & Results

4.1 Model selection

After identifying an underlying ARCH process in the residual series, we must establish the order of our final (G)ARCH model. As noted in Section 2, the ARCH portion of our process is modeled separately from the mean and thus we must also identify the order of the underlying mean ARMA process. While finding the order of the mean series is straightforward with traditional correlogram techniques, finding the order of ARCH processes is not as easy. In practice, once ARCH effects have been identified, low-order models have been found to perform just as well (or better) than high-order models during forecasting (Hansen & Lunde, 2001; Jafari et al., 2007). We point the reader to literature that proposes alternate techniques (Bollerslev, 1986; Hughes et al., 2004).

In Section 3.1, we noted that our log-returns data does not exhibit any serial correlation and is stationary. We noted that isolating a mean process was not necessary and accordingly analyzed the ARCH effects directly. This means we do not need an ARMA process to isolate the mean and thus our model will be of the form ARMA(0,0)-ARCH(m) or ARMA(0,0)-GARCH(m, s).

Toward a more nuanced analysis of the ARCH portion of the model, we consider both ARCH and GARCH models with varying orders. We establish ARCH(1), ARCH(2), and ARCH(3) models based on the high serial correlation in the squared log returns, along with GARCH(1,1), GARCH(2,1), and GARCH(1,2) models. It has been found that even first-order GARCH models like GARCH(1,1) perform considerably better than high-order alternatives (Hansen & Lunde, 2001; Jafari et al., 2007).

Parameter estimation In this report, we do not outline estimation techniques for ARCH models. We point to existing literature that provides a comprehensive summary of the likelihood algorithms frequently entertained (Bauwens et al., 2003). We do note that ARCH models are frequently instantiated with varying assumptions on the standardized residuals of the final model. Most commonly these are the normal distribution and Student's t-distribution (Chatfield & Xing, 2019; Tsay, 2010). In our analysis, we compare model results for both residual distributions and report our results for both distribution assumptions. Additionally, we would like to emphasize that for data where a mean process is to be isolated, parameter estimation is performed jointly over the mean (ARMA) and residual (ARCH) processes. That is to say, we typically obtain weaker results from two-step estimates that compute the mean and residual model parameters separately (Tsay, 2010).

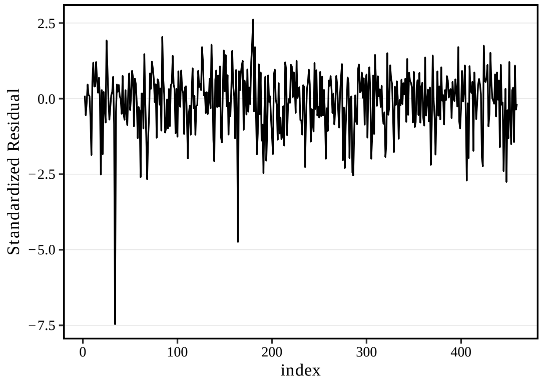

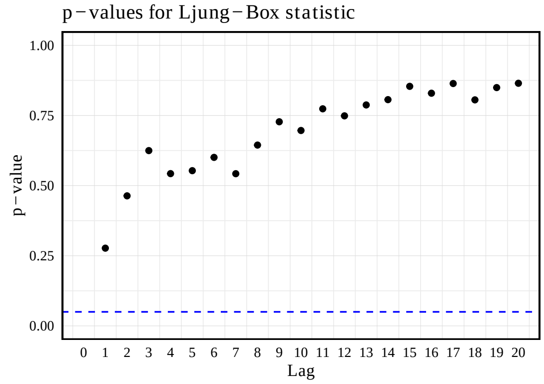

Figure 4: Standardized residuals of ARMA(0,0)-GARCH(1,1) model under Student's t-distribution: (left) standardized residual time series, (right) Ljung-Box statistic to indicate lack of significant serial correlation in the standardized residual series. Further analysis is provided in Table 5.

