Monte-Carlo Simulations of GARCH, GJR-GARCH and constant volatility on NASDAQ-500 and the 10 year treasury

Lee, Dongkeun
Liu, David
Duke University

April 21, 2014

Contents

0.1 Abstract 1
0.1.1 Acknowledgements 2
0.1.2 Paper Organization 2
0.2 Motivation for GARCH 2
0.2.1 Testing Serial Correlation 3
0.3 GARCH Model 3
0.4 Estimation of GARCH parameters 4
0.5 GJR-GARCH 6
0.6 Analysis 6
0.6.1 Parameter Estimates 6
0.6.2 t-test statistic 7
0.6.3 News Impact Curves 8
0.6.4 Monte Carlo Simulation Results 8
0.6.5 Goodness-of-fit metrics 10
0.6.6 Analysis of the Simulation Results 11
0.7 Conclusion 13

0.1 Abstract

Our paper asks and attempts to answer the following two questions: Does an ARMA(1,1) Model with a constant Gaussian volatility parameter estimate better predict financial data than an ARMA(1,1) Model with a GARCH volatility parameter estimate? Does the GJR-GARCH volatility parameter estimate improve on the GARCH parameter estimate?
We show that the constant volatility parameter estimate does a better job with the NASDAQ-500 index but worse with the 10-year treasury. Under hypothesis testing, we show that the GJR-GARCH(1,1) is a better model to use for the NASDAQ while GARCH(1,1) should be used to model the treasury. Finally, using various goodness-of fit metrics, we determine that GARCH(1,1) fits better than the GJR-GARCH(1,1) model on the NASDAQ-500 and the 10-year treasury. [1] [2]

[1]Dongkeun Lee Duke University Department of Economics dl163@duke.edu
[2]David Liu Duke University Department of Economics dl110@duke.edu

0.1.1 Acknowledgements

This paper would not have been possible without the invaluable help and advice of colleagues and professors. We want to thank Dr. Arlie Petters for encouraging us to study GARCH in his financial derivatives course. We also want to thank Dr. Edward Tower, Dr. Charles Becker and Dr. Andrew Patton for their helpful advice and comments. Finally, we want to thank all our friends at Duke as well as our families for their friendship and unwavering support.

0.1.2 Paper Organization

The paper is organized as follows. In the first section (Section 0.2-0.5), we motivate GARCH, introduce the Ljung-Box Q test and the GARCH Model (Bollerslev (1986)), derive the log-likelihood form of GARCH(1,1) and finally, motivate and introduce the GJR-GARCH(1,1) model (Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993)) model.
In the second section (Section 0.6-0.7), we present our analysis. We show that under 20 different Monte-Carlo simulation trials, the GARCH Model misses an average of 7.95 and 6.4 times, the GJR-GARCH Model misses an average of 10.95 and 35.6 times while the constant volatility misses an average of 4.6 and 27.95 times for the NASDAQ and the treasury, respectively. We show that GARCH is preferable to GJR-GARCH using different goodness-of-fit tests, but not necessarily in hypothesis testing. Finally, we conclude by discussing which model we would use to forecast financial data in the future.

0.2 Motivation for GARCH

Let WN denote a White Noise process. Consider a standard ARMA(1,1) model for an asset return:

$$\begin{aligned} &Y_{t+1} = \phi_0 + \phi_1 Y_t + \epsilon_{t+1} + \theta\epsilon_t \\ &\epsilon_{t+1} \sim WN(0, \sigma^2) \end{aligned}$$

which implies that

$$\begin{aligned}V_t[Y_{t+1}] &= V_t[\phi_0 + \phi_1 Y_t + \epsilon_{t+1} + \theta\epsilon_t] \\ &= V_t[\epsilon_{t+1}], \text{ since the other terms are known at time t} \\ &= \sigma^2 \text{ by assumption that } \epsilon_{t+1} \sim WN(0, \sigma^2)\end{aligned}$$

Thus, in standard models $V_t[Y_{t+1}] = \sigma^2$, a constant, which means that Y_{t+1} is (conditionally) homoscedastic.

