A Time Series Analysis of the Volatility in Three Financial Sector Large Caps

A Senior Honors Thesis

Presented in Partial Fulfillment of the Requirements for Graduation *with Research Distinction* in the Undergraduate Colleges of The Ohio State University

by

David Garewal

The Ohio State University
Spring 2008

Project Advisor: Professor Peter F. Craigmile, Department of Statistics

I. Abstract

The financial sector has been at the center of attention for many investors in 2008 due to price volatility, extreme trading volumes, and innumerable problems within the sub-prime, banking, and brokerage industries. The collapse of Bear Stearns after missing earnings estimates for the fourth quarter of 2007 is one of the more notable catastrophes of 2008. This thesis will analyze three large cap financial stock companies in the investment banking and brokerage industry over the time period from January 2002 to the end of March 2008 to determine whether other companies will undergo similar future downtrends. These companies are Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS).

II. Introduction

The Goldman Sachs Group Inc. (GS), Merrill Lynch & Co., Inc. (MER), and Morgan Stanley (MS) are all current large cap financial sector leaders in the investment banking and brokerage sub-industry of the Standard & Poor's 500 Index Fund. The financial sector was of great interest to economists in the first quarter of 2008. In the first three months of the year, GS, MER, and MS all saw massive declines in their stock's value. We will investigate what relationship, if any, is there between the three stocks, whether the variability in the closing prices and volume has been constant, and if this decline was predictable.

The Goldman Sachs Group Inc. (GS) was founded in 1869 and is currently the world's largest investment bank[1]. GS became the largest investment bank partly because of the investment strategies of Goldman traders, such as Michael Swenson and Josh Birnbaum who acquired large financial profits by shorting the sub-prime mortgage-backed market.

This may be a reason behind Goldman avoiding writing off any sub-prime loans in 2007[2]. As of May 1st 2008, GS was trading at $199.05 per share and had a 52-week range of $140.27 - $250.70 per share[3]. The average volume for Goldman over the last 3 months has been 13,224,600 shares per day.

Merrill Lynch & Co., Inc. (MER) was founded in 1914 and headquartered in New York City[4]. By the end of 2007, MER wrote down $14.1 billion in sub-prime loans and recorded a deficit of $9.95 in earnings per share for the year[2]. As of May 1st 2008, MER was trading at $52.39 per share and had a 52-week range of $37.25 - $95.00 per share[3]. The average volume for Merrill over the last 3 months has been 28,002,600 shares per day.

Morgan Stanley (MS) was founded in 1935 and also has its global headquarters in New York City[5]. In the fourth quarter of 2007, MS wrote down $9.4 billion in assets and announced a deal it had made with the China Investment Corporation to receive $5 billion in capital infusion in exchange for securities convertible in 2010[2,5]. As of May 1st 2008, MS was trading at $50.33 per share and had a 52-week range of $33.56 - $90.95 per share[3]. The average volume for Morgan over the last 3 months has been 17,825,900 shares per day.

Currently, all of these three stocks are showing larger than the market betas (β=1.00), as well as an increase in average volume. A *beta value* is a measure of volatility, or systematic risk, of a security or a portfolio[7]. If the value is zero, no correlation to the stock market is shown. A positive beta shows a direct correlation, whereas a negative beta shows an inverse correlation to the stock market[6]. According to Yahoo! Finance's key statistics, the betas for GS, MER, and MS are 2.24, 1.82, and 1.75 respectively as of May 1st 2008[3].

In order to better analyze the weekly closing prices and volumes, we will construct statistical models such as an autoregressive conditional heteroskedascity model of order p (ARCH(p)) and assess their goodness of fit. In this model, the variance of current error terms is related to the squares of the past error terms[8]. Time series plots, plots of the autocorrelation function (ACF) and partial autocorrelation function (PACF), as well as Quantile-Quantile (Q-Q) plots will be used to assess goodness of fit under different model choices.

Due to the marked ongoing recent volatility in the price of stocks in the financial sector, we will take a closer look over a longer period between January 2002 and the end of March 2008. We will also search for clusters over the same time period for GS, MER, and MS with increased variability in weekly closing prices and increased trading volume.

Finally, we will evaluate the joint distribution functions between MS, MER, and GS to see whether or not any correlations exist between their price movements. In addition, we will check to see whether the price movement of any two of these stocks can be reasonably accurate in predicting the movement of the third stock.

After a thorough time series analysis of the weekly closing price and volume of Goldman Sachs, Merrill Lynch, and Morgan Stanley, we will try to answer as best we can whether the movement in the three stocks correlated at all and if so, what was the strength of the correlation? Was there any increase in variability of weekly closing stock prices and/or volume? Was the down trend in 2008 in the financial investment banking and brokerage sector foreseeable?

III. Exploratory Time Series Analysis (Univariate)

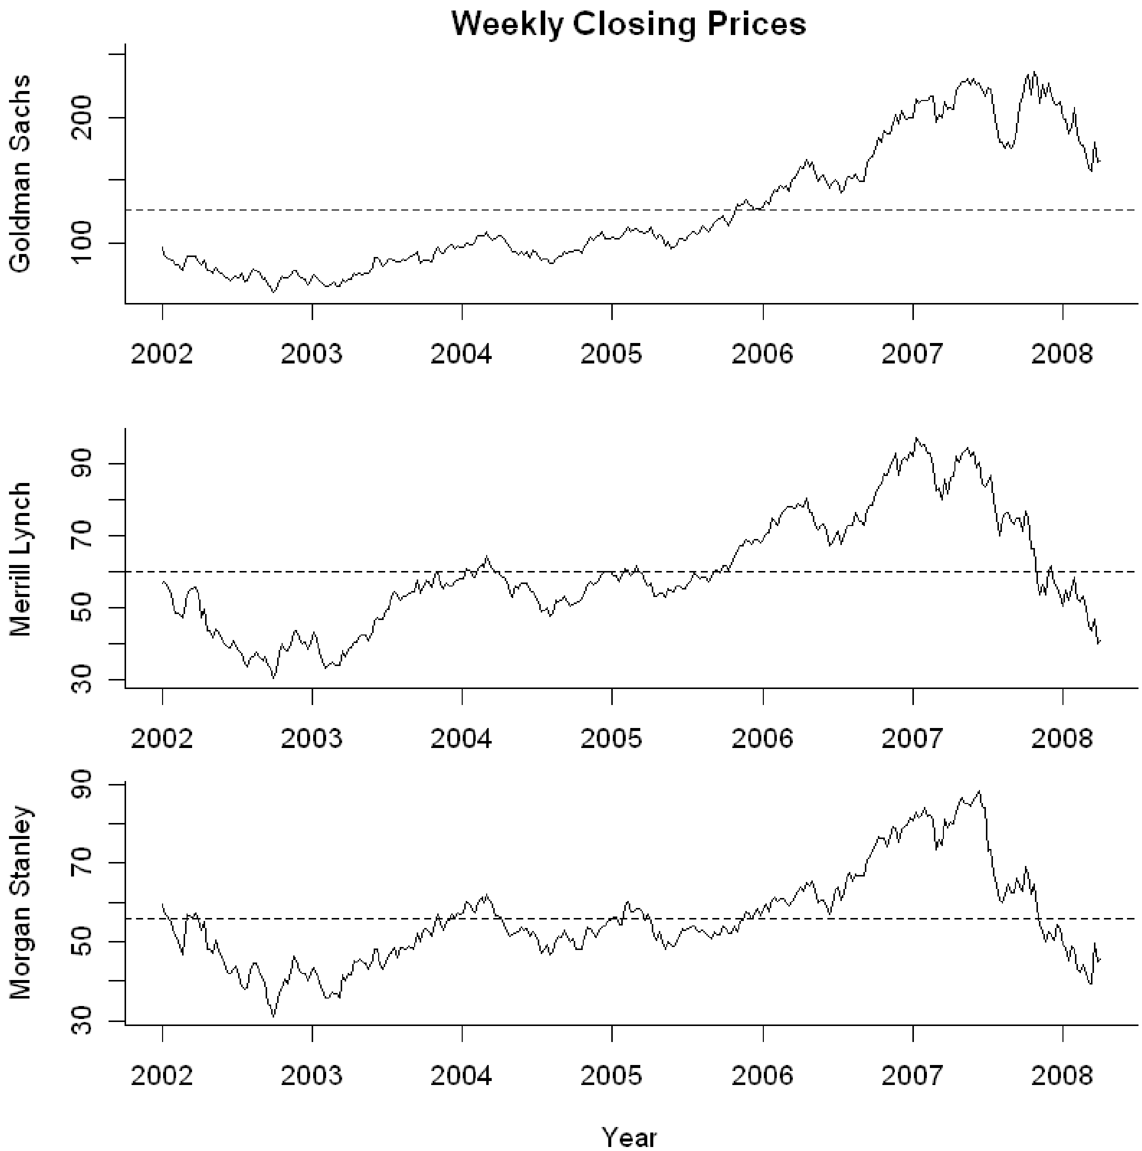

Figure 1: Weekly Closing Prices for Goldman Sachs, Merrill Lynch, and Morgan Stanley from January 2002 to March 2008 in the Standard & Poor's 500 Fund Index

Figure 1 shows the weekly closing prices of Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS) from the beginning of January 2002 to the end of March 2008. The graphs of these three financial companies appear to show some correlation. From 2002 to 2005, all three companies show a similar pattern in closing prices; all decreased in price from 2002 to 2003, increased from 2003 to 2004, and remain stagnant from 2004 to 2005. From 2005 to 2007, all showed an increase, but GS increased at a faster rate than MER and MS. During 2007 and 2008, all three stocks began decreasing in value, but at different rates. MER closed at a new weekly high in January 2007 ($97.02), but then fell after attaining this high. MS closes at its weekly high in May 2007 ($87.10), but also falls for the remainder of 2007 and start of 2008. GS performed differently in the same time period. In April 2007 GS reached a new weekly high ($230.71), but then showed a sharp decline of around 25% in July 2007 ($175.00). After this large decline, it erased all losses with a quick movement to a new high in September 2007 ($235.92). Following this rapid increase, GS began a gradual decrease in value which continued through March 2008.

Looking at the average daily volumes for Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS), GS averaged a little over 5 million shares traded per day, while MER and MS both averaged a little over 7 and 6 million shares per day, respectively. We will plot a base 10 log-transformation of the volumes against the year, because of evidence of a mean-variance relationship in the volumes (Figure 2). This will allow us better visibility of the volatility in all three financial companies over the time period. In June 2002, GS showed a great deviation away from mean volume while MER and MS remained fairly stable. Other noticeable spikes in volume away from the mean before 2007 are seen in Morgan Stanley between March 2005 and May 2005. In 2007 and 2008, all three stocks begin to show an increase in the average daily volume. GS and MS increased moderately from January 2007 to mid-2007 and then begin to level off into March 2008. ML increased much more than the other two stocks in the same time period and showed a continuous increase in volume from mid-2007 into March 2008.