4.2 Model Diagnostics & Residual Analysis

In Table 2 we report the performance of our six candidate models. Based on these preliminary diagnostics, we select our GARCH(1,1) as the most parsimonious model with the best overall performance. We specifically choose the GARCH(1,1) model under an assumed Student's t-distribution, as this model performed notably better than those under a normally distributed residual assumption.

A summary of this fitted model is shown in Table 3, indicating the specific parameters used in the model. We provide standardized residual analysis in Figure 4, indicating there is no serial correlation. We provide further analysis in Table 5 that confirms our model is appropriately fitted; we show that the standardized residuals exhibit no ARCH effects and fit the assumed Student's t-distribution well. Combined with the constant mean process we identified in Section 3.1, we use the final ARMA(0,0)-GARCH(1,1) model under Student's t-distribution to forecast future volatility in the series of log-returns of the NASDAQ Composite Index. Our results are shown in Table 4.

Table 3: Estimation for GARCH(1,1) in log-returns of the NASDAQ Composite Index, fitted under Student's t-distribution.

Parameter	Estimate	Std. Error	t-value	p-value
μ	0.0132687	0.0021841	6.075	< 0.001
α_0	0.0002773	0.0001342	2.066	0.038802
α_1	0.2218715	0.0706947	3.138	0.001698
β_1	0.7235482	0.0805776	8.980	< 0.001
ν (shape)	6.3768236	1.6699582	3.819	0.000134

Table 4: Forecast of NASDAQ log-returns with a ARMA(0,0)-GARCH(1,1) model under Student's t-distribution and a Normal distribution.

Month	Mean Forecast	Mean Error	Std.	Lower	Upper
3-2023	0.0132687	0.06466891	0.06466891	-0.1159727	0.1425101
4-2023	0.0132687	0.06504719	0.06504719	-0.1167287	0.1432661
5-2023	0.0132687	0.06540282	0.06540282	-0.1174394	0.1439768
6-2023	0.0132687	0.06573727	0.06573727	-0.1181078	0.1446452
7-2023	0.0132687	0.0660519	0.0660519	-0.1187366	0.145274
8-2023	0.0132687	0.06634799	0.06634799	-0.1193283	0.1458657

5 Summary and Conclusions

In this paper, we consider the historical prices and returns of the NASDAQ Composite Index (from 1985-2023) as a demonstrative case study of ARCH modeling. More broadly, we studied how to identify and classify ARCH effects within asset prices that we couldn't capture with traditional autoregressive processes. By looking at the connection between risk and volatility, we discussed how financial time series are a particularly interesting target for the analysis of conditional-level fluctuations in variance. From there, we described how the ARCH framework introduced by Engle (1982) fits into the literature and why his breakthroughs found such successful applications in econometric models for the past forty years.

We start by breaking down our series into a mean and residual innovation process. Though stationary, the series of monthly log returns of the NASDAQ itself does not exhibit any significant autocorrelation, so our mean process is found to be redundant; with no meaningful autocorrelation effects in the data, we consider the log-return series as our base residual series. While we expect this series to behave like white noise (and indeed, it does appear to have no significant serial correlation) there is a clear dependence at the conditional level of the data. There is a high serial correlation in the squared log-returns series, suggesting there is an underlying dependence and fluctuation in conditional variance. By identifying this conditional heteroskedasticity, clusters of volatility, and leptokurticity, we find that our log-returns series exhibits ARCH effects. Having identified these ARCH effects, we discuss the selection of both ARCH and Generalized ARCH (GARCH) models to study this behavior. We focus on low-order ARCH and GARCH models and go on to identify a GARCH(1,1) model to be the most performant. We additionally find that fitting to the residuals of this model under the Student's t-distribution is more appropriate than the Normal distribution. This is seen in the performance of all 6 fitted models we study. Using our results, we find that our final ARMA(0,0)-GARCH(1,1) model under the Student's t-distribution fits well to the historical log-returns data. Residual analysis indicates a strong fit while being sufficiently parsimonious and forecasting well.