0.2.1 Testing Serial Correlation

McLeod and Li (1983) suggest using the Ljung-Box test on the squared residuals to test jointly for evidence of serial correlation. If the squared residuals of the data is serially correlated, Y_{t+1} cannot be (conditionally) homoscedastic. The test is as follows. Assume that the the stock prices are distributed IID logNormal (so returns are Normally distributed). Let ρ_j denote the jth autocorrelation i.e.

$$\rho_j := Corr(Y_t, Y_{t-j}) = \frac{Cov(Y_t, Y_{t-j})}{\sqrt{Y_t}\sqrt{Y_{t-j}}}$$

Then the Ljung-Box Q-statistic tests the following.

$$\begin{aligned}H_0 &: \rho_1 = \rho_2 = ... = \rho_L = 0 \\ H_1 &: \rho_j \neq 0 \text{ for some j=1,2...L}\end{aligned}$$

The Ljung-Box Q statistic, denoted $Q_{LB}(L)$, is:

$$Q_{LB}(L) = T(T+2)\sum_{j=1}^{L}(\frac{1}{T-j})\hat{\rho_j^2}$$

Under the null hypothesis, the $Q_{BL}(L)$ statistics is distributed as χ^2_L with L degrees of freedom. The weakness of this test is that the test is not robust; we need to assume that stock prices are distributed iid logNormal in order to use the test.

0.3 GARCH Model

Let η_{t+1} be defined as:

$$\eta_{t+1} = \epsilon^2_{t+1} - E_t[\epsilon^2_{t+1}],\ \eta_{t+1} \sim WN(0)$$

GARCH(1,1) (Bollerslev (1986)) applies the ARMA(1,1) model for the conditional variance.

$$\epsilon_{t+1}^2 = \omega + \gamma\epsilon_t^2 + \gamma\eta_t + \eta_{t+1}\eta_{t+1} \;\sim WN(0)$$

Then, the conditional variance is

$$\begin{aligned}
\sigma_{t+1}^2 &:= V_t[Y_{t+1}] = E_t[\epsilon_{t+1}^2] \\
&= E_t[\omega + \gamma\epsilon_t^2 + \gamma\eta_t + \eta_{t+1}] \\
&= \omega + \gamma\epsilon_t^2 + \gamma\eta_t \\
&= \omega + \gamma\epsilon_t^2 + \gamma(\epsilon^2 - E_{t-1}[\epsilon_t^2]) \\
&= \omega + \gamma\epsilon_t^2 + \gamma(\epsilon^2 - \sigma^2) \\
&= \omega + (\gamma + \lambda)\epsilon_t^2 - \gamma\sigma_t^2 \\
&= \omega + \alpha\epsilon_t^2 + \beta\sigma_t^2
\end{aligned}$$

where $\alpha = (\gamma + \lambda)$ and $\beta = -\lambda$. The equation states that conditional variance of tomorrow's return is equal to a constant, plus today's residual squared, plus today's known variance.

0.4 Estimation of GARCH parameters

GARCH models are estimated using MLE. MLE requires us to make certain assumptions about the distribution of ϵ_{t+1}. The most common distributional assumption is that of normality:

$$\epsilon_{t+1}|\mathcal{F}_t \sim \mathcal{N}(0, \sigma_{t+1}^2)$$

We assume that the conditional mean follows an ARMA(1,1) process. Then, this implies that the time series is conditionally normally distributed.

$$\begin{aligned}
Y_{t+1} &= \mu_{t+1} + \epsilon_{t+1} \\
\mu_{t+1} &= E_t[Y_{t+1}] = \phi_0 + \phi_1 Y_t + \lambda\epsilon_t \\
\sigma_{t+1}^2 &= V_t[Y_{t+1}] = \omega + \alpha\epsilon_t^2 + \beta\sigma_t^2
\end{aligned}$$

Thus, under properties of the Normal Distribution:

$$Y_{t+1}|\mathcal{F}_t \sim \mathcal{N}(\phi_0 + \phi_1 Y_t + \lambda\epsilon_t, \omega + \alpha\epsilon_t^2)$$

The likelihood function is then

$$\begin{aligned}\mathcal{L}(\theta|y_1, y_2...y_T) &= f(y_1...y_T)\\ &= f_{y_1}(y_1) \times \prod_{t=2}^{T} f_{y_t|y_{t-1}...y_1}(y_t|y_{t-1}...y_1)\end{aligned}$$

