



Seminar in finance

Modeling and Forecasting Volatility in Financial Time Series: An Econometric Analysis of the S&P 500 and the VIX Index.

Jon Lieberkind
Graduate Student, Department of Economics,
University of Copenhagen.

Submitted: 19/05/2009

| | |
|--|-----------|
| 1.1 INTRODUCTION..... | 3 |
| 1.2 METHODOLOGY | 3 |
| 1.3 DELIMITATION | 3 |
| 2.1 THE THEORY OF (G)ARCH..... | 4 |
| 2.2 INFORMAL TEST FOR (G)ARCH | 4 |
| 2.4 ESTIMATING AN AUTOREGRESSIVE UNIVARIATE MODEL FOR S&P 500 | 5 |
| 2.5 TESTING FOR ARCH | 6 |
| 3.1 GENERALIZED ARCH (GARCH) MODELS | 7 |
| 3.2 TESTING FOR GARCH | 7 |
| 4.1 THE VOLATILITY INDEX..... | 8 |
| 5.1 VOLATILITY FORECASTING IN PRACTICE | 10 |
| 5.2 COMPARING THE GARCH FORECASTS, ESTIMATED VOLATILITY AND THE VIX..... | 10 |
| 6.1 TRADING STRATEGIES AND THE VIX..... | 11 |
| 6.2 PORTFOLIO HEDGING USING VIX OPTIONS | 12 |
| 6.3 THE BENEFITS OF USING VIX OPTIONS AS A HEDGING TOOL | 13 |
| 7. CONCLUSION..... | 13 |
| 8. REFERENCES..... | 15 |
| APPENDIX 1 | 17 |
| APPENDIX 2 | 18 |
| APPENDIX 3 | 19 |

1.1 Introduction

During the current financial crisis it has become increasingly important for financial institutions as well as for the average retail investor to learn about volatility in the financial markets and know how to analyze and manage financial risk. This paper explores the utility of modeling and forecasting financial volatility. The aim is to provide investors with knowledge of the possibilities of hedging their portfolios by applying the aforementioned measures. By exploring some of the weaknesses of the ARCH(p) model it will be argued why a GARCH(1,1) model is useful in volatility forecasting. The GARCH(1,1) model is easily applied with only three unknown parameters to estimate and the model fits data relatively good. Furthermore, the GARCH(1,1) model is a widely used model specification when indexes on daily returns are modeled. The S&P 500 and the volatility index (VIX) are examined in order to analyze the robustness of VIX as an indicator for future volatility in the markets. The analysis will result in an evaluation of how investors can use VIX options as a hedging tool.

1.2 Methodology

By applying the much celebrated theory of (G)ARCH models pioneered by Engle (1982) and modified by Bollerslev and Taylor (1986), the paper quantifies the nature of S&P 500 stock returns¹ and the VIX from 1.1.1990-1.2.2009. The financial time series are streamed from the EcoWin Database. The econometric estimation methods; OLS, Maximum Likelihood and the Breusch-Pagan LM test for heteroskedasticity are applied on the autoregressive univariate models, ARCH(p) and GARCH(1,1). The econometric software used is OxMetrics/PcGive.

1.3 Delimitation

Due to the restricted size of this paper, it has not been possible to extend the econometric analysis. Otherwise it would have been beneficial to solve some of the weaknesses of the (G)ARCH models. Nelson (1990) demonstrates that the *IGARCH* model is strictly stationary, thereby solving the potential problem of unit roots in the ARMA process of the GARCH(1,1) model. Another problem with the (G)ARCH model is that positive and negative shocks have the same effect on volatility because it depends on the square of the previous shocks. In practice, the price of a financial asset responds differently to positive and negative shocks². Glosten et al (1993) suggest an asymmetric model that treats positive and negative shocks differently. In PcGive this is known as the *threshold*

¹ The paper uses closing prices. The S&P 500 is the most commonly used benchmark for the U.S. equity market, and the most popular underlying for U.S. equity derivatives. Therefore the S&P 500 is chosen as the market proxy.