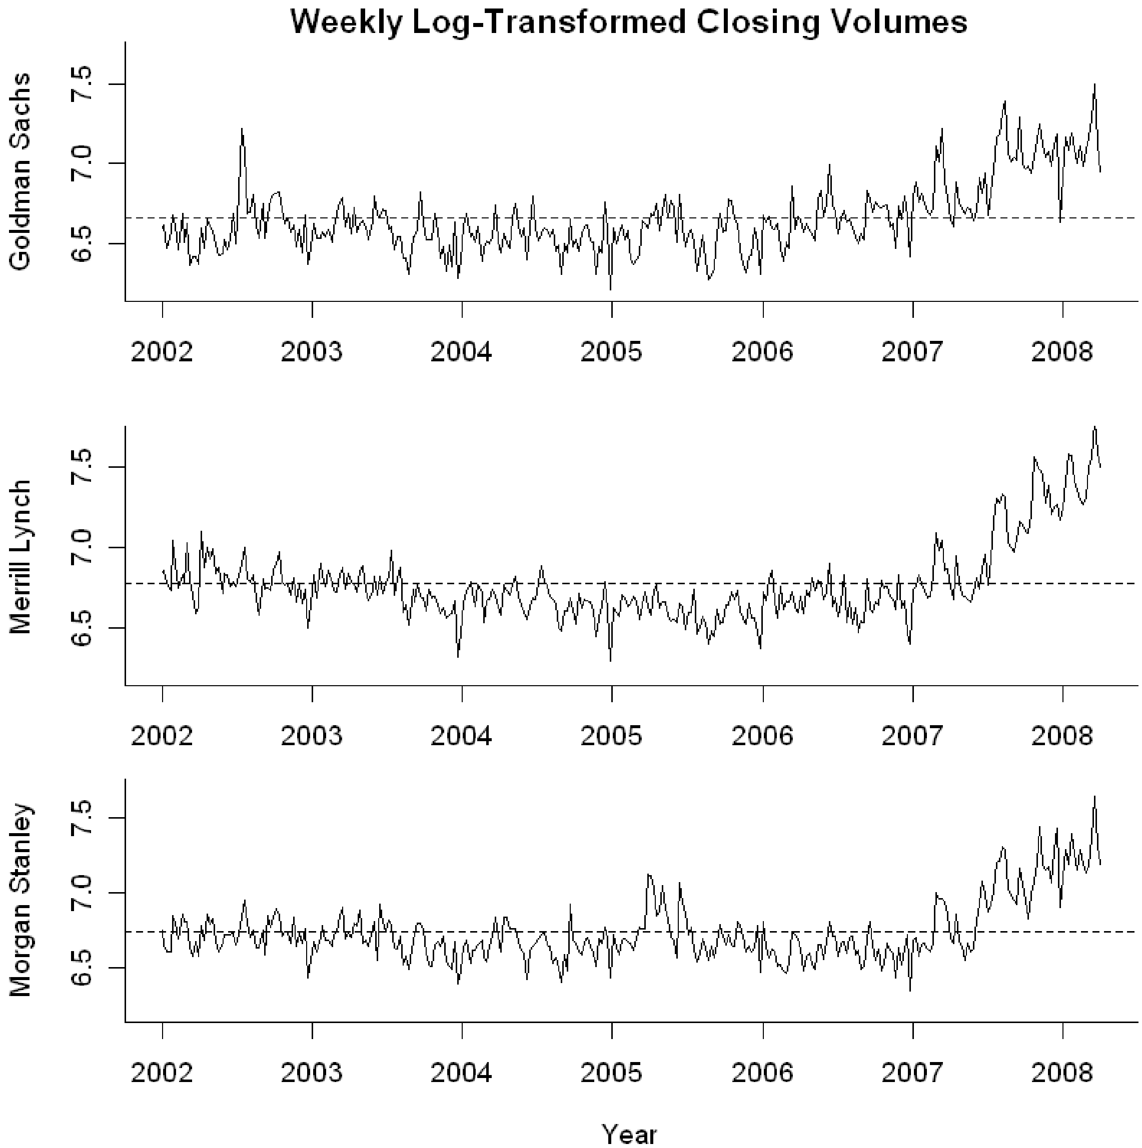

Figure 2: Average Daily Base 10 Log Volumes for Goldman Sachs, Merrill Lynch, and Morgan Stanley from the Weekly Closing Volume Data from January 2002 to March 2008 in the Standard & Poor's 500 Fund Index

We also perform a base 10 log-transformation of the closing prices data in order to reduce the mean-variance relationship for the three financials (Figure 3). From the beginning of 2003 until the end of 2006, there is evidence of an upward trend in the log closing prices for the three financial companies. From 2002 to 2003 and 2007 to the end of March 2008, the log closing prices of Goldman, Merrill, and Morgan tended to decrease, but at different rates. From 2007 to March 2008, Goldman Sachs decreased at a slower rate than the other two financials. To remove these trends and focus solely on the volatility, we will difference each log price series.

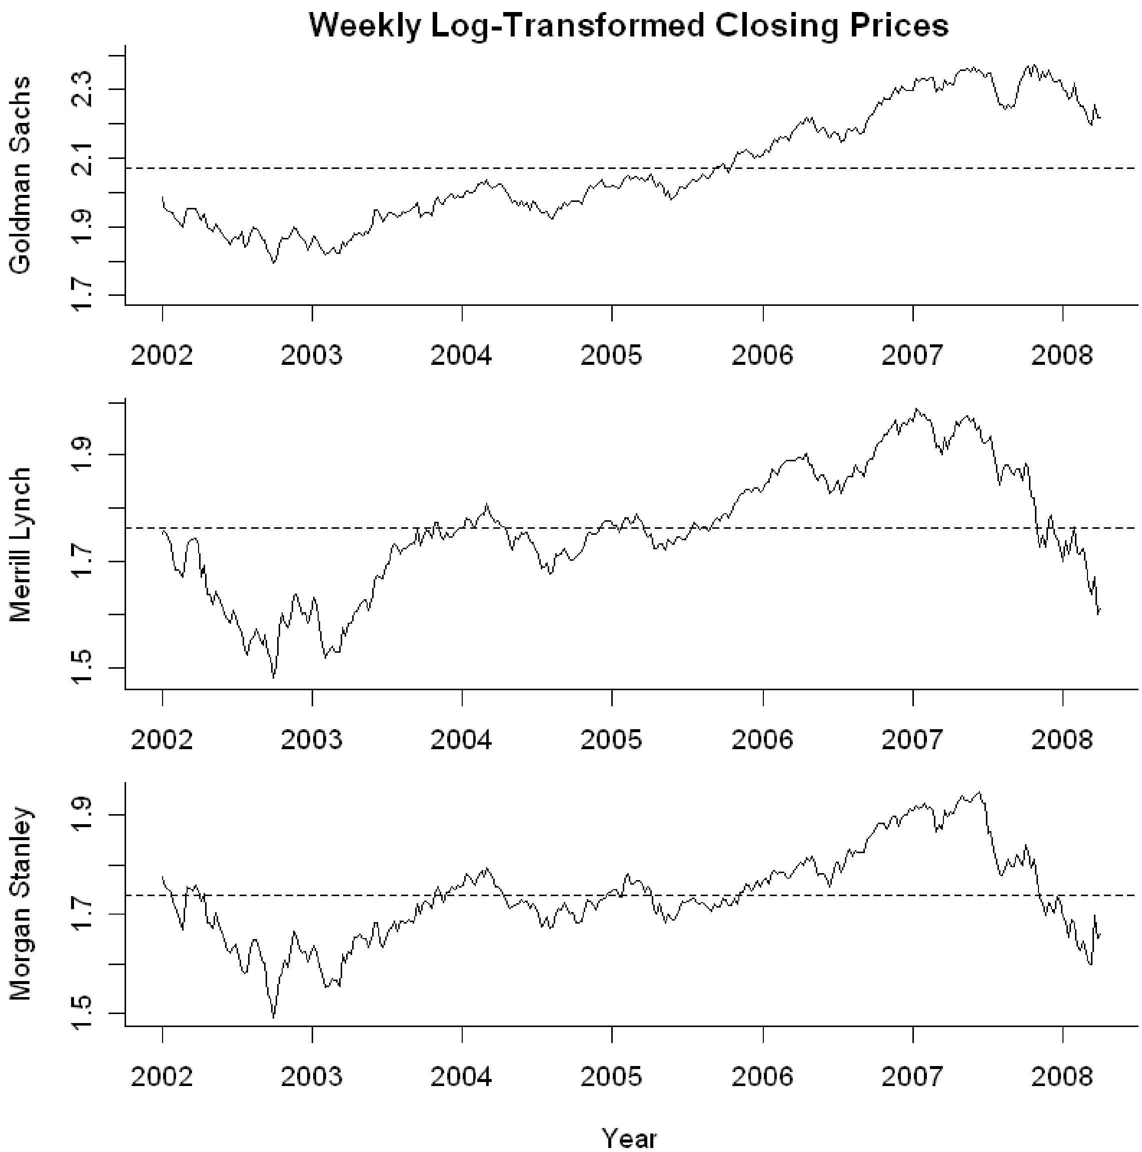

Figure 3: Base 10 Log Weekly Closing Prices of Goldman Sachs, Merrill Lynch, and Morgan Stanley from January 2002 to March 2008 under the Standard & Poor's 500 Fund Index

Figure 4 shows the weekly base 10 differenced log closing prices for Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS). Looking at this figure, we can now see the differenced log closing prices of GS, MER, and MS centered around a mean of zero with a variability that appears to change with time. MER and MS indicate possible increased volatility in their differenced log prices relative to GS. From January

2007 to March 2008, increased variability is seen in Goldman Sachs. Merrill Lynch shows more variability in 2002 and 2003, as well as 2007 and 2008. In 2007, Merrill Lynch's variability starts out quite small, but slowly increases as the year continues into the end of March. Morgan Stanley also has increased variability in 2002, 2003, 2007, and 2008. Notice how all three stocks have a great amount of variability towards the end of our dataset (March 2008).