The distribution of $y_2|y_1...y_T|y_{T-1}...y_1$ are all Normal densities with different means and variances. We do not know what the unconditional distribution of y_1 is. Thus, we will just instead maximize the conditional likelihood, which is

$$f(y_2..y_T|y_1;\theta) = \prod_{t=2}^{T} f_{y_t|y_{t-1}...y_1}(y_t|y_{t-1}...y_1) = \prod_{t=2}^{T} \frac{1}{\sqrt{2\pi\sigma_t^2}} exp\{-\frac{\epsilon_t^2}{2\sigma_t^2}\}$$

where

$$\begin{aligned}\epsilon_t &= Y_t - \phi_0 - \phi_1 Y_{t-1} - \lambda\epsilon_{t-1}\\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2\end{aligned}$$

The conditional log-likelihood is thus:

$$\begin{aligned}\frac{1}{T-1} logf(y_2...y_T|y_1;\theta) &= \frac{1}{T-1} log\mathcal{L}(\theta|y_1, y_2...y_T)\\ &= \frac{-1}{2} log(2\pi) - \frac{1}{2(T-1)}\sum_{t=2}^{T} log\sigma_t^2 - \frac{1}{2(T-1)}\sum_{t=2}^{T}\frac{\epsilon_t^2}{\sigma_t^2}\end{aligned}$$

The MLE is thus

$$\begin{aligned}\hat{\theta} &= \arg\max_{\theta} log\mathcal{L}(\theta|y_1...y_T)\\ &= \arg\max_{\theta}\{-\frac{1}{2}log(2\pi) - \frac{1}{2(T-1)}\sum_{t=2}^{T} log\sigma_t^2 - \frac{1}{2(T-1)}\sum_{t=2}^{T}\frac{\epsilon_t^2}{\sigma_t^2}\}\end{aligned}$$

There is no closed-form solution for the MLE estimates so a common approach is to use numerical methods to maximize the likelihood. For this paper, we will use EstMdl function in MATLAB to estimate the parameters.

0.5 GJR-GARCH

Black (1976) observed that stock returns are negatively correlated with respect to changes in volatility i.e. lower stock returns than expected followed with higher levels of volatility and higher stock returns than expected followed with lower levels of volatility. This is intuitively explained by what is called the "leverage effect". If a firm uses both debt and equity to finance, as stock prices decline, its debt to equity ratio will increase, which increases equity return volatility. Thus, lower stock returns than expected leads to higher future volatility and higher stock returns than expected leads to lower future volatility. The GJR-GARCH Model is a model within the class of models that captures this "leverage effect."
The standard GARCH Model, which shows that tomorrow's volatility is quadratic to today's residual, so the sign of the residual does not matter.

$$\sigma_{t+1}^2 = \omega + \beta\sigma_t^2 + \alpha\epsilon_t^2$$

Glosten, Jagannathan and Runkle (1993) (called GJR-GARCH) extends the GARCH Model to capture this effect.

$$\text{GJR-GARCH: } \sigma_{t+1}^2 = \omega + \beta\sigma_t^2 + \alpha\epsilon_t^2 + \delta\epsilon_t^2\mathbb{1}\{\epsilon_t < 0\}$$

If $\delta > 0$ then the impact of tomorrow's volatility is negative if today's residual is negative.

0.6 Analysis

We take the NASDAQ Composite (IXIC) and the CBOE Interest Rate 10-year T-No (TNX) from 2/7/2000 to 11/6/2013. The Ljung-Box Test (done in MATLAB) shows that we reject H_0 for NASDAQ but cannot reject H_0 for the 10 year treasury. Thus, for the treasury, there isn't evidence that the squared residuals are serially auto-correlated.