² Negative shocks tend to increase volatility more than positive shocks do.

asymmetric model. *ARCH in Mean* is also an interesting extension of the (G)ARCH model as it investigates to which degree investors may require a higher return (risk premium) in periods of high volatility, Glosten et al (1993). Applying *EGARCH*³, Nelson (1991), may have improved the forecasts, as this model (unlike GARCH) has an asymmetric relationship between the error terms and future variances.

2.1 The Theory of (G)ARCH

Financial time series is often characterized by having a non-constant variance, also known as *heteroskedasticity*, and this particular form of time varying heteroskedasticity is known as *Autoregressive Conditional Heteroskedasticity* (ARCH), Engle (1982). Heteroskedasticity arises, as the stock markets are typically characterized by periods of high volatility and periods of low volatility. This is especially the case for high frequency data, i.e. S&P 500, but less clear at lower frequencies. Empirically stock returns are approximately uncorrelated in the long run, and it is difficult to predict the conditional mean of an asset return, Nielsen (2007). One can interpret this as a weak form of market efficiency, because perfect foresight would probably allow easy speculative gains. A stylized fact for many financial time series is a pattern of non-constant volatility, more precisely;

“...large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes.” Mandelbrot (1963).

Vís a vís Mandelbrot, one often observes what is referred to as *volatility clustering* in financial time series. Volatility clustering is often seen in data, which contains ARCH effects.

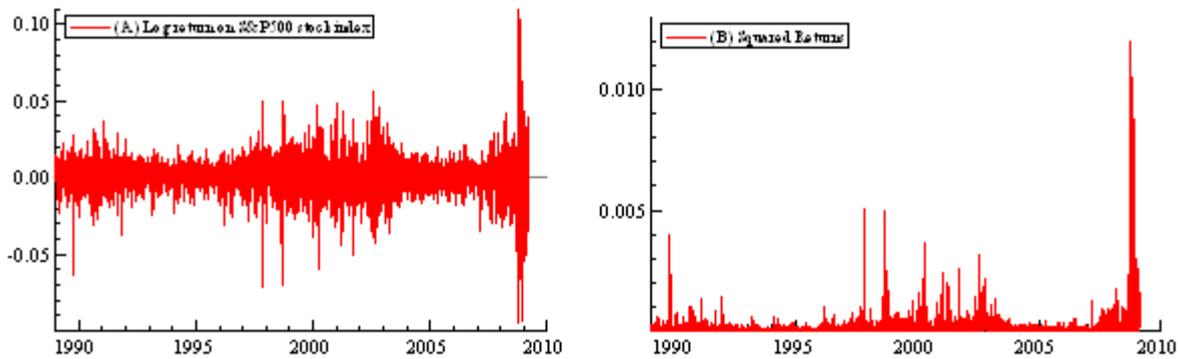
2.2 Informal test for (G)ARCH

The ongoing financial crises reflects high market uncertainty and Figure 1 (A) illustrates this by demonstrating enormous clustering from 2008-2009. Moreover, we see that the periods of large variation entail both positive and negative returns, so whereas the mean seems to be constant, the variance is non-constant. In other words, we observe clear ARCH effects. Figure 1 (B) depicts the squared returns for the S&P 500, and the time dependence is clearly visible in this graph - reflecting the clusters of high variance. It should be emphasized that whereas the stock returns themselves may not be autocorrelated, it is clear that the ARCH effects will make the squared returns highly autocorrelated. An *informal test* for ARCH effects is therefore employed to examine the squared returns and to calculate the autocorrelation function (ACF) for squared returns. Figure 2 (A) depicts

³ The News Impact Curve.