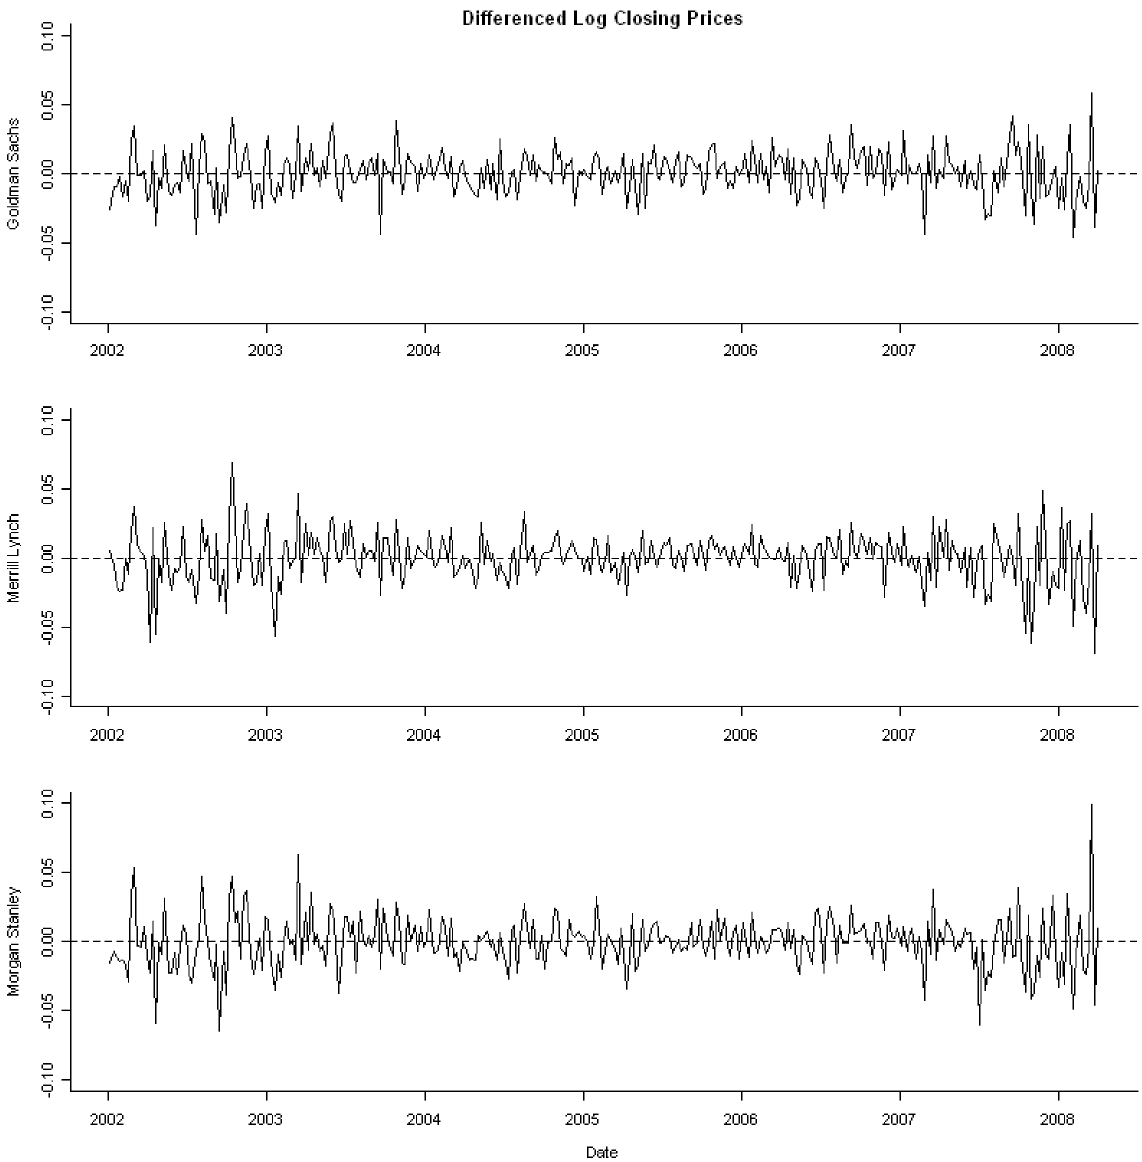

Figure 4: Weekly Base 10 Differenced Log Closing Prices for Goldman Sachs, Merrill Lynch, and Morgan Stanley from January 2002 to March 2008 under the Standard & Poor's 500 Fund Index

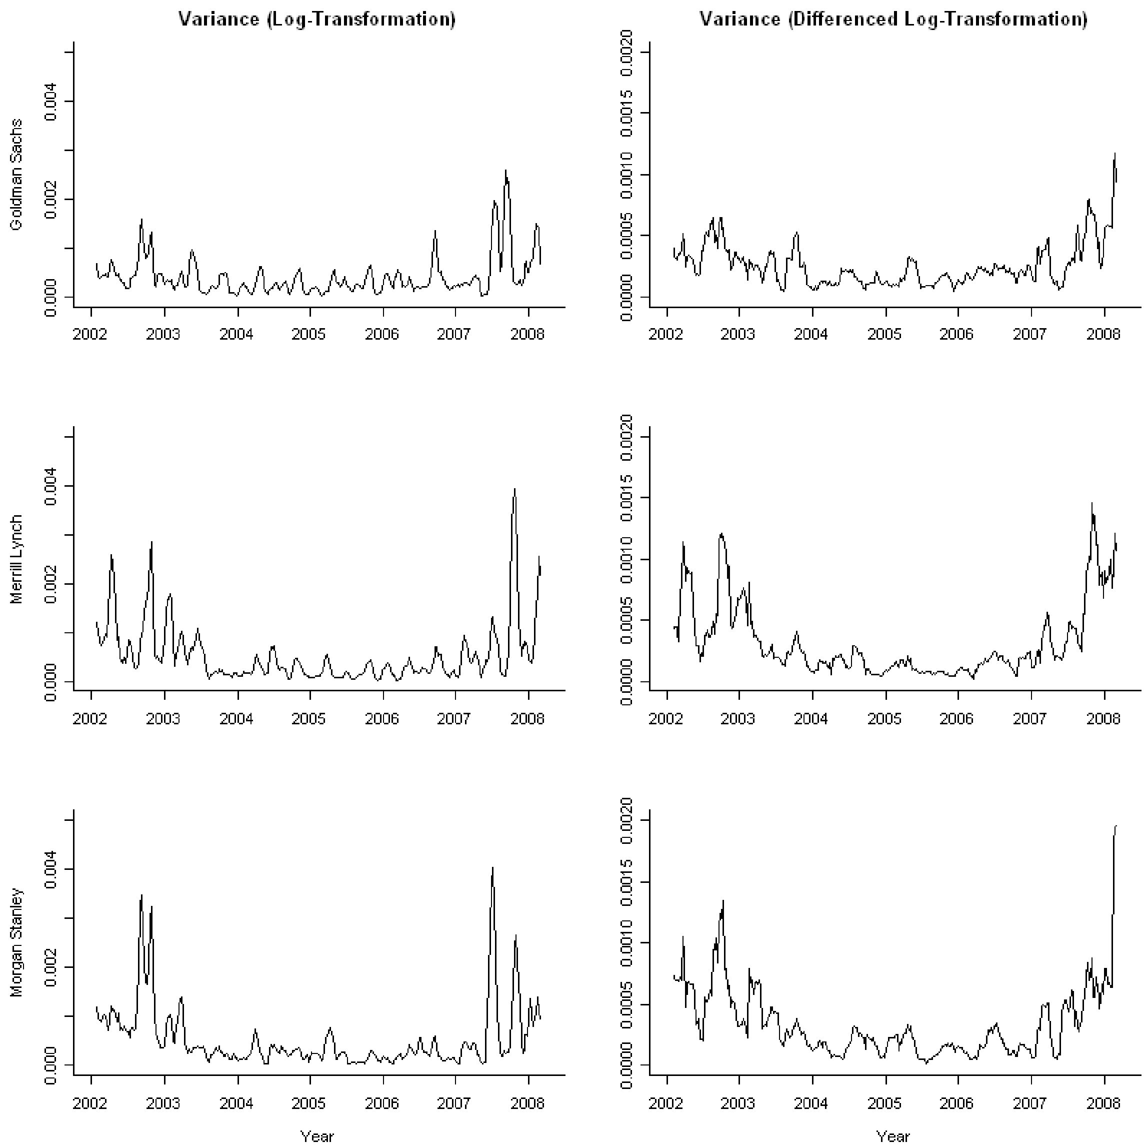

Figure 5: Goldman Sachs, Merrill Lynch, and Morgan Stanley's Variances from January 2002 to March 2008 Using Log-Transformation and Differenced Log-Transformation with Closing Price Data Recorded Every 4 Weeks

We will now plot the variance of the closing price data by time, over appropriate time windows for both the log and differenced log series of closing prices to indicate whether or not the closing price variability is constant or not. Once we have done so, we will try to smooth out the variance without losing any important signs of changes in the variance over time. A choice of interval length of 9 weeks gives a good estimate of the time-varying variance. This is shown for Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS) in Figure 5. Looking for significant changes in the variability of each series (GS, MER, and MS), we notice that all financials show a larger variance towards the end of 2002, as well as at the end of 2007 and into 2008. In addition, we see that the variance graphs of GS contain less variability than MER and MS. Moreover, all variance graphs appear fairly constant from January 2004 to December 2006. The volume graph (Figure 2) also shows an increase in variability in 2007 and 2008, which could be a reason for the increased variance in this period for all financials. Any statistical model should take these changes of variance into account.

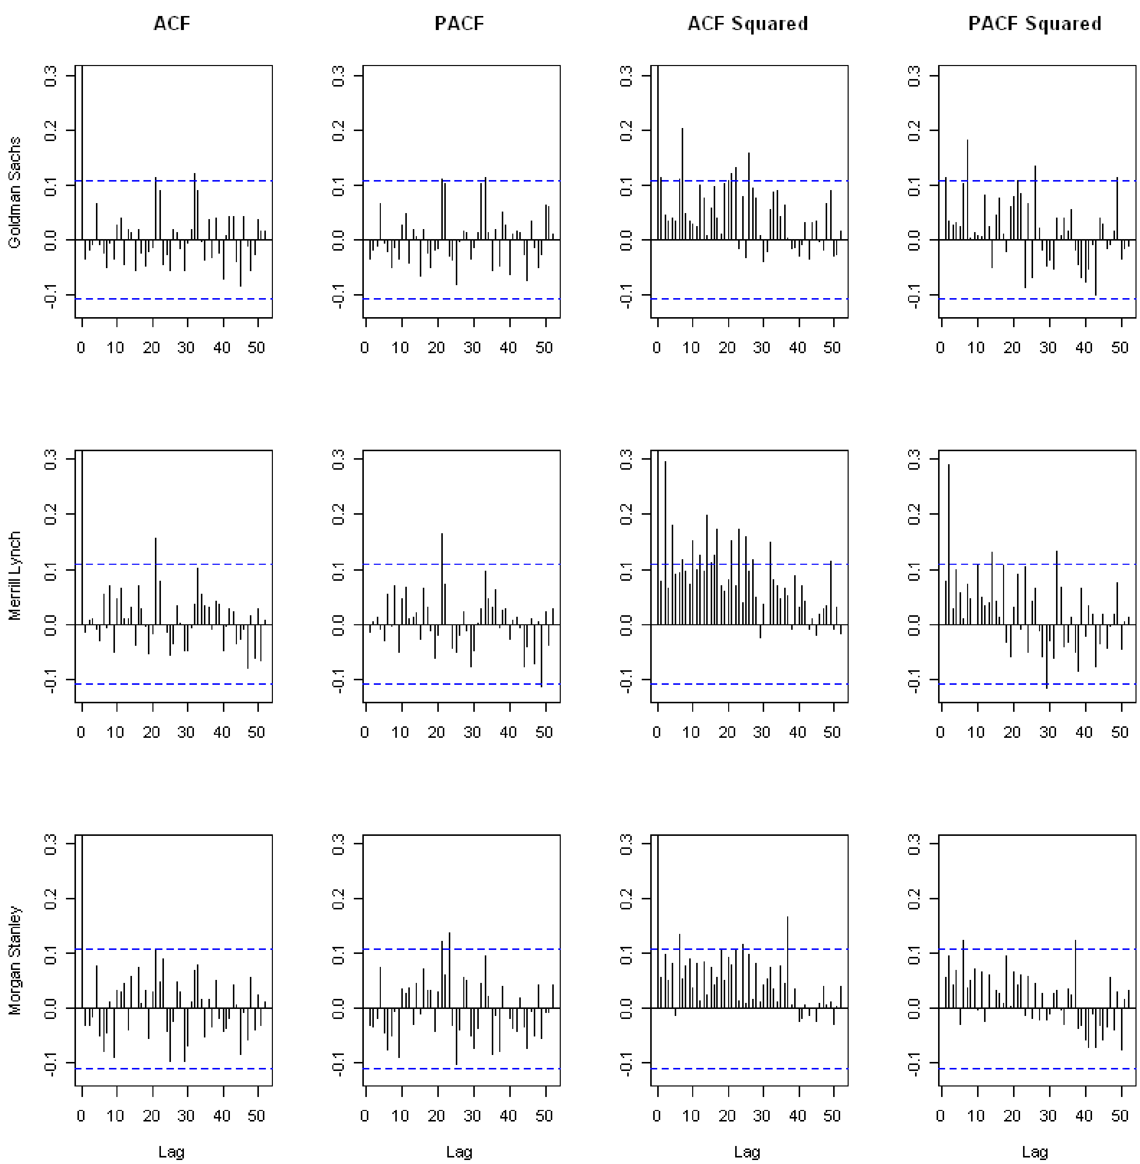

Figure 6: ACF, PACF, ACF Squared, and PACF Squared of Goldman Sachs, Merrill Lynch, and Morgan Stanley's Weekly Differenced Log Base 10 Closing Prices from January 2002 to March 2008 Under the Standard and Poor's 500 Fund Index