0.6.1 Parameter Estimates

For the NASDAQ Composite, the GARCH(1,1) Conditional Variance Model parameter estimates are (in MATLAB)

Parameter	Value	Standard Error	t statistic
Constant	1.5477e-06	5.8247e-07	2.65713
GARCH{1}	0.917646	0.0074822	122.644
ARCH{1}	0.0763399	0.00661047	11.5483

and under the GJR-GARCH(1,1) Conditional Variance Model

Parameter	Value	Standard Error	t statistic
Constant	2e-05	2.08912	9.57339
GARCH{1}	0.742542	0.0199526	37.2154
ARCH{1}	0.0535206	0.0135482	3.67884
Leverage{1}	0.267255	0.0289146	9.24292

The ARMA(1,1) parameters are

Parameter	Value	Standard Error	t statistic
Constant	-1.4158e-05	0.000149655	-0.094601
AR{1}	0.473452	0.113503	4.17129
MA{1}	-.525954	0.108609	-4.84317
Variance	0.000303766	4.01074e-06	75.7381

For the 10 year Treasury, the GARCH(1,1) Conditional Variance Model parameter estimates are (in MATLAB)

Parameter	Value	Standard Error	t statistic
Constant	7.23721e-07	4.70219e-07	1.53911
GARCH{1}	0.947127	0.00437276	216.597
ARCH{1}	0.0523367	0.00468205	11.1782

and under the GJR-GARCH(1,1) Conditional Variance Model

Parameter	Value	Standard Error	t statistic
Constant	2e-05	2.04482e-06	9.78082
GARCH{1}	0.800138	0.0131241	60.9673
ARCH{1}	0.131747	0.017005	7.72757
Leverage{1}	0.030674	0.0200255	1.53175

The ARMA(1,1) parameters are

Parameter	Value	Standard Error	t statistic
Constant	-0.000433322	0.0005393	-0.803296
AR{1}	-0.623894	0.436365	-1.42975
MA{1}	0.636994	0.430507	1.47958
Variance	0.000360598	4.98113e-06	72.3928

0.6.2 t-test statistic

In the following two paragraphs, we will use the t-test statistic to compare GARCH(1,1) with GJR-GARCH(1,1). Later on in our analysis, we will use goodness-of-fit metrics to compare GARCH(1,1) with GJR-GARCH(1,1).

For NASDAQ, the t-test statistic for the leverage is 9.24292, which illustrates there is indeed a significant 'leverage effect.' This implies that the GJR-GARCH(1,1) is a better model to use than the GARCH(1,1).
However, for the treasury, the t test statistic for the leverage is 1.53175, which is lower than the 95% (with 4 degrees of freedom) significant value of 2.132. Thus, since the treasury bill does not exhibit this 'leverage effect,' the t-test indicates that GARCH(1,1) is the appropriate model to use for the treasury.

0.6.3 News Impact Curves

The following 'News Impact Curves' (Figure 1) graphs the impact of today's residual squared term on tomorrow's variance against different values of ϵ_t. In the GARCH-Model the impact of today's residual squared term on tomorrow's variance is invariant to the residual; thus the equations are governed by

$$y_t = 0.0763399\epsilon_t^2$$

while in the GJR-GARCH Model, the equations are governed by

$$y_t = \begin{cases} 0.320776\epsilon_t^2, & \text{if } \epsilon_t < 0 \\ 0.053521\epsilon_t^2 & \text{if } \epsilon_t \geq 0 \end{cases}$$

In accordance to the motivation of GJR-GARCH, volatility is higher following a better than expected return to the stock and lower following a lower than expected return.

0.6.4 Monte Carlo Simulation Results

The figures below (Figure 2 and 3) show one particular Monte-Carlo simulation of returns on the NASDAQ and 10-Year US Treasury indexes. The prediction model assumes that

$$Y_{t+1} = \mu_{t+1} + \epsilon_{t+1}$$
$$\text{where } \epsilon_{t+1}|\mathcal{F}_t \sim \mathcal{N}(0, \sigma_{t+1}^2)$$

with μ_{t+1} driven by an ARMA(1,1) model and σ_{t+1}^2 driven by a GARCH(1,1) model. Because GARCH gives us a model for the conditional variance, we