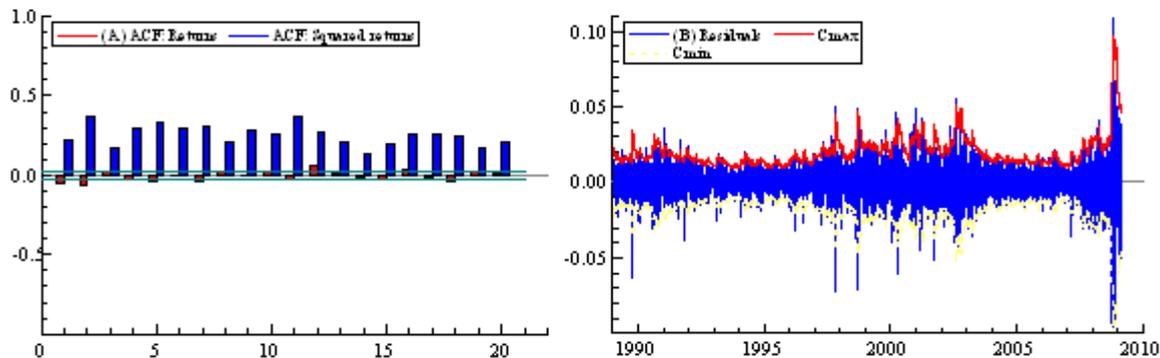
that the ACF is very small for the returns, while it is much larger for the squared returns, indicating ARCH effects. Figure 2 (B) is important, because it illustrates a fundamental principle in ARCH models;

Figure 1 – (A) Log return on S&P 500 and (B) Squared returns



In absence of shocks the conditional variance shrinks, and in the presence of shocks the confidence bands expands, as a large shock tend to be followed by a large shock. In absence of ARCH effects the conditional variance should be constant.

Figure 2 – (A) ACF for returns & ACF for squared returns, (B) Estimated residuals with confidence bands⁴.



2.4 Estimating an autoregressive univariate model for S&P 500

By applying OLS, the log return on S&P 500 stock index is estimated and tested for ARCH. An AR(5) model is estimated in PcGive and Table 2 in appendix 1 indicates that the lagged values are generally insignificant and they can be removed one by one. At last the model will consist of only one constant. The R-squared is very low, $R^2 = 0.0093$, meaning that the statistical model explains less than one percent of the variation in returns. The P-value for the F – statistic is 0 and the error

⁴ Estimated residuals, $\hat{\varepsilon}_t$, and the conditional std. (i.e. 95% confidence bands as $\pm 2\sigma$).

terms are not normally distributed⁵, indicating ARCH effects. Table 4 denotes that the null hypothesis of no autocorrelation in the residuals cannot be rejected, as the (high) p – value shows no evidence against H_0 . This is convenient, because we can only test for ARCH in the residuals if there is no autocorrelation. Autocorrelation means that the residuals depend on previous residuals. Thereby the squared residuals also depend on the lagged squared residuals, indicating ARCH. That is, autocorrelation implies ARCH, but ARCH does not imply autocorrelation. The ARCH test is said to have power against autocorrelation and therefore one always need to test for autocorrelation prior to the ARCH test.

2.5 Testing for ARCH

When testing for ARCH effects up to order p in the data set, one can use a standard Breuch-Pagan LM test for no-heteroskedasticity using the following auxiliary regression model:

$$\hat{\varepsilon}_t^2 = \gamma_0 + \gamma_1 \hat{\varepsilon}_{t-1}^2 + \gamma_2 \hat{\varepsilon}_{t-2}^2 \dots + \gamma_p \hat{\varepsilon}_{t-p}^2 + error, \quad (1)$$

where $\hat{\varepsilon}$ is the estimate residual from the regression of interest, Nielsen (2007). We test simultaneously for no ARCH effects by imposing the null hypothesis of no ARCH:

$$H_0 : \gamma_1 = \gamma_2 = \dots = \gamma_p = 0$$

If the null hypothesis cannot be rejected, the variance of the entire data set is constant (conditional variance = unconditional variance) and there is no evidence of ARCH effect. Table 5 in the appendix denotes that the p-value for the test of no ARCH is 0, which implies that the null hypothesis of no ARCH(7) can be rejected on a 1% significance level. It can also be rejected that all the coefficients corresponding to the lagged variables are equal to zero:

$$\varepsilon_t^2 = \beta_0 + \underset{(0.014)}{0.0146} \varepsilon_{t-1}^2 + \underset{(0.014)}{0.2203} \varepsilon_{t-2}^2 - \underset{(0.014)}{0.0344} \varepsilon_{t-3}^2 + \underset{(0.014)}{0.1168} \varepsilon_{t-4}^2 + \underset{(0.014)}{0.2089} \varepsilon_{t-5}^2 + \underset{(0.014)}{0.1262} \varepsilon_{t-6}^2 + \underset{(0.014)}{0.1193} \varepsilon_{t-7}^2 \quad (2)$$