References

Andersen, T., Bollerslev, T., Christoffersen, P., & Diebold, F. (2005). *Volatility forecasting* (tech. rep.). National Bureau of Economic Research. https://doi.org/10.3386/w11188

Bauwens, L., Laurent, S., & Rombouts, J. (2003). Multivariate GARCH models: A survey. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.411062

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654. https://doi.org/10.1086/260062

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Chatfield, C., & Xing, H. (2019). *The analysis of time series*. Chapman; Hall/CRC. https://doi.org/10.1201/9781351259446

Engle, R. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, *50*(4), 987. https://doi.org/10.2307/1912773

Engle, R. (2004). Risk and volatility: Econometric models and financial practice. *American Economic Review*, *94*(3), 405–420. https://doi.org/10.1257/0002828041464597

Hansen, P. R., & Lunde, A. (2001). A forecast comparison of volatility models: Does anything beat a GARCH(1, 1)? *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.264571

Hughes, A. W., King, M. L., & Kwek, K. T. (2004). Selecting the order of an ARCH model. *Economics Letters*, *83*(2), 269–275. https://doi.org/10.1016/j.econlet.2003.05.003

Jafari, G. R., Bahraminasab, A., & Norouzzadeh, P. (2007). Why Does The Standard GARCH(1,1) Model Work Well? *International Journal of Modern Physics C*, *18*(07), 1223–1230. https://doi.org/10.1142/s0129183107011261

Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business*, *36*(4), 394. https://doi.org/10.1086/294632

Mandelbrot, B. B. (1997). The variation of certain speculative prices. In *Fractals and scaling in finance* (pp. 371–418). Springer New York. https://doi.org/10.1007/978-1-4757-2763-0_14

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, *19*(3), 425. https://doi.org/10.2307/2977928

Tsay, R. S. (2010). *Analysis of financial time series*. John Wiley & Sons, Inc. https://doi.org/10.1002/9780470644560

Wang, Y., Xiang, Y., Lei, X., & Zhou, Y. (2021). Volatility analysis based on GARCH-type models: Evidence from the chinese stock market. *Economic Research-Ekonomska Istraživanja*, *35*(1), 2530–2554. https://doi.org/10.1080/1331677x.2021.1967771

Xiao, L., & Aydemir, A. (2007). Volatility modelling and forecasting in finance. In *Forecasting volatility in the financial markets* (pp. 1–45). Elsevier. https://doi.org/10.1016/b978-075066942-9.50003-0

Appendices

A Further Model Diagnostics

Table 5: Summary of ARMA(0,0)-GARCH(1,1) model performance and standardized residuals tests under a Student's t-distribution. With no serial correlation in the standard residuals (and matching the distribution assumption), we find that the model is a good fit of the data. This is further supported by the principle of parsimony, as we achieve these results with few parameters and low ARCH order.

Performance		
Log Likelihood		680.5589
AIC		-2.950039
BIC		-2.904986
SIC		-2.950274
HQIC		-2.932295
Standardized Residuals Tests		
Jarque-Bera Test	863.7822	p-value = 0
Shapio-Wilk Test	0.9344609	p-value = 2.657985e-13
Ljung-Box Test (R, Q(10))	7.304218	p-value = 0.6964466
Ljung-Box Test (R, Q(15))	9.433403	p-value = 0.8537886
Ljung-Box Test (R, Q(20))	13.28696	p-value = 0.8647351
Ljung-Box Test (R^2, Q(10))	1.527112	p-value = 0.9988484
Ljung-Box Test (R^2, Q(15))	6.581835	p-value = 0.9682328
Ljung-Box Test (R^2, Q(20))	7.04916	p-value = 0.9965207
LM Arch Test (R, TR^2)	2.372653	p-value = 0.9985833