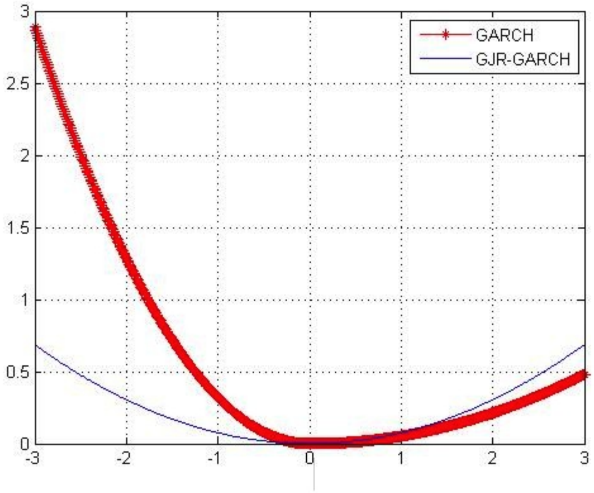

Figure 1: News Impact Curves for NASDAQ-500

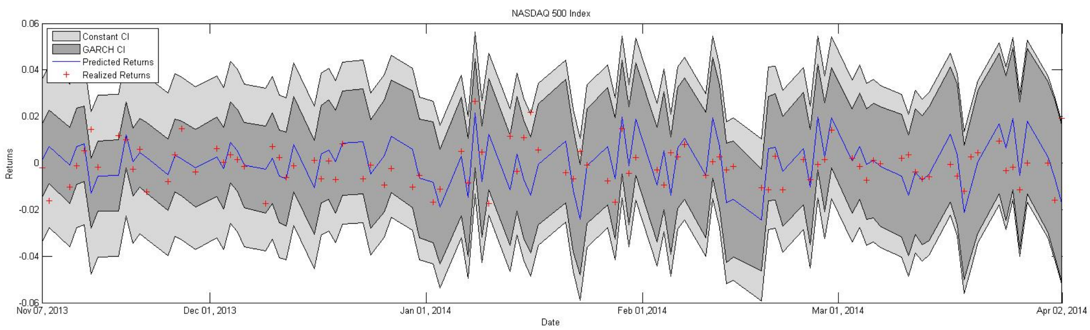

Figure 2: NASDAQ 500 Simluation

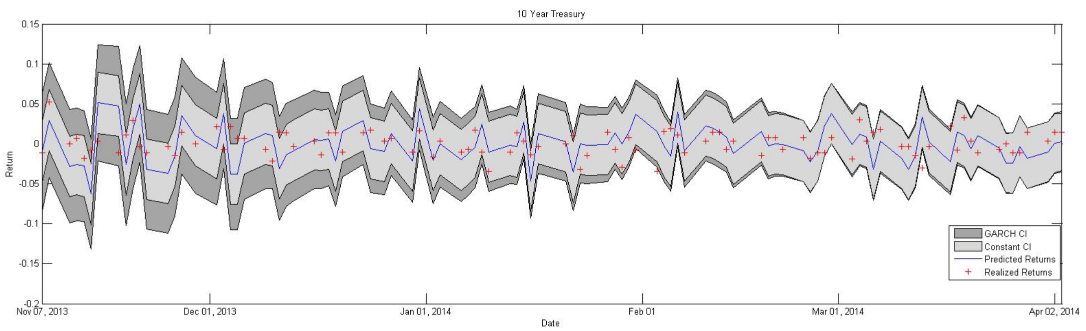

Figure 3: Treasury Simulation

can construct a two standard deviation confidence interval around our predicted value $\hat{\mu}_{t+1}$ whose width varies with time:

$$[\hat{\mu}_{t+1} - 2\hat{\sigma}_{t+1}, \hat{\mu}_{t+1} + 2\hat{\sigma}_{t+1}]$$

In contrast, a model that assumes constant conditional variances has a confidence interval of constant width around the predicted return.

The simulations in Figure 2 and Figure 3 use historical data from 2/7/2000 to 11/6/2013 to estimate the model parameters, then based on those estimates we forecast returns on the indexes for the next 100 days and compare the results to the realized returns from 11/7/2013 to 4/2/2014. In both figures we are able to observe how the width of the GARCH confidence interval changes through time. Thus, we see that by capturing the time-varying volatility of the assets, we are able to better estimate the uncertainty of our forecasts than with a constant volatility model.

Looking at Figure 2, we see that the 95% CI for the constant volatility is almost always larger than the 95% CI for GARCH. Thus, all the 'misses' with constant volatility were also 'misses' with GARCH plus GARCH 'misses' all the realized returns outside the 95% CI of GARCH but within the 95% CI of constant volatility. This begs the question: Is the CI for constant volatility always larger than the CI for GARCH?