With standard errors in the parentheses. 7 lags are chosen because it should be sufficient to state some of the weaknesses of the ARCH model; it is difficult to precisely pin down shape of the memory structure and the estimated coefficients are often relatively unstable, i.e. equation (2). Instead of expanding the equation to a more complicated model, we can use a *Generalized Autoregressive Conditional Heteroskedasticity* (GARCH) model to resolve the problems with many lags in ARCH models.

⁵ Table 3 denotes that the return on S&P 500 is not normally distributed, as it displays excess *kurtosis*. $K(x) = 3$ for a normally distribution, TSAY (2005).

3.1 Generalized ARCH (GARCH) models

ARCH models have been generalized in several different ways, Bollerslev (1986) proposed a useful extension known as the GARCH model. The simplest case is the very popular GARCH(1,1) model defined by the equation:

$$\sigma_t^2 = \bar{\omega} + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2. \quad (3)$$

Equation (3) is an equation for the *conditional* variance. The intuition is that if the shock at time $t-1$ is large (measured as a large residual ε_{t-1}) then the variance at time t is high, implying a large probability of a large shock (positive or negative). The ARCH effect comes in via the lagged squared residual. The GARCH model also includes the lagged conditional variance, σ_{t-1}^2 . It can be noted that a GARCH model is a parsimonious (and restricted) representation of an ARCH model with many lags. The GARCH model implies an ARMA structure for the squared residuals, see appendix 2. Note that the equation only has three unknown parameters to estimate, despite this the GARCH(1,1) specification often performs very well, Verbeek (2004). One explanation for this efficiency is that the lagged dependent variable allows for a more smooth development in the variance and a longer memory without including too many parameters, Nielsen (2007).

3.2 Testing for GARCH

Using PcGive a GARCH(1,1) model is estimated by maximum likelihood. The GARCH model does well in catching the ARCH effects and is considerably more convenient to operate with than an ARCH model with many lags. The test results for equation (3) are reported in Table 1. The (G)ARCH effects are clearly significant, with t-ratios for α and β of 6.89 and 119, respectively.

Figure 3 illustrates the interpretation of the GARCH model. When a large shock hits the market, the variance increases and the confidence bands for the following shocks also increase. The current financial crisis is clearly depicted, accelerating in the last quarter of 2008. Recall that the GARCH(1,1) model implies an ARMA(1,1) structure for the squared error terms. Stationarity of the ARMA process requires that the autoregressive root is smaller than unity, i.e. $\alpha + \beta < 1$. Table 1 denotes that the estimated $\hat{\alpha} + \hat{\beta} = 0.994242$. This is still in the stationary region, but it is very close to a unit root⁶. In other words, we almost have a unit root in the ARMA(1,1) process for squared error terms. This implies that the shocks to the squared residuals have persistent effects, and the market will only very slowly revert to a normal state (however, one may argue that a normal state does not exist in stock markets). Figure 2 (A) also illustrates this fact, as the autocorrelation

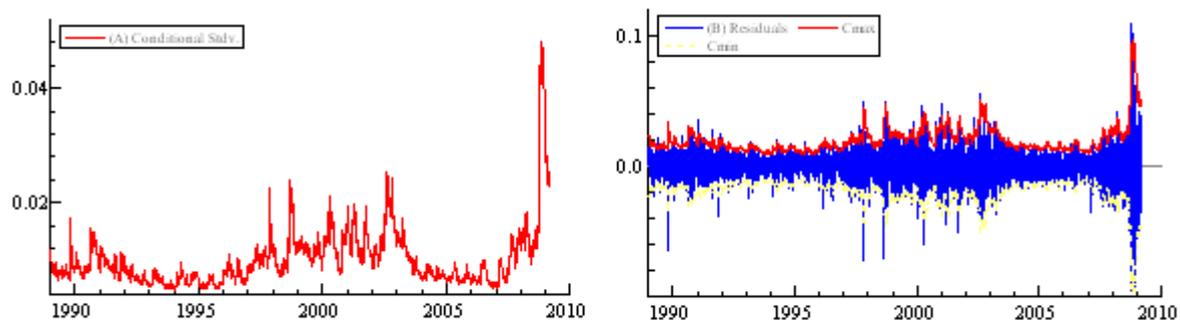
⁶ This is also referred to as a random walk.

function shows high persistence for the squared returns. Figure 3 (A) depicts the estimated conditional standard deviation⁷ of the S&P 500 stock index, which enables us to see how the stock market is affected by financial turmoil. Figure 3 illustrates that the periods in which the conditional standard deviation increase is the same periods for which the confidence bands are notable expanded and where the residuals are high. Again underlining the fundamental principal in (G)ARCH models, that extreme market movements tend to be followed extreme movements. That is exactly what investors have witnessed in the current turmoil on the financial markets.

Table 1 – Modeling log return on S&P 500 stock index by GARCH(1,1).

| | Coefficient | Std.Error | Robust-SE | t-value | t-prob |
|---------------------|-------------|-----------|--------------------|---------|--------------------|
| Constant | 0.000464351 | 0.0001132 | 0.0001167 | 3.98 | 0.000 |
| ω | 0.0757976 | 0.0001482 | 0.0002304 | 3.29 | 0.001 |
| α | 0.0575055 | 0.005445 | 0.008345 | 6.89 | 0.000 |
| β | 0.936737 | 0.005648 | 0.007899 | 119 | 0.000 |
| log-likelihood | 16.595,4206 | | | | |
| mean(h_t) | 0.00012845 | | var(h_t) | | 4,32E-03 |
| no. of observations | 5080 | | no. of parameters | | 4 |
| alpha(1)+beta(1) | 0.994242 | | alpha_i+beta_i>=0, | | alpha(1)+beta(1)<1 |

Figure 3 – (A) Conditional Standard deviation on the S&P 500 and (B) Estimated residuals with confidence bands.



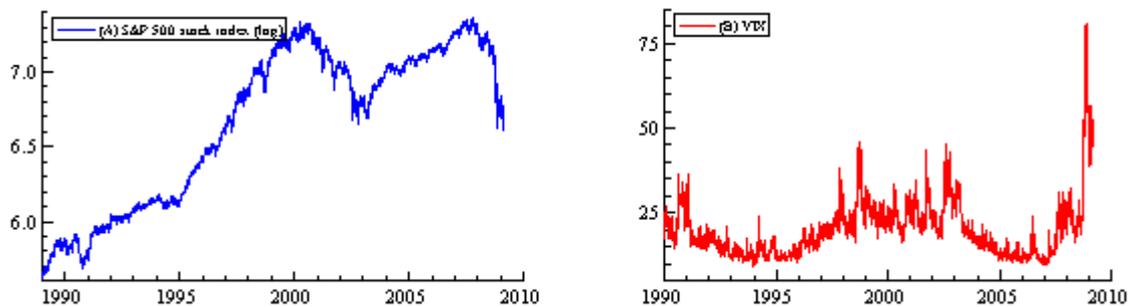
4.1 The Volatility Index

Chicago Board Options Exchange (CBOE) introduced the volatility index (VIX) in 1993, and it quickly became the benchmark for stock market volatility. The VIX is widely followed by financial practitioners and is extensively used in the financial literature, CBOE (2003). The VIX measures

⁷ The conditional standard deviation of the S&P 500 stock index can be seen as a measure of uncertainty in the stock market and hereby also the systematic risk that the investors are exposed to.

market expectations of the volatility in the S&P 500 index over the next 30 days, which is the *implied volatility* in the market. The VIX is derived⁸ from bid/ask quotes of options on the S&P 500 index rather than solving it out of an option pricing formula. The VIX has the convenient feature that, $VIX = \sigma * 100$ making it very easy to interpret.

Figure 4 - (A) S&P 500 Stock index (log) and (B) The VIX, 1990-2009.



A VIX value of 50 means, that the standard deviation in the market is approximately 50 pct. We can think of options as representing a form of insurance and the more expensive the insurance is the more fear in the market, all else equal. The idea is that the options trading are used as a proxy on how the actual trade in the underlying asset is going to be like. Volatility is typically an indication of financial turmoil, and therefore the VIX is often referred to as the “*investor fear gauge*”, as the VIX increases in periods of financial distress, CBOE (2003). For this reason, the VIX also reflects investor sentiment and risk aversion. Figure 4 depicts that the VIX tends to rise, when the S&P 500 declines and vice versa. One explanation for this might be that when the S&P 500 declines, investors are generally more eager to buy out-of-the-money put options⁹. This will drive up implied volatility, and thus the value of these options. This increase in implied volatility is reflected in an increase in the level of the VIX. Arbitrageurs who sell puts must sell stocks in order to hedge their position and remain delta neutral¹⁰, as the prices of puts are driven up by investors who are seeking to hedge downside risk, Siegel (2008). Effectively, this may deepen the decline of S&P 500 and potentially increase the magnitude of the negative correlation between the VIX and S&P 500. The inverse relationship between S&P 500 and the VIX is more evident in Figure 6. Another explanation for the inverse correlation is that volatility has been more excessive in bear markets than in bull markets. Investor expectations based on historical volatility will therefore cause implied volatilities to exhibit similar behavior, Sloyer and Tolking (2008). In our sample period, the VIX

⁸ See appendix 2 for the formula.

⁹ Out-of-the-money put option: A put option where the asset price is greater than the strike price, Hull (2008).

¹⁰ A portfolio with a delta of zero so that there is no sensitivity to small changes in the price of the underlying asset, Hull (2007).

has been all time high during the current financial crisis and thereby emphasizing how extremely volatile the financial markets has been since the last quarter of 2008.

5.1 Volatility forecasting in practice

An important application of (G)ARCH models is prediction of future volatility. By applying the theory of volatility forecasting¹¹ a time series from the returns on the S&P 500 has been constructed. By programming¹² in the OxRun module in OxMetrics a time series consisting of 6.485 forecasts that continuously predicts the variance of the stock returns on S&P 500 from 30.1.1991-30.1.2009 is constructed. By running the forecast code in OxMetrics one will obtain thousands of forecasts on 1-30 days horizon, but only the estimate that predicts the variance 30 days ahead has been kept, in order to make it comparable with the VIX¹³. One can think of this time series as being backward looking, because the forecasted values of the conditional variance will always converge towards the unconditional variance. Therefore the forecasts will not provide an investor with the opportunity to directly pin out future events, unless they have already been counted for in the stock prices. Nevertheless, (G)ARCH models do provide relatively good forecasts, see Andersen and Bollerslev (1997). To achieve a measure that can be compared to the VIX, which indicates the standard deviation in the market, the variance forecasts are turned into standard deviations. Effectively, the forecasts and the VIX are moved forward 30 days, so that the prediction made at i.e. the 1. November 2008 will tell us how the volatility will look like on the 1. December 2008.

5.2 Comparing the GARCH forecasts, estimated volatility and the VIX

Figure 5 (A) illustrates that the trend in the forecasts, estimated volatility and the VIX follow each other very closely. The time series are highly correlated, indicating that all three measures seem to respond on the same market movements. The GARCH(1,1) forecast appears to fit the estimated volatility extremely well. The VIX also seems to be a good indicator of the future volatility in the stock market. Furthermore, Figure 5 (A) illustrates that from the beginning of 2004 the markets just recovered from a period with high volatility. This explains why the GARCH(1,1) forecast slightly overshoots the estimated standard deviation (2004-2005), as the forecast has some memory. At the end of 2008, where the financial markets were extremely nervous, the estimated variance increased dramatically. However, this change in the stock market volatility is observed with some delay by

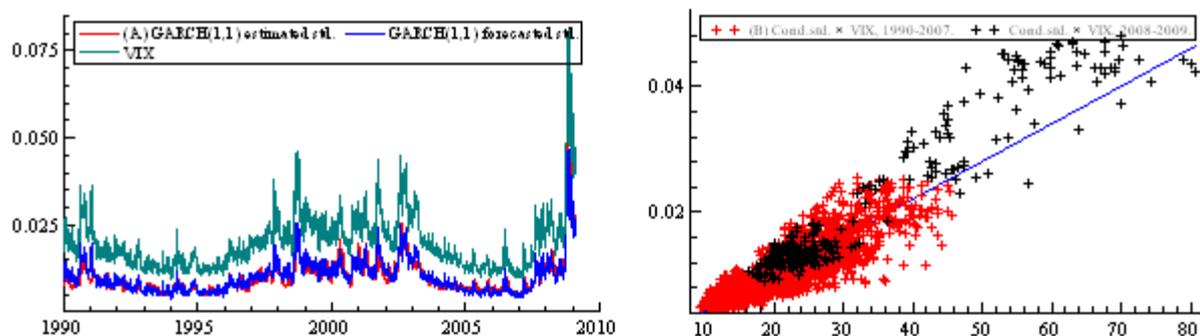
¹¹ See appendix 3.

¹² One can contact me personally for the programming code and the comprehensive forecast output.

¹³ Recall that the VIX is based on options with 30 days to expiry.

the forecasts, as they are backward looking. It is notable that when the market is hit by a shock and the volatility increases dramatically, some persistence occurs in the forecasts while the VIX seems to react more rapidly and to a larger extent. Indicating that the VIX might be a better measure for the *day trader* who always needs to be up to date with the changes in the market. However, it should be emphasized that even though the VIX more rapidly incorporates an increase in the risk level on the stock markets, Figure 5 (A) also illustrates that the VIX tends to overestimate the market risk¹⁴. This might be due to the risk premium built into equity options and that options empirically are more sensitive than stocks, Hull (2008). Another reason might be that the investors are too bearish on the market during periods of financial turmoil. The simple reason why implied volatility and actual volatility differs is that market participants do poorly in forecasting actual volatility. Figure 5 (B) is a cross plot of the estimated standard deviation against the VIX from 1990-2009. The black crosses explain the dependence between the two measures from 2008-2009, and thereby focuses on the financial crisis. Also indicating that the VIX tends to overestimate market risk (volatility) in periods of financial distress, as several black crosses are plotted North West in the graph and above the regression line.

Figure 5 – (A) Estimated std., Forecasted std., and the VIX, (B) Conditional std. vs. VIX 1990-2009



6.1 Trading strategies and the VIX

There are several reasons why investors should care about the VIX; this paper will only touch briefly upon some of them in order to give the investor an idea of how to take advantage of this fear gauge. Many investors follow the VIX because it provides important information about investor sentiment, which can be helpful in evaluating potential market turning points. More specialized investors use VIX options and futures to speculate on the future direction of the market¹⁵. The inverse relationship between the VIX and the S&P 500 (i.e. Figure 6) implies that when the VIX

¹⁴ CBOE's analysts also find that the VIX tends to overestimate the market risk, CBOE (2009).

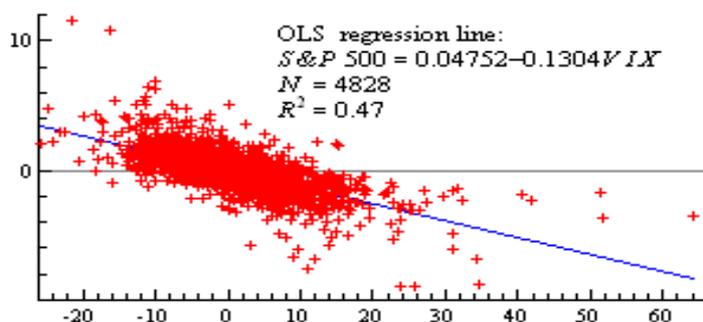
¹⁵ See Ahoniemi (2006) for an interesting discussion.

increases, investors who are long S&P 500 tend to lose money. There is an old adage saying that, when the VIX is high it is time to buy and when the VIX is low it is time to go, CBOE (2009). The investor should not build his trading platform on this alone, but it holds true to some degree. The VIX is a *mean reverting* index, meaning that if it goes very low or high the probability of it going back towards the average is increased, CBOE (2009).

6.2 Portfolio hedging using VIX options

Exchange-traded VIX options give investors the ability to trade market volatility. Trading VIX options can be a useful tool for investors wanting to hedge their portfolios against sudden market declines. Figure 6 illustrates that the S&P 500 and the VIX are strongly inverse correlated, making the VIX very appropriate as a hedging tool. When the VIX is low, the inverse correlation of the highly volatile VIX to the S&P 500 enables the investor to use VIX options as a hedge to protect a portfolio against a market crash. The investor may implement such a hedge by purchasing near-term slightly out-of-the-money VIX calls¹⁶. Simultaneously the investor sells slightly out-of-the-money VIX puts¹⁷ of the same expiration month. This strategy is known as the *reverse collar*, CBOE (2009). In the event of a stock market crash, as the one we experienced in November 2008 it is liable that the VIX will rise high enough so that the VIX call options gain sufficient value to offset the losses in the portfolio.

Figure 6 – Daily changes in S&P 500 vs. Daily changes in the VIX



That is the basic idea of the reverse collar, and if the investor can execute this strategy he may effectively protect the value of his portfolio against market crashes as the one we have witnessed during the current crisis. If the investor uses VIX options to hedge a portfolio, he needs to ensure that the portfolio is highly correlated to the S&P 500 index, that is, a beta close to one. The

¹⁶ Out-of-the-money call option: A call option where the asset price is less than the strike price, Hull (2008).

¹⁷ Out-of-the-money put option: A put option where the asset price is greater than the strike price, Hull (2008).

challenge for the investor is to determine the amount of VIX calls he needs to purchase to protect his portfolio. Several simplified examples can be found on CBOE's website and other trading sites.

6. 3 The benefits of using VIX options as a hedging tool

The VIX offers investors the opportunity to trade equity volatility. Due to the fact that S&P 500 is the underlying index, the VIX can be employed as a hedging vehicle on the most liquid stock index in the financial markets. Investors can hedge their portfolios using several other types of options, but these options are not pure trades on volatility. Generally option positions have exposure to market direction in addition to volatility. Therefore, option traders must spend considerable resources on delta hedging¹⁸ their positions if they wish to minimize directional risk. The VIX on the contrary, provide investors with the ability to take a position on implied volatility independent of the level and directions of stock prices. Investors are often long S&P 500 and thereby implicitly short volatility due to the inverse correlation between equities and equity volatility. The VIX therefore provides a powerful diversification tool. One unique feature of the VIX is that investors can reduce their market risk due to volatility by including volatility (VIX) in a diversified portfolio. Several empirical studies support this claim, concluding that adding VIX options to an S&P 500 portfolio reduces risk without significantly affecting return, Moran and Dash (2007). Sloyer and Tolking (2008) finds that the inclusion of the VIX to S&P 500 portfolios improves the risk-return profile of respective portfolios. Furthermore, they find that the risk-adjusted performance (Sharpe Ratio) of the VIX enhanced portfolios are more profitable than non-VIX equity portfolios. However, it should be emphasized that the high cost of holding a long volatility position in a low volatility market may affect portfolio returns negatively.

7. Conclusion

The purpose of this paper has been to model and forecast volatility in the financial markets and analyze ways in which investors can use VIX options to hedge their portfolios during periods of financial turmoil. The graphical analysis revealed large autocorrelations and high persistence for the squared returns, indicating ARCH effects. The estimated residuals with confidence bands illustrated a fundamental principle in (G)ARCH models, namely that large market movements tend to be followed by large market movements, of either sign, and vice versa. The applied econometric method has been (G)ARCH modeling and the empirical findings indicated that the