To explore the statistical dependencies within all three financial stocks, we can look at their autocorrelation (ACF) and partial autocorrelation (PACF) plots of the series, as well as ACF and PACF plots of the series of the squared values (Figure 6). The ACF and PACF plots of the series explore remaining dependencies in the mean of the differenced log closing price, whereas the ACF and PACF plots of the squared series examine dependencies in the variance of the differenced log closing price. Looking at the ACF and PACF graphs using the weekly log differenced closing price data from January 2002 to March 2008 for the financials, we can see multiple significant lags for all stocks. In the top row of Figure 6, Goldman Sachs (GS) shows weak significant lag 21 and 32 in the ACF graph, as well as stronger significant lags 7 and 26 in both the ACF and PACF squared graphs. In Merrill Lynch's (MER) set of graphs found in the middle row of Figure 6, strong significant lag 21 is seen in both the ACF and PACF graphs. MER shows multiple significant lags in its ACF squared graph as early as lag 2 and 4. The PACF squared graph of MER also shows a strong significant lag 2. In the bottom set of graphs in Figure 6, two moderately significant lags are seen at 21 and 23 in the PACF graph of Morgan Stanley (MS). MS also shows significant lag 6 and 37 in both its ACF and PACF squared graphs. A noticeable feature in all of these graphs is the significant lag 21 appearing in one or more graphs of each of the three financials. If we look at the ACF and PACF graphs from mid-2003 to the end of 2006, which according to Figure 5 is a period of lower variability, we may see fewer significant lags in the graphs. Reproducing Figure 6 over this time period, we notice several changes. In all of the ACF and PACF graphs from mid-2003 to the end of 2006 for the financials, significant lags are seen at 4. Also, Morgan Stanley shows a significant lag 2 in its ACF and PACF graphs. Since the significant lags over this period have a stronger correlation than the longer January 2002 – March 2008 period (ρ=0.15 vs. ρ=0.10 in GS and MER and ρ=0.20 vs. ρ=0.10 in MS), more dependence on the mean is likely. In the ACF and PACF squared graphs of the financials, Goldman Sachs shows no significant lags in either graph. Merrill Lynch shows significant lag 24 in both the ACF and PACF squared graphs. Morgan Stanley shows no significant lag in the ACF squared graph, but does show significance at lag 24 in the PACF squared graph. Since the ACF and PACF squared graphs of the shorter period show weaker correlations than the longer periods for all financials except Morgan Stanley (ρ=0.10 vs. ρ=0.20 in GS, ρ=0.15 vs. ρ=0.30 in MER, and ρ=0.15 vs. ρ=0.15 in MS), the dependency between variances of closing prices is smaller for GS and MER over the smaller period compared to the original period (Jan. 2002 – March 2008), while MS has equal dependency between variances of closing prices over both periods. Over the longer period, more appropriate models for the financial stocks are necessary to better fit the financial companies' closing prices due to the significant lags seen in both the ACF and PACF.

IV. Methodology

To find a good model for the volatility of our financial data, namely the differenced log weekly closing price, it will be necessary to fit multiple models and check the goodness of fit for each model. To take into account the changes in the variance of the differenced log weekly closing prices (Figure 5), the first model we will use to analyze the data is the ARCH(p) model. This model will relate the variance of current closing prices with the variances of past closing prices.

To define the ARCH(p) model, we shall let $X_t = r_t - \mu_t$ represent differenced log return at week t with r_t equal to a differenced weekly log closing price and μ_t equal to the mean of all of the differenced weekly log closing prices[9]. If $\{X_t\}$ follows an ARCH(p) model, we have that

$$X_t = \sigma_t \varepsilon_t, \qquad \sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i X_{t-i}^2,$$

where $\{\varepsilon_t\}$ contains random independent and identically distributed (iid) random variables with a mean of 0 and variance of 1, α_0>0, and α_i≥0 for all i >0. This model implies that $\{X_t\}$ looks like Gaussian white noise, but $\{X_t^2\}$ does not. The X_t^2 value depends on past differenced weekly log closing prices.

For example, consider the following ARCH(7) model (which may be an appropriate model for Goldman Sachs), where

$$X_t = \sigma_t \varepsilon_t, \quad \sigma_t^2 = \alpha_0 + \alpha_1 X_{t-1}^2 + \alpha_2 X_{t-2}^2 + \alpha_3 X_{t-3}^2 + \alpha_4 X_{t-4}^2 + \alpha_5 X_{t-5}^2 + \alpha_6 X_{t-6}^2 + \alpha_7 X_{t-7}^2,$$

and $\{\varepsilon_t\}$ is iid normal with a mean of 0 and variance of 1, α_0>0, and α_i≥0 for all i >0. This model says that the variance at week t depends linearly on the square values of the past 7 weeks.

We can fit this model in the R software package using the `garch` function in the `tseries` R library[10]. For some series there will be evidence that a subset of the $\{\alpha_i\}$ coefficients (i = 1,…, 7 in this case) will be zero. In that case we need to fit a subset ARCH model. This is not possible using the `garch` function, so we need to calculate the maximum likelihood estimates by other means. In this case we use a conditional likelihood method, which we now describe. Suppose for example that we wish to fit the following subset ARCH(7) model:

$$X_t = \sigma_t \varepsilon_t, \qquad \sigma_t^2 = \alpha_0 + \alpha_7 X_{t-7}^2,$$

where $\{\varepsilon_t\}$ is iid normal with a mean of 0 and variance of 1. For a time series of length N, let our data vector be $\underline{X} = (X_1, \ldots, X_N)$. Then the likelihood for $\underline{X}$ with parameters α_0 and α_7 is equal to

$$\begin{aligned} f(\underline{X} \mid \alpha_0, \alpha_7) = \prod_{t=8}^{N} f(X_t \mid X_{t-1}, \ldots, X_1, \alpha_0, \alpha_7) \bullet \\ f(X_7 \mid X_6, \ldots, X_1, \alpha_0, \alpha_7) \bullet \\ f(X_6 \mid X_5, \ldots, X_1, \alpha_0, \alpha_7) \bullet \\ \vdots \\ f(X_1 \mid \alpha_0, \alpha_7), \end{aligned} \tag{1}$$

In practice, it can be hard to provide the distributions of $f(X_7 \mid X_6, \ldots, X_1, \alpha_0, \alpha_7)$, $f(X_6 \mid X_5, \ldots, X_1, \alpha_0, \alpha_7)$, … , $f(X_1 \mid \alpha_0, \alpha_7)$. Instead, we ignore these terms, and work with the conditional likelihood, conditional on $X_{1, \ldots,} X_7$, which is given by

$$f(\underline{X} \mid X_1, \ldots, X_7, \alpha_0, \alpha_7) = \prod_{t=8}^{N} f(X_t \mid X_{t-1}, \ldots, X_1, \alpha_0, \alpha_7). \tag{2}$$

For each t, $X_t \mid X_{t-1}, \ldots, X_1, \alpha_0, \alpha_7$ is independent normal with mean 0 and variance of $\alpha_0 + \alpha_7 X_{t-7}^2$. Thus,

$$f(X_t \mid X_{t-7}) = \frac{1}{\sqrt{2\pi(\alpha_0 + \alpha_7 X_{t-7}^2)}} \exp\left(\frac{-(X_t - 0)^2}{2(\alpha_0 + \alpha_7 X_{t-7}^2)} \right). \tag{3}$$

After taking the logs of both sides,

$$\log\ f(X_t \mid X_{t-7}) = -½ \log(2\pi) - ½ \log(\alpha_0 + \alpha_7 X_{t-7}^2) - \left(\frac{X_t^2}{2(\alpha_0 + \alpha_7 X_{t-7}^2)} \right). \tag{4}$$

Thus, the conditional log likelihood, which is the log of (2) is given by

$$\sum_{t=8}^{N} \log f(X_t \mid X_{t-7}) = -½(\text{N-8})\log(2\pi) - ½ \sum_{t=8}^{N} \log(\alpha_0 + \alpha_7 X_{t-7}^2) - ½ \sum_{t=8}^{N} \frac{X_t^2}{(\alpha_0 + \alpha_7 X_{t-7}^2)}. \tag{5}$$

We calculate the maximum conditional likelihood estimates of α_0 and α_7 by maximizing the log conditional likelihood with respect to α_0 and α_7 (we did this in the R software package using the `nlm` function). In our analyses, we condition on the same amount of data at the start of the series (in this case seven values). Based on these maximum conditional likelihood estimates, we can calculate the residuals using the formula

$$\hat{\varepsilon}_t = \frac{X_t}{\hat{\sigma}_t} = \frac{X_t}{\sqrt{\hat{\sigma}_t^2}} = \frac{X_t}{\sqrt{\hat{\alpha}_0 + \hat{\alpha}_7 X_{t-7}^2}}, \quad \text{where t > 7.} \qquad (6)$$

Additionally, we can add a volume variable to better fit our models. An ARCH(p) model containing both weekly closing price and average daily volume will be defined as follows. Again, we shall let $X_t = r_t - \mu_t$ represent the differenced log weekly closing price, where r_t equals a differenced weekly log closing price and μ_t equals the mean of all of the differenced weekly log closing prices and we will also let V_t = the average daily volume for week t = 1,…, N. Then our model becomes

$$\mathrm{X}_t = \sigma_t \varepsilon_t, \qquad \sigma_t^2 = \alpha_0 + \alpha_7 \mathrm{X}_{t-7}^2 + \beta \log(\mathrm{V}_t), \qquad (7)$$

where $\{\varepsilon_t\}$ contains random independent and identically distributed random variables with mean of 0 and variance of 1, and the set of $\{\alpha_i\}$, β along with $\{\log(V_t)\}$ guarantee that σ_t^2 >0 for all t. The likelihood for this model is

$$\begin{aligned} f(\underline{\mathrm{X}} \mid \alpha_0, \alpha_7, \beta) = \prod_{t=8}^{N} f(\mathrm{X}_t \mid \mathrm{X}_{t-1}, \ldots, \mathrm{X}_1, \alpha_0, \alpha_7, \beta) \bullet \\ f(\mathrm{X}_7 \mid \mathrm{X}_6, \ldots, \mathrm{X}_1, \alpha_0, \alpha_7, \beta) \bullet \\ f(\mathrm{X}_6 \mid \mathrm{X}_5, \ldots, \mathrm{X}_1, \alpha_0, \alpha_7, \beta) \bullet \\ \vdots \\ f(\mathrm{X}_1 \mid \alpha_0, \alpha_7, \beta). \end{aligned} \qquad (8)$$

Again, we will work with the conditional likelihood to rewrite our function as

$$f(\underline{X} \mid X_1,...,X_7,\alpha_0,\alpha_7,\beta) = \prod_{t=8}^{N} f(X_t \mid X_{t-1},...,X_1,\alpha_0,\alpha_7,\beta), \qquad (9)$$

where each t in $X_t \mid X_{t-1},...,X_1,\alpha_0,\alpha_7,\beta$ is independent normal with mean 0 and variance of $\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t)$. Thus,

$$f(X_t \mid X_{t-7}) = \frac{1}{\sqrt{2\pi(\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t))}} \exp\left(\frac{-(X_t - 0)^2}{2(\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t))} \right). \qquad (10)$$

After taking the logs of both sides, we obtain

$$\log\ f(X_t \mid X_{t-7}) = -½ \log(2\pi) - ½ \log\left(\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t)\right) - \left(\frac{X_t^2}{2(\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t))} \right). \qquad (11)$$