The answer is no, as Figure 3 demonstrates. In Figure 3, the GARCH 95% CI is almost always larger than the constant volatility 95% CI. Thus, in the treasury, GARCH 'misses' a lot less than constant volatility.

Under the normality assumption of $\epsilon_{t+1}|\mathcal{F}_t$, we would expect to see about 5 observations outside of our two standard deviation confidence interval over the course of 100 observations. Using GARCH volatility, we observed 7 and 4 points outside of our interval, while the constant volatility model had 1 and 14 observations outside the interval. In 20 different simulation trials, we observed that the GARCH interval missed an average of 7.95 and 6.4 times for the NASDAQ and Treasury respectively, while the constant volatility model missed an average of 4.6 and 27.95 times. The GJR-GARCH on the other hand, missed 10.95 times and 35.6 times on the NASDAQ and Treasury, respectively. Thus, GARCH performed the best in the Monte-Carlo Simulation Results for the treasury while the constant volatility performed the best for NASDAQ-500. Figure 4 is the complete summary statistic for the simulation results misses that lied outside the two standard deviation confidence interval.

0.6.5 Goodness-of-fit metrics

We compared the goodness-of-fit of GARCH and GJR-GARCH on the NASDAQ and the treasury using four different metrics (ln L, AIC, HQIC, BIC). Our results are displayed in Figure 5. * denotes which model the specific test says is a better fit. The tests are ordered (left to right) by the severity the test penalizes extra parameters. According to our tests, the GARCH is a better fit than GJR-GARCH in all metrics. Since the AIC, HQIC and BIC are most commonly used, we conclude that the GARCH is preferred for both the NASDAQ and the treasury.

0.6.6 Analysis of the Simulation Results

Constant volatility performed best in out-of-sample simulations on NASDAQ while GARCH performed best in out-of-sample simulations on the treasury. This is a surprising result because the treasury is thought to be less volatile than NASDAQ whilst GARCH is considered to be a more accurate measure of volatility than constant volatility. One explanation on why this may not be as surprising as first thought could be because the time frame associated with the estimated parameters included the 2008 financial crises, and the first 100 sample days of NASDAQ was nothing like the returns from the financial crises. Thus, it may be the case that the crises 'skewed' the GARCH parameters enough that constant volatility became a better measure for dates in which the NASDAQ returns were pretty normal, as in the past 100 days. However, this 'skewing' effect would be smaller in the treasury because the treasury wasn't as affected by the financial crises as NASDAQ was. Thus, this could be the reason why GARCH performs better than constant volatility on the treasury but does not perform as well on NASDAQ.

Additionally, the t-test indicated that GJR-GARCH was a better model to use for NASDAQ but GARCH was a better model for the treasury. The goodness-of-fit metrics pointed to GARCH as the model to use. Additionally, GARCH performed best in out-of-sample simulations on the treasury while the constant volatility did best on NASDAQ. How does one elucidate these seeming contradictions? The simulation results are an out-of-sample analysis of the models while the goodness-of-fit tests analyzes how the models fit for the on-the-sample data. Thus, they measure two completely different things. Secondly, the simulations were performed during one specific interval. The results are unlikely to hold true if we were given another time interval. For example, it would not surprise us if GJR-GARCH convincingly beat GARCH and constant volatility in a Monte Carlo Simulation performed on the next hundred days. Lastly, the interval length may be too small for the models to work. GARCH(1,1) and GJR-GARCH(1,1) are often used to analyze between a year to two years worth of out-of-sample data. Thus, we could have suffered from a lack of sample size of sorts.

	Mean	SD	25-tile	50-tile	75-tile
NASDAQ					
Constant	4.6	3.330876	2.5	4	6
GARCH	7.95	2.781045	6	8	10
GJR-GARCH	10.95	6.747124	6.5	9.5	15
Treasury					
Constant	29.45	10.99509	21	31	35
GARCH	6.4	1.902906	6	6	8
GJR-GARCH	35.6	9.582879	29.5	37	41.5

Figure 4: Summary Statistic for 'misses'

	ln L	AIC	HQIC	BIC
NASDAQ 500				
GARCH	9.9332e03*	-5.7416*	-5.7397*	-5.7383*
GJR-GARCH	9.8683e03	-5.7035	-5.7010	-5.6964
Treasury				
GARCH	9.3862e03*	-5.4333*	-5.4314*	-5.4279*
GJR-GARCH	9.2663e03	-5.3632	-5.3607	-5.3561

Figure 5: Goodness of Fit Tests

0.7 Conclusion

In our paper, we took two financial time series models that modelled conditional variance, did a Monte-Carlo simulation using the estimated parameters and studied how the predicted returns matched up with the realized returns in a specified time interval. We wanted to study how the different conditional variance models compared to its predecessor, the constant conditional variance model. Finally, we performed goodness-of-fit tests as well as the t-test statistic to compare GARCH and GJR-GARCH.