B R Code

⋆ Additional plotting functions and helper methods can be found in the following repository: https://github.com/sim15/ma464X_ts_report_code

```r
library(zoo)
library(dse)
library(LSTS)
library(fGarch)
library(extrafont)
library(ggplot2)
library(tseries)
font_install("fontcm")
loadfonts()

setPdf <- function(fname, w, h) {
  pdf(file=fname, height=h, width=w, family="CM Roman")
  par(
    family="CM Roman"
  )

  # pdf(file=fname, height=h, width=w)
}

# plotting functions
source("plotter.R")

# load raw data
nasdaq <- read.csv("^IXIC (1).csv", header=T)
ts.nasdaq <- ts(nasdaq$Adj.Close)
nasdaq.logreturns <- diff(log(ts.nasdaq))

# plot of monthly prices and returns (two graphs)
{
setPdf("./nasdaqFigs/monthlyPricesLogReturns.pdf", 8, 3)
dailyPricesReturns(ts.nasdaq, nasdaq.logreturns)
dev.off()
}

# plot acf and pacf corelogram
{
  setPdf("./nasdaqFigs/acfpacf.pdf", 1.2*9, 1.2*3)
  acfPacfPlot(nasdaq.logreturns, maxLag=40)
  dev.off()
}

# plot summary statistics + Q-Q plot
{
  setPdf("./nasdaqFigs/summaryStats.pdf", 1.7*8, 1.7*3)
  summaryPlot(nasdaq.logreturns)
  dev.off()
}

# plot combined returns and price of data
```

```r
{
  setPdf("./nasdaqFigs/combinedReturns.pdf", 1.2*9, 1.2*5)
  oldCombined(ts.nasdaq, nasdaq.logreturns,
              ylim1=c(0,23000),
              ylim2=c(-1,0.2),
              ylab1=seq(0, 15000, by = 3000),
              ylab2=seq(-0.3, 0.3, by = 0.15),
              legendLabel=c("NASDAQ", "NASDAQ Returns"),
              mainn="Monthly Closing Price and Returns of NASDAQ (19852023)",
              roundDigits=3)
  dev.off()
}

# ADF test results for stationarity
adf.test(nasdaq.logreturns)

# all model fits (normal and Student's t-distribution)
fit.arch1.norm=garchFit(
  nasdaq.logreturns~garch(1,0),
  data=nasdaq.logreturns,trace=F)
fit.arch1.std=garchFit(nasdaq.logreturns~garch(1,0),
                       data=nasdaq.logreturns,trace=F,
                       cond.dist="std")

fit.arch2.norm=garchFit(nasdaq.logreturns~garch(2,0),
                        data=nasdaq.logreturns,
                        trace=F)
fit.arch2.std=garchFit(nasdaq.logreturns~garch(2,0),
                       data=nasdaq.logreturns,trace=F,
                       cond.dist="std")

fit.arch3.norm=garchFit(nasdaq.logreturns~garch(3,0),
                        data=nasdaq.logreturns,trace=F)
fit.arch3.std=garchFit(nasdaq.logreturns~garch(3,0),
                       data=nasdaq.logreturns,trace=F,
                       cond.dist="std")

fit.garch11.norm=garchFit(nasdaq.logreturns~garch(1,1),
                          data=nasdaq.logreturns,trace=F)
fit.garch11.std=garchFit(nasdaq.logreturns~garch(1,1),
                         data=nasdaq.logreturns,trace=F,
                         cond.dist="std")

fit.garch12.norm=garchFit(nasdaq.logreturns~garch(1,2),
                          data=nasdaq.logreturns,trace=F)
fit.garch12.std=garchFit(nasdaq.logreturns~garch(1,2),
                         data=nasdaq.logreturns,trace=F,
                         cond.dist="std")

fit.garch21.norm=garchFit(nasdaq.logreturns~garch(2,1),
                          data=nasdaq.logreturns,trace=F)
fit.garch21.std=garchFit(nasdaq.logreturns~garch(2,1),
                         data=nasdaq.logreturns,trace=F,
```

```
                         cond.dist="std")


# plot residual summary of model
{
  setPdf("./nasdaqFigs/standardResidualSummary.pdf", 0.7*14, 0.7*5)
  residualPlots(residuals(fit.garch11.std, standardize=TRUE))
  dev.off()
}

# fit model and predictions results
summary(fit.garch11.std)

predict(fit.garch11.norm, n.ahead=6, plot = T)
predict(fit.garch11.std, n.ahead=6, plot = T)

plot(fit.garch11.std)
fit.garch11.std
```