Thus, the conditional log likelihood, which is the log of (9) is given by

$$\sum_{t=8}^{N} \log f(X_t \mid X_{t-7}) = -½(\text{N-8})\log(2\pi) - ½ \sum_{t=8}^{N} \log(\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t)) - \sum_{t=8}^{N} \left(\frac{X_t^2}{\left(\alpha_0 + \alpha_7 X_{t-7}^2 + \beta \log(V_t)\right)} \right). \qquad (12)$$

We calculate the maximum conditional likelihood estimates of α_0, α_7, and β by maximizing the log conditional likelihood with respect to α_0, α_7, and β (again we did this in the R software package using the `nlm` function). Again for each model we condition on the same amount of data at the start of the series (seven values). Based on these maximum conditional likelihood estimates, we can calculate the residuals using the formula

$$\hat{\varepsilon}_t = \frac{X_t}{\hat{\sigma}_t} = \frac{X_t}{\sqrt{\hat{\sigma}_t^2}} = \frac{X_t}{\sqrt{\hat{\alpha}_0 + \hat{\alpha}_7 X_{t-7}^2 + \hat{\beta} \log(V_t)}}, \quad \text{where } t > 7. \qquad (13)$$

V. Results

Goldman Sachs (ARCH(7))

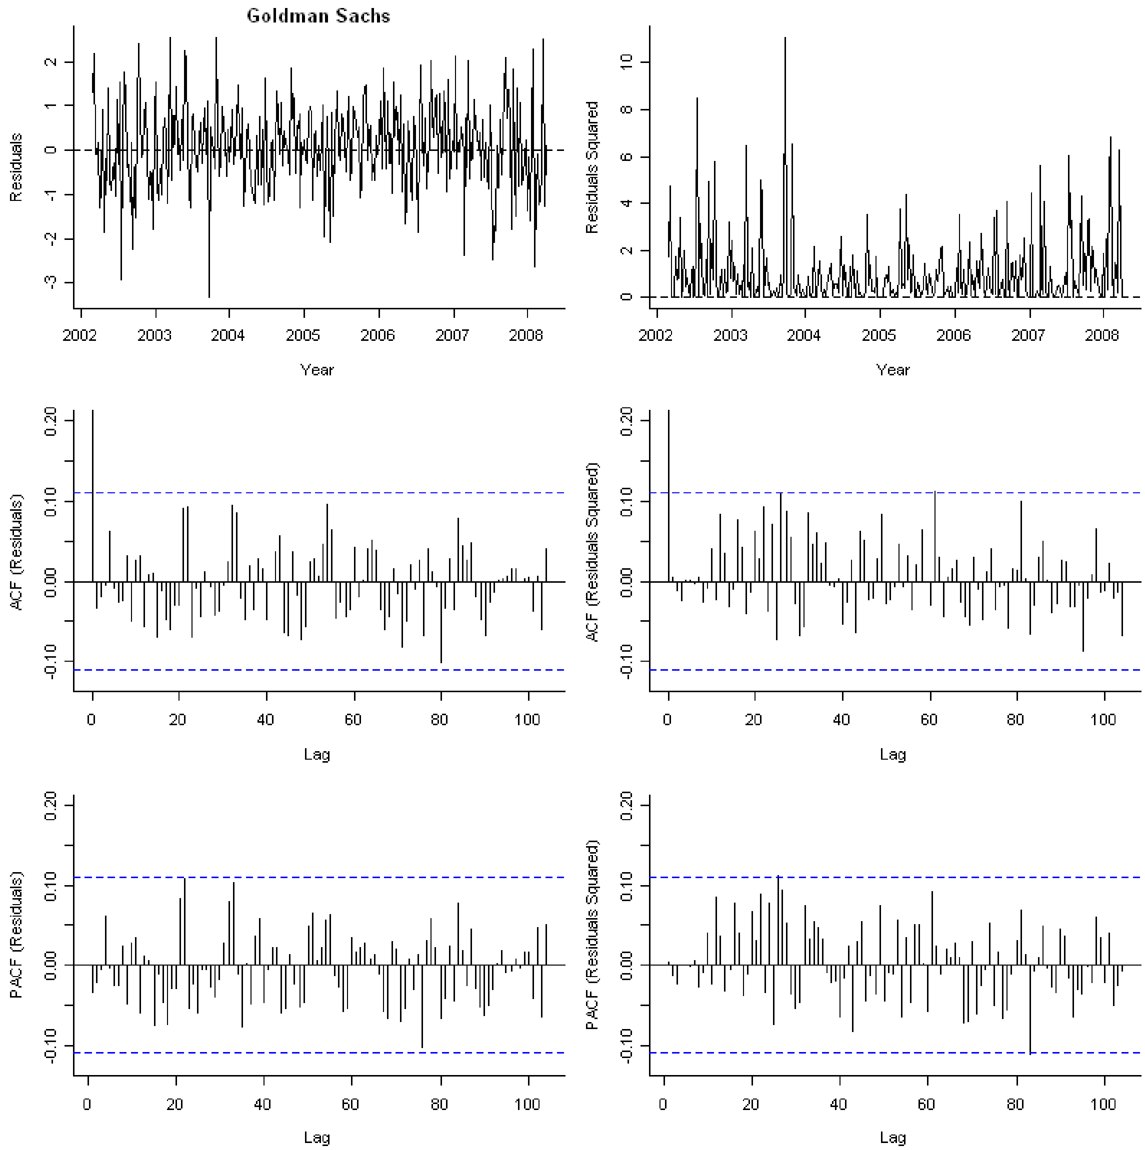

Figure 7: Residuals, Residuals Squared, ACF, and PACF Using Residuals and Residuals Squared of Goldman Sachs under the ARCH(7) model from January 2002 to March 2008

Seeing the earliest significant lag at 7 for Goldman Sachs in the ACF and PACF squared graphs of the differenced log weekly closing prices (Figure 6), an ARCH(7) model will be our first model used to try to interpret Goldman Sachs' past volatility. Under the ARCH(7) model, we see that Goldman Sachs fits fairly well. To illustrate this, we can look at the residual and residual squared, as well as ACF and PACF of the residuals and residuals squared under this model (Figure 7). The residuals graph appears to be independent, but not identically distributed. A changing variance over time can be seen in the residuals squared graph. In the graph, two rather large negative residuals are seen between 2002 and 2004 and increasing residuals are noted from 2006 to 2008. Under the ACF using residuals, all lags (except 0) appear to fall within the white noise bounds (which matches with what we saw in our exploratory data analysis). In addition, the PACF using residuals shows no strong significant lags outside of the bounds. This indicates evidence for white noise. In the ACF for the residuals squared, no nonzero lags are strongly significantly different from zero. Also, the PACF using residuals squared shows no lags outside of the bounds. This indicates no dependence between any two differenced log closing prices over the weeks. Except for the evidence of variance changes with time, the fit of the ARCH model looks reasonable.

Goldman Sachs (ARCH(7))

Parameter	Coefficient	t-statistic	Pr(>|t|)
$\hat{\alpha}_0$	0.0001436	4.930	8.24 e-07
$\hat{\alpha}_1$	0.09525	1.535	0.1247
$\hat{\alpha}_2$	0.03660	0.563	0.5732
$\hat{\alpha}_3$	0.01346	0.262	0.7937
$\hat{\alpha}_4$	0.03972	0.595	0.5522
$\hat{\alpha}_5$	0.03665	0.897	0.3695
$\hat{\alpha}_6$	0.04917	0.721	0.4710
$\hat{\alpha}_7$	0.1975	2.434	0.0150

Table 1: ARCH(7) Summary for Goldman Sachs

To estimate the parameters of the ARCH(7) process for Goldman Sachs, we used the `garch` command from the `tseries` library in R. A summary of the coefficients in this model is given in Table 1. To determine if any of the parameters are unnecessary, we will perform a hypothesis test on each parameter $\{\alpha_j\}$ to test its significance in the model. For example, for our first parameter α_0, we test H_0: $\alpha_0 = 0$ versus H_1: $\alpha_0 \neq 0$. With the weekly differenced log closing prices in Goldman Sachs, the R software estimates the coefficients given in Table 1, using maximum likelihood methods. The *t-statistic* in the table is equal to the estimated coefficient divided by its standard error and is used in testing the significance of each coefficient[11]. After obtaining a t-statistic, we can calculate a p-value, which is equal to the probability of obtaining a t-statistic value as extreme as the one observed in the dataset, given the null hypothesis (H_0) is true. Assume for our α_0 hypothesis test, our significance level is 5%. Then, if the probability of receiving a value larger than the absolute value of the t-statistic is less than 5%, we will reject H_0 and conclude H_1. However, if the probability of receiving a value larger than the absolute value of the t-statistic is greater than 5%, we will fail to reject H_0 and conclude H_0. Because the probability of receiving a value larger than the absolute value of the t-statistic is less than 5% (p-value = 8.24 e-07), we can reject H_0 and conclude H_1. This indicates that α_0 will remain in our model. Of the 8 parameters, α_0 and α_7 fall within the 5% significance level and will remain in Goldman Sach's ARCH(7) refined model; the other parameters do not. We refit the model using conditional maximum likelihood. The estimated subset ARCH(7) model we obtain is:

$$X_t = \hat{\sigma}_t \varepsilon_t, \qquad \hat{\sigma}_t^2 = 0.0001436 + 0.1975 X_{t-7}^2, \tag{14}$$

where $\{\varepsilon_t\} \sim$ iid N(0,1).

Merrill Lynch (ARCH(2))

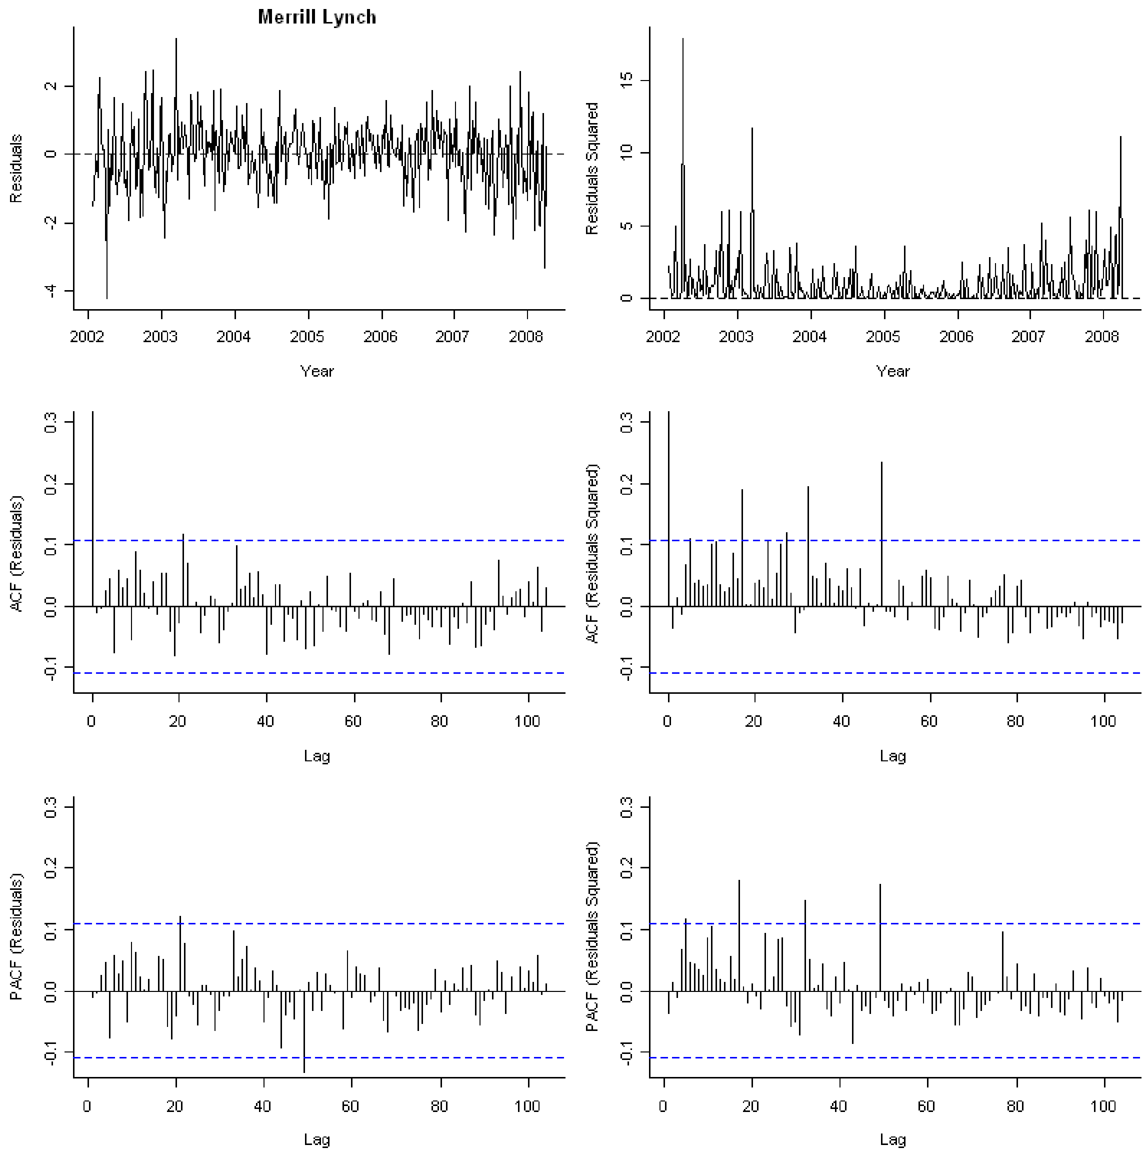

Figure 8: Residuals, Residuals Squared, ACF, and PACF Using Residuals and Residuals Squared of Merrill Lynch under ARCH(2) model from January 2002 to March 2008

Because of a significant lag 2 in Figure 6 for Merrill Lynch in the ACF squared graph, we will fit an ARCH(2) model for the differenced log closing price. The residual diagnostic plots for Merrill Lynch are shown in Figure 8. Looking at the residual plots, slight negative skewness is seen beginning in 2008. Despite this, the residuals appear to be independent, but not identically distributed. Looking at the residuals squared graph, three residuals squared are seen above 10 as well as increasing residuals from 2006 to

2008. Under the ACF and PACF graphs using the residuals squared, three large significant lags are noted at lag 17,32, and 49. Lag 21 is significantly different from zero for the ACF and PACF of the residuals, indicating no evidence of white noise. Because of the significant lags and large squared residuals, a better model may exist to exhibit Merrill Lynch's dataset.

Merrill Lynch (ARCH(2))

Parameter	Coefficient	t-statistic	Pr(>|t|)
$\hat{\alpha}_0$	0.0001804	9.735	< 2e-16
$\hat{\alpha}_1$	0.1675	2.329	0.019839
$\hat{\alpha}_2$	0.3001	3.446	0.000568

Table 2: ARCH(2) Summary for Merrill Lynch

Using the R software, a summary of the coefficients for Merrill Lynch's dataset under the ARCH(2) model is given in Table 2. Looking at the p-values in column four of Table 2, all fall within the 5% significance level so there is no need to refine this model. Therefore, the estimated ARCH(2) model for Merrill Lynch is

$$X_t = \hat{\sigma}_t \varepsilon_t, \qquad \hat{\sigma}_t^2 = 0.0001804 + 0.1675X_{t-1}^2 + 0.3001X_{t-2}^2, \tag{15}$$

where $\{\varepsilon_t\}$ ~ iid N(0,1). In this model, the lag 2 volatility is greater than the lag one volatility.

Morgan Stanley (ARCH(6))

Due to the significant lag 6 found in Figure 6 for Morgan Stanley under the ACF and PACF graphs using residuals squared, we will attempt to fit the Morgan Stanley differenced log closing prices with an ARCH(6) model. To test the goodness of fit for Morgan Stanley under the ARCH(6) model, we shall again analyze the residuals, residuals squared, and ACF, and PACF of the residuals and squared residuals (Figure 9). The residuals appear independent but not equally distributed. Looking at the residuals

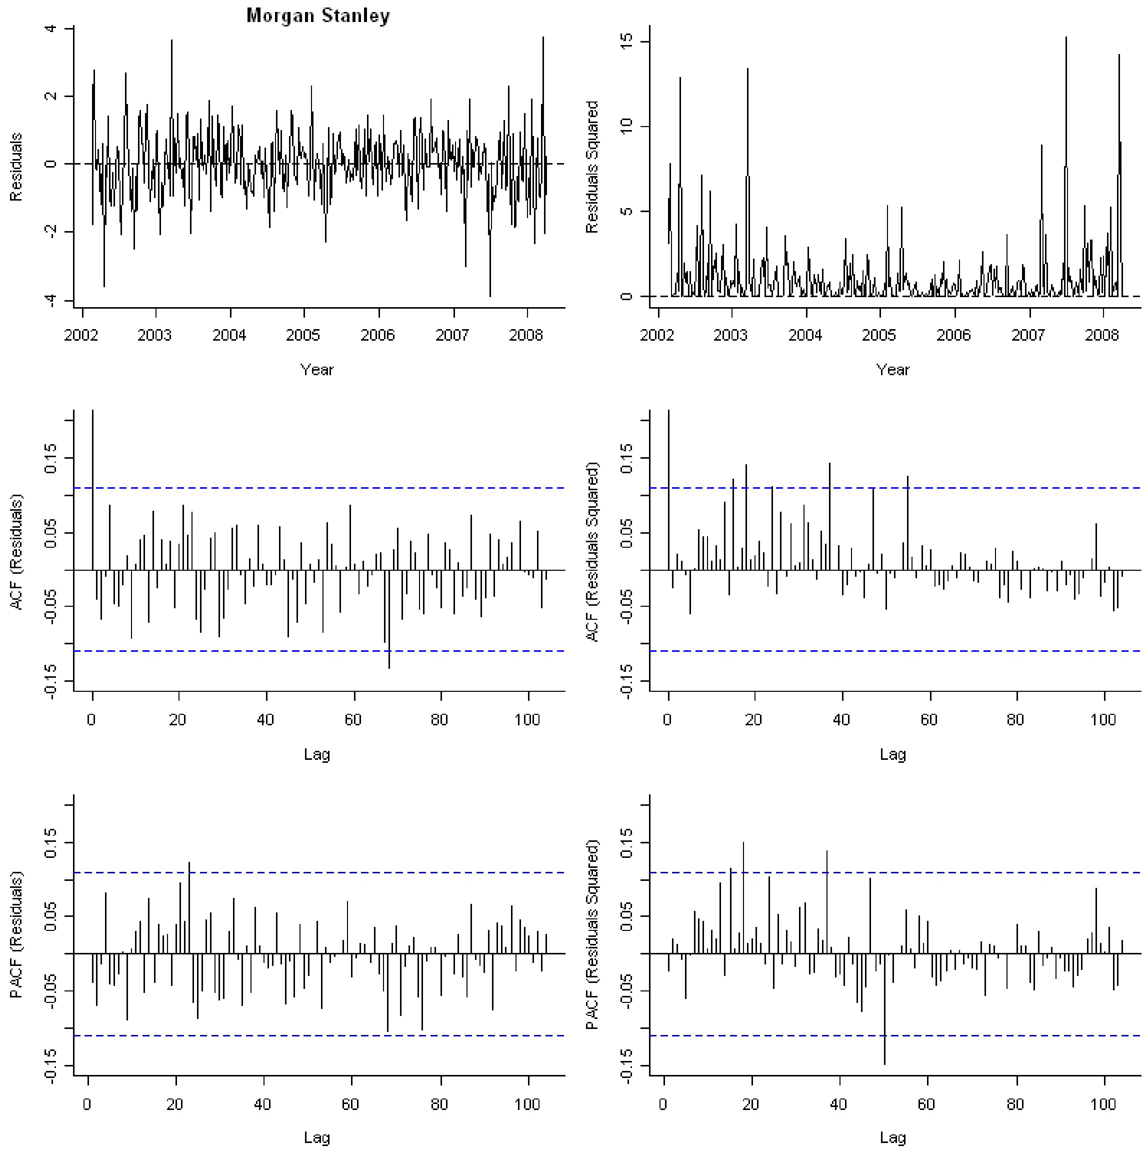

Figure 9: Residuals, Residuals Squared, ACF, and PACF Using Residuals and Residuals Squared of Morgan Stanley under ARCH(6) model from January 2002 to March 2008

squared, increased variability is noted between 2002 and 2003, as well as 2007 to the end of March 2008. In addition, there is a small increase in variability in the beginning of 2005. The ACF using residuals shows a significant lag 68, but the remainder of the series appears as white noise. Under the ACF graph using residuals squared, significant lags at 15, 18, 24, 37, 47, and 55 are seen with the remaining lags fairly close to zero. Using residuals, the PACF graph shows a significant lag 23, as well as other lags appearing close to significance. Like the ACF graph using residuals, the PACF graph using residuals demonstrates white noise. Significant lags 15, 18, 37, and 50 also exist in the PACF graph using residuals squared. Due to numerous significant lags in the residuals squared graphs and increased variability, it is necessary to test another model for better fit in Morgan Stanley's dataset.

Morgan Stanley (ARCH(6))

Parameter	Coefficient	t-statistic	Pr(>|t|)
$\hat{\alpha}_0$	0.0001731	6.45	1.12e-10
$\hat{\alpha}_1$	0.02427	0.789	0.4299
$\hat{\alpha}_2$	0.1368	2.198	0.0279
$\hat{\alpha}_3$	0.08179	1.188	0.2349
$\hat{\alpha}_4$	0.08139	0.856	0.3920
$\hat{\alpha}_5$	8.125e-15	1.25e-13	1.000
$\hat{\alpha}_6$	0.1613	4.015	0.0000594

Table 3: ARCH(6) Summary for Morgan Stanley

In the R software, we estimate the parameters for Morgan Stanley's process under the ARCH(6) model. A summary of the coefficients in this model is given in Table 3. Dropping α_1, α_3, α_5, and α_6 from the model (they are not significant at the 5% level), we obtain the refined subset ARCH(6) model for Morgan Stanley fit using conditional likelihood of:

$$\mathrm{X}_t = \hat{\sigma}_t \varepsilon_t, \qquad \hat{\sigma}_t^2 = 0.0001731 + 0.1368\mathrm{X}_{t-2}^2 + 0.1613\mathrm{X}_{t-6}^2, \tag{16}$$

where $\{\varepsilon_t\} \sim$ iid N(0,1).

We now add the volume variable to our ARCH(p) model to try and account for the changes in variance that we observe in each series. To create our new variance equation, we will use the `nlm` command in R to estimate the α and β values of equation (7) for each of the financials. We will assume the same parameters for each ARCH(p) subset model and only add the β parameter in our new variance calculations.

Goldman Sachs (ARCH(7))

Fitting model (7) to GS in R, we estimate our parameters to be

Parameter	Coefficient
$\hat{\alpha}_0$	-0.0440
$\hat{\alpha}_7$	0.3
$\hat{\beta}$	0.0071

Table 4: Goldman Sachs Subset ARCH(7) model + Base 10 Log(Volume) Estimate Using R

Thus,

$$X_t = \hat{\sigma}_t \varepsilon_t, \qquad \hat{\sigma}_t^2 = -0.0440 + 0.3\, X_{t-7}^2 + 0.0071 \log(V_t), \tag{17}$$

where $\{\varepsilon_t\}$ ~ iid N(0,1). Note that although the α_0 estimate is negative, the minimum log (V_t) value (6.216298) multiplied by the β estimate (0.0071) will always be large enough to yield a positive variance ($\hat{\sigma}_t^2$ = 0.000136 at our minimum log (V_t) value). In order to find whether or not the new model creates a better fit for the dataset, we will need to check the ACF and PACF graphs (as shown in the next section).

Merrill Lynch (ARCH(2))

Fitting model (7) to MER using R software, we estimate our parameters to be

Parameters	Coefficients
$\hat{\alpha}_0$	-0.0433
$\hat{\alpha}_1$	0.3
$\hat{\alpha}_2$	0.3
$\hat{\beta}$	0.0069

Table 5: Merrill Lynch Subset ARCH(2) model + Base 10 Log(Volume) Estimate Using R

Thus,

$$X_t = \hat{\sigma}_t \varepsilon_t, \quad \hat{\sigma}_t^2 = -0.0433 + 0.3\, X_{t-1}^2 + 0.3\, X_{t-2}^2 + 0.0069 \log(V_t), \tag{18}$$

where $\{\varepsilon_t\}$ ~ iid N(0,1). Again, despite the negative α_0 estimate, the minimum log (V_t) value (6.292167) multiplied by the β estimate (0.0069) will always be large enough to yield a positive variance ($\hat{\sigma}_t^2$ = 0.000116 at our minimum log (V_t) value). To check if the new variance estimate better fits the dataset, we will need to check the ACF and PACF graphs for Merrill Lynch (shown in the next section).

Morgan Stanley (ARCH(6))

Fitting model (7) to MS, the R software estimates our parameters to be

Parameters	**Coefficients**
$\hat{\alpha}_0$	-0.0435
$\hat{\alpha}_2$	0.3
$\hat{\alpha}_6$	0.3
$\hat{\beta}$	0.0069

Table 6: Morgan Stanley Subset ARCH(6) model + Base 10 Log(Volume) Estimate Using R

Thus,

$$X_t = \hat{\sigma}_t \varepsilon_t, \qquad \hat{\sigma}_t^2 = -0.0435 + 0.3\,X_{t-2}^2 + 0.3\,X_{t-6}^2 + 0.0069 \log(V_t), \qquad (19)$$

where $\{\varepsilon_t\}$ ~ iid N(0,1). Again, despite the negative $\hat{\alpha}_0$, the minimum log (V_t) value (6.35112) multiplied by the β estimate (0.0069) will always be large enough to yield a positive variance ($\hat{\sigma}_t^2$ = 0.000323 at our minimum log (V_t) value). To see if the new variance estimate better fits Morgan Stanley's dataset, we will need to check its ACF and PACF graphs. This is demonstrated in the next section.

Diagnostic Checks for the Models Including Log Volume

Looking at the ACF and PACF graphs using residuals and residuals squared of the differenced log base 10 weekly closing prices after removal of the first 7 closing prices and addition of log base 10 volume to the model of Goldman Sachs (top row), Merrill Lynch (middle row), and Morgan Stanley (bottom row) in Figure 10, we notice the

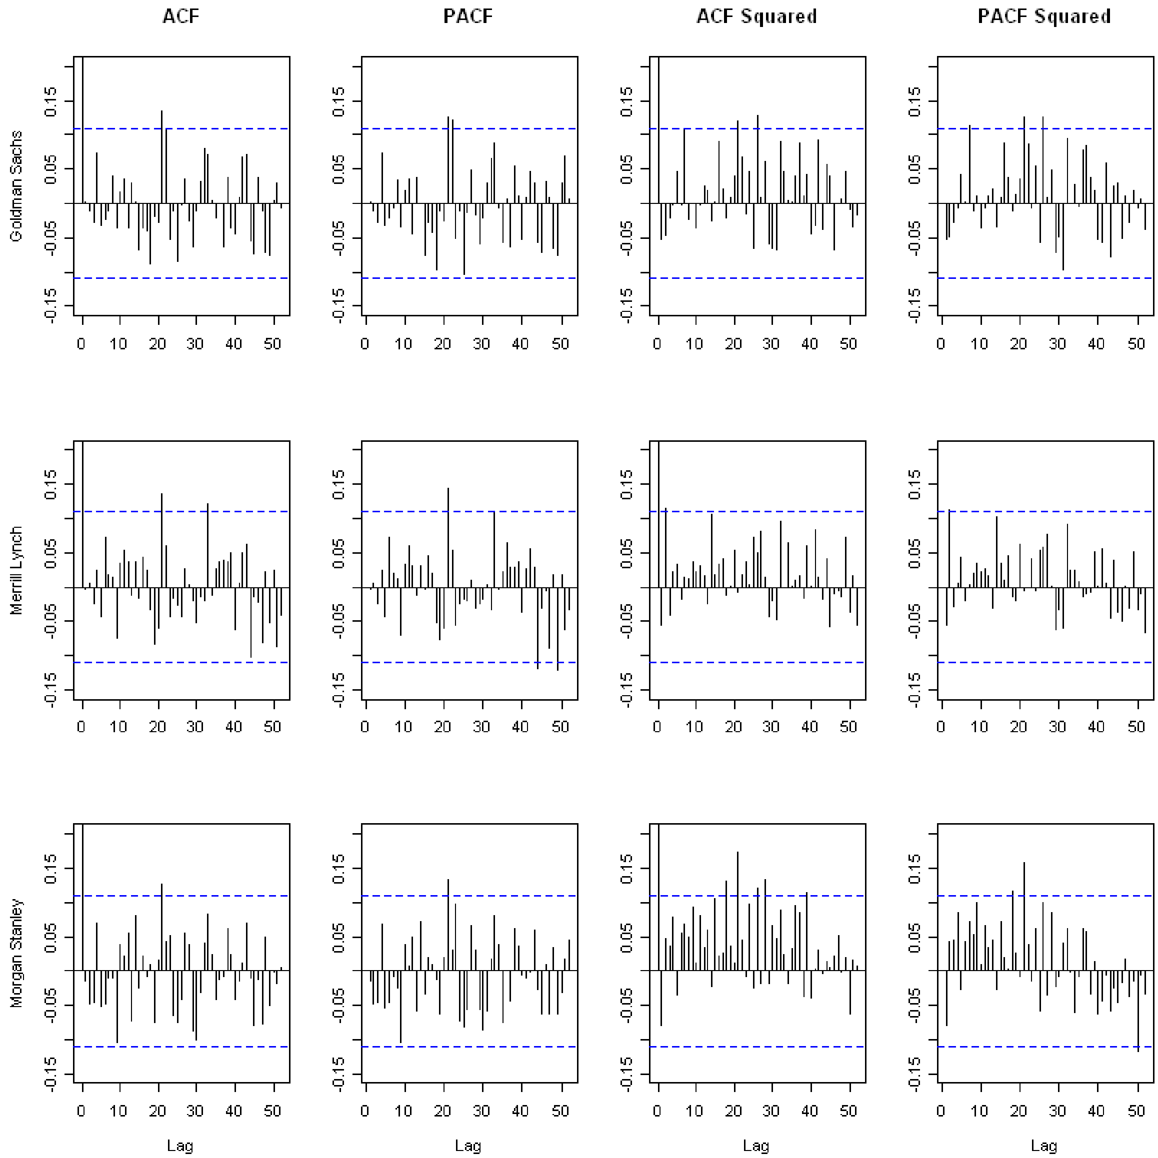

Figure 10: Estimated ACF and PACF of Differenced Log Base 10 Weekly Closing Prices With First Seven Dates Removed and Log Base 10 Volume Added into the Variance Equation; Goldman Sachs (top row), Merrill Lynch (middle row), and Morgan Stanley (bottom row)

correlations and number of significant lags reduced when compared to Figure 6, which did not incorporate volume into its model. The ACF and PACF graphs using residuals both show significant lags 21 and 22. Since these lags are fairly close to the bounds, the ACF and PACF graphs of GS using residuals both show white noise. Both the ACF and PACF graphs using residuals squared show lags 7, 21, and 26 for Goldman Sachs. Although these lags are significant, their correlation values are reduced from Figure 6. For Merrill Lynch, both the ACF and PACF graphs using residuals show significant lags

21 and 33. Also, the PACF graph shows lags 44 and 49 as significant. In the ACF and PACF graphs using residuals squared, lag 2 appears significant in both, but at a correlation value much lower than in Figure 6. In Morgan Stanley's ACF and PACF graphs using residuals, lag 21 shows significance in both. Both graphs appear to show white noise in the remaining lags. Multiple significant lags are seen in both the ACF and PACF graphs of Morgan Stanley using residuals squared. In the ACF graph, lags 18, 21, 26, 28, and 39 show significance while lags 18, 21, and 50 appear as significant in the PACF graph. Moreover, the ACF and PACF squared graphs of Morgan Stanley produce larger significant lags (at 21) than were produced in Figure 6. Another time series analysis may be necessary to evaluate the reason behind the increased correlation in this lag.

VI. Exploratory Time Series Analysis (Multivariate)

In order to check for any relationships between Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS), we will now look at pairwise scatterplots of the three weekly differenced log series and their residuals (Figure 11). Looking at the graphs, a moderate correlation between all stocks becomes visible (GS vs. MER: Adj. $R^2 = 0.5643$, GS vs. MS: Adj. $R^2 = 0.548$, MER vs. MS: Adj. $R^2 = 0.5772$). Comparing the distributions, a lower variability is seen between Goldman Sachs and Morgan Stanley and a slightly larger variability between Goldman Sachs and Merrill Lynch. The residual graphs appear to have a mean of zero for all three of the financials, but show unequal distributions at the beginning and end of the time period. Larger residuals are seen from 2002 to 2003 and mid-2007 into 2008 for all companies.

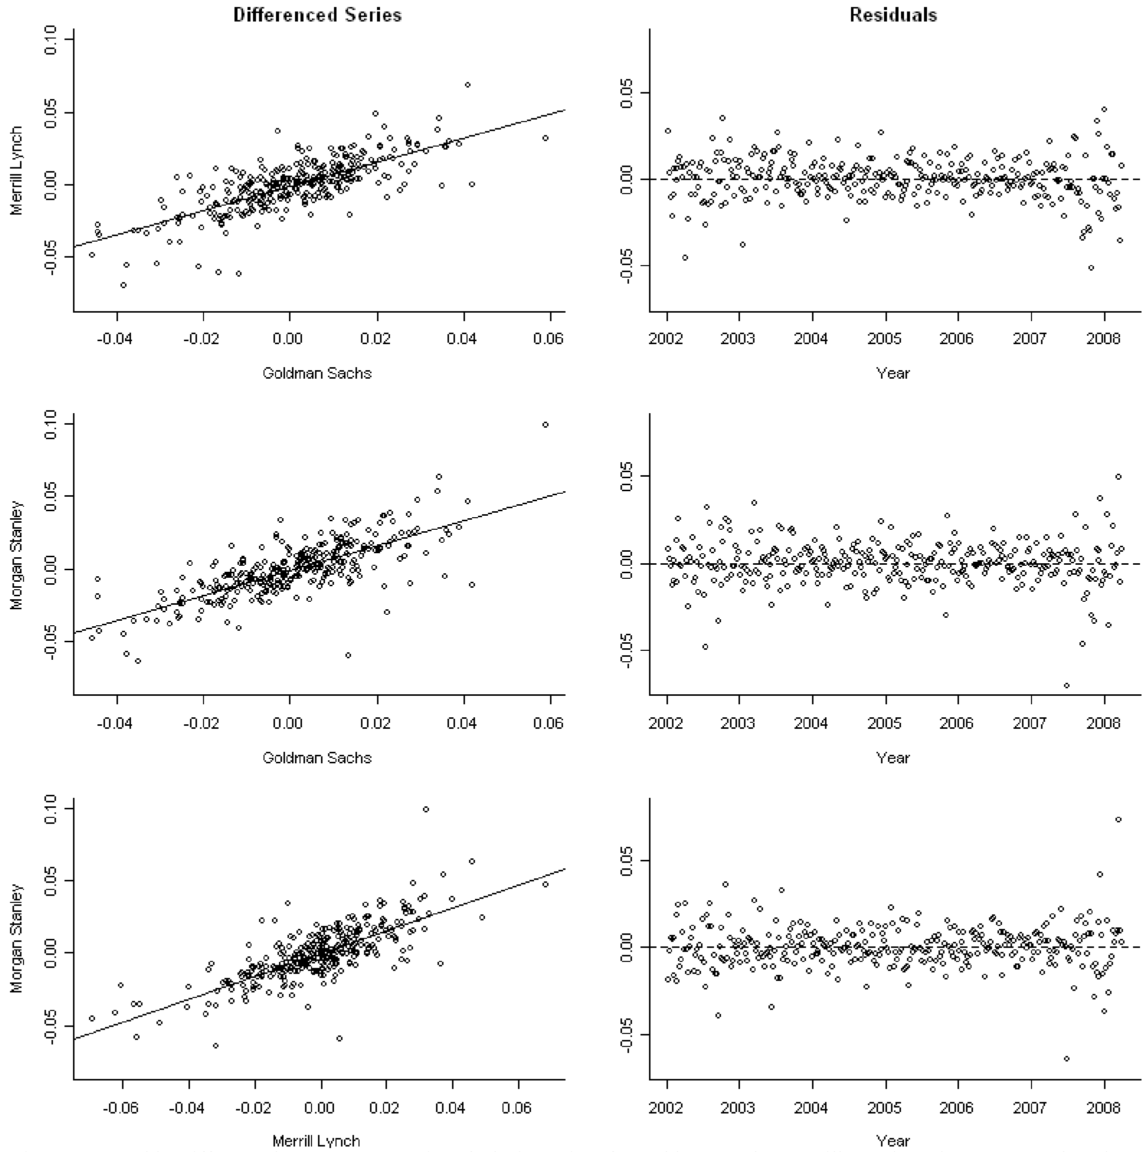

Figure 11: Weekly Differenced Log Base 10 Series of Closing Prices for Goldman Sachs, Merrill Lynch, and Morgan Stanley Plotted Against One Another and Residuals from January 2002 to March 2008

Looking closer at the relationship between the companies, we plot differenced log non-squared and squared cross correlation functions (CCF) to look for any potential significant lags (Figure 12). For example, the CCF of lag *h* between series $\{X_t\}$ and $\{Y_t\}$ is given by:

$$\rho_{xy}(h) = \text{corr}(X_{t+h}, Y_t). \qquad (20)$$

A cross correlation function is useful in determining whether there is a relationship in closing prices between any two series across different lags.

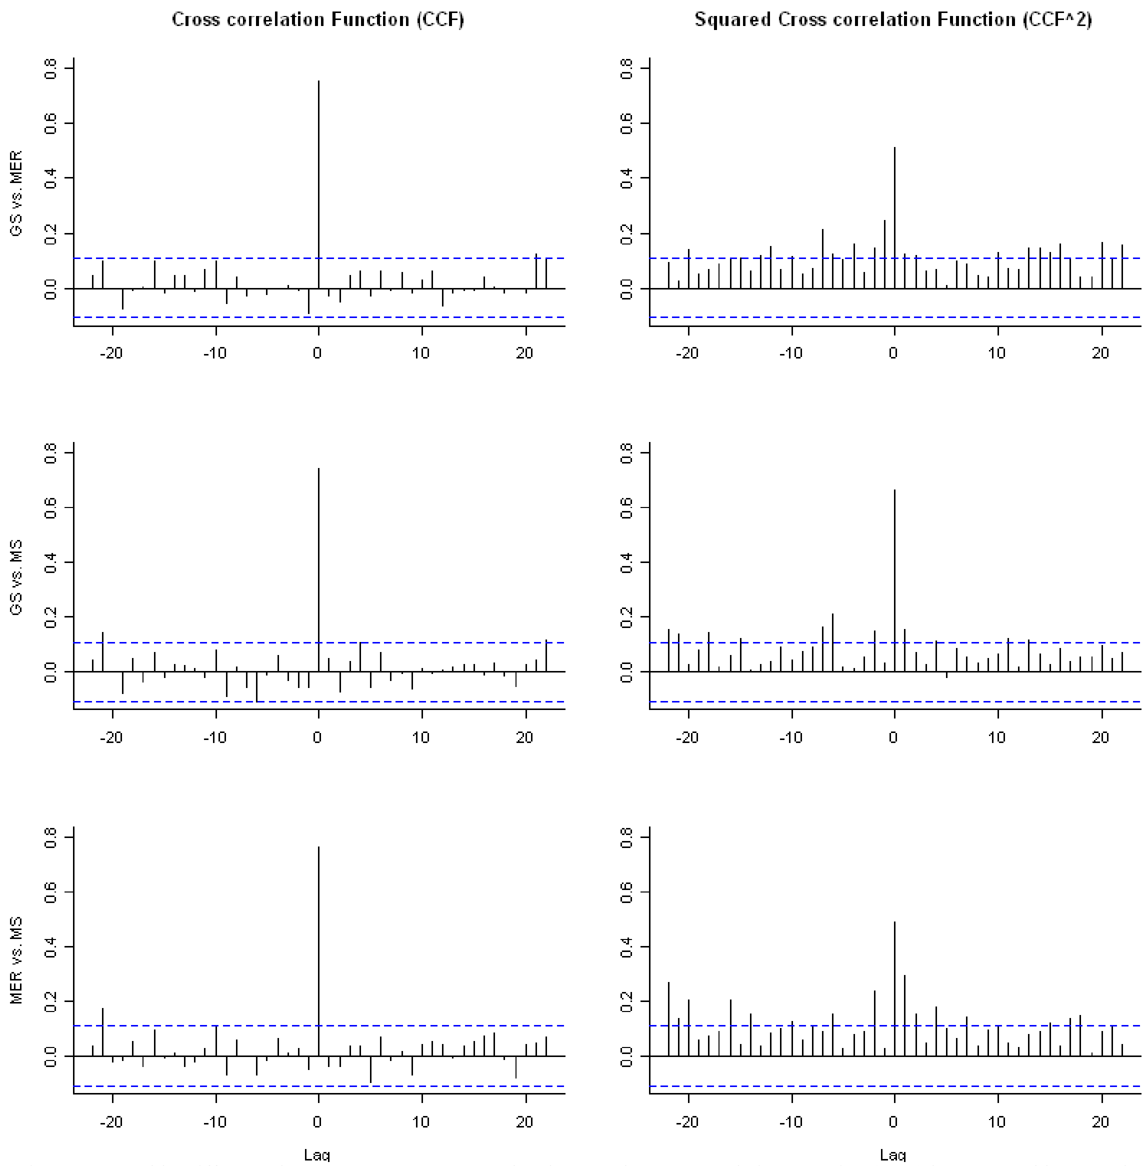

Figure 12: Weekly Differenced Log Base 10 Non-Squared and Squared Cross Correlation Function (CCF) between Goldman Sachs, Merrill Lynch, and Morgan Stanley from January 2002 to March 2008

Looking at the three CCF graphs using weekly differenced log base 10 closing prices on the left hand side of Figure 12, we see that lag 0 is most significant in all models, indicating a strong correlation between the three series. In addition, Goldman Sachs (GS) versus Merrill Lynch (MER) shows a significant lag at positive 21 while both GS versus Morgan Stanley (MS) and MER versus MS show significance at negative lag 21. GS versus MS also shows a significant lag at positive 22. Squaring the differenced log series and plotting the squared CCF will give us a better idea if there is a lag relationship in the closing prices. Despite all three CCF squared graphs (on the right hand side of Figure 12) showing lag 0 as most significant, multiple significant lags of weaker strengths than lag 0 also become present. This shows evidence of codependence between the variances of the weekly closing prices. Despite this, there is evidence that a significant correlation between the series at different lags exists and this should be explored in a future time series model.

VII. Conclusion

After performing both univariate and multivariate analysis of the closing prices and volumes for Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS) from January 2002 to March 2008, we can now answer the questions posed prior to the time series analysis.

After plotting the financials against each other in the cross correlation function (CCF) graphs, plotting the companies against each other, and examining the R^2 values, we can see that there was a correlation between GS, MER, and MS over January 2002 to March 2008. The strength of this correlation appears to be moderate as the R^2-value suggests in Figure 17. Lag 0 in the CCF of the differenced values and the CCF of the squared differenced values appears as the largest lag in all graphs. But, there also exists other significant lags in the CCF. Finally in the prediction graphs, we note significant residuals, kurtosis in the Q-Q plots, as well as a few small significant lags in the ACF and PACF graphs.

The variability of the closing prices for Goldman Sachs (GS), Merrill Lynch (MER), and Morgan Stanley (MS) does not appear constant over the time period under analysis. In 2002, 2003, 2007, and 2008, Figure 4 shows all three financial companies showing larger periods of variability than in 2004 to 2006. Also, increasing volume variability is seen in Figure 2 from 2007 to 2008. We found that incorporating the log volume as a covariate in the ARCH model substantially improved the fit.

Looking at the data from January 2002 to December 2007 does not appear to give us evidence of a downtrend in 2008 for any of the financials. However, increased variances in closing prices and volumes in 2007 may have provided investors an idea of the potential volatility in the year ahead.

Works Cited

[1] "Goldman Sachs." Wikipedia. 1 May 2008 <http://en.wikipedia.org/wiki/Goldman_Sachs>.

[2] Nikic, Tom, and Christopher Robertson. Value Line, 2008. 6 Apr. 2008 <http://www.valueline.com>.

[3] Yahoo! Finance. 1 May 2008 <http://finance.yahoo.com>.

[4] "Merrill Lynch." Wikipedia. 1 May 2008 <http://en.wikipedia.org/wiki/Merrill_Lynch>.

[5] "Morgan Stanley." Wikipedia. 1 May 2008 <http://en.wikipedia.org/wiki/Morgan_Stanley>.

[6] "Beta." Investopedia. 20 Apr. 2008 <http://www.investopedia.com/terms/b/beta.asp>.

[7] "Beta (Finance)." Wikipedia. 20 Apr. 2008 <http://en.wikipedia.org/wiki/Beta_%28finance%29>.

[8] "Autoregressive Conditional Heteroskedasticity." Wikipedia. 6 May 2008 <http://en.wikipedia.org/wiki/ARCH>.

[9] Tsay, Ruey S. Analysis of Financial Time Series. John Wiley & Sons, Inc., 2002.

[10] Trapletti, Adrian. "tseries." The R Project. <http://cran.r-project.org/web/packages/tseries/index.html>.

[11] Downing, Douglas, and Jeffrey Clark. Statistics the Easy Way. 3rd ed. Barron's Educational Series, Inc., 1997. 285.