In out-of-sample data, the GJR-GARCH and the constant variance produces horrendous results in simulation against the treasury yield. The GARCH does fairly well, but is not perfect. On the other hand, constant volatility clearly beats GARCH and GJR-GARCH on the NASDAQ. We suspected that the 2008 Financial Crises 'skewed' the GARCH parameters enough that it does worse on NASDAQ but not as bad in the treasury. However, according to goodness-of-fit tests, GARCH is the model to use (so consequently is the logical choice model to use to predict upcoming data). The t-test tells us that GARCH is better on the NASDAQ but GJR-GARCH is better on the treasury. And yet, with all these seeming contradictions, the result is nothing too surprising in the always confusing field of finance. GARCH may have been a refinement on the constant volatility and GJR-GARCH a further refinement on the GARCH, but it doesn't necessarily translate into a better predictive model. In fact, if we had to forecast financial data for the next hundred days, we would use the GJR-GARCH model, despite it being the worst of the three models in both NASDAQ and the treasury to use in the past hundred days! The reasons we would prefer GJR-GARCH over GARCH or constant volatility is that

1) The Ljung-Box test shows that asset returns are serially auto-correlated, which implies that constant volatility is the 'wrong' model to use.

2) The leverage effect is real and present in stocks and GARCH ignores this.

3) Although we haven't shown it, because of (1) and (2), we believe that as we increase the number of times we do out-of-sample simulations on different dates, GJR-GARCH will rise to be the better predictive model over GARCH or constant volatility.

This leads us to two points: finance is still very unpredictable, as an analogy, it is similar to how unpredictable the NCAA basketball tournament is nowadays even with all the advanced sports metrics in place and the 24/7 coverage of sports plastered online and on T.V. Secondly, if you develop a model that does a good job in forecasting financial data, you could make a lot of money if you kept it a secret (one such strategy would be to sell far out-of-money calls and puts if you knew that the returns skewed far more than predicted).

Future Plans of the Authors

Dongkeun Lee will graduate with a Masters of Arts in Economics in Spring of 2014. Afterwards, he will work as a consultant for CarMax in Richmond, VA. His proudest achievement at Duke was placing 1st in the 2013 Fuqua Poker Tournament after having learned the strength of each poker hand minutes before.

David Liu will graduate with a Bachelors of Arts in Economics and a Bachelor of Science in Mathematics in Spring of 2014. Afterwards, he will work at BlackRock in New York, NY. He almost agreed to eat 8 McDoubles minutes before the financial derivatives final so that he could earn $5.

Bibliography

[1] Black, Fischer (1976). The pricing of commodity contracts, Journal of Financial Economics, 3, 167-179.

[2] Bollerslev, T. ”Generalized Autoregressive Conditional Heteroskedasticity.” Journal of Econometrics. Vol. 31, 1986, pp. 307327.

[3] Engle, R. F. ”Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation.” Econometrica. Vol. 50, 1982, pp. 9871007.

[4] Glosten, L. R., R. Jagannathan, and D. E. Runkle. ”On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks.” The Journal of Finance. Vol. 48, No. 5, 1993, pp. 17791801.

[5] McLeod, A. I. and W. K. Li. ”Diagnostic Checking ARMA Time Series Models Using Squared-Residual Autocorrelations.” Journal of Time Series Analysis. Vol. 4, 1983, pp. 269273.

[6] Patton, Andrew. Forecasting Financial Markets Part 2. Econ 413: Forecasting Financial Markets. Duke University. Durham, NC. Notes.

[7] Wold, H. A Study in the Analysis of Stationary Time Series. Uppsala, Sweden: Almqvist & Wiksell, 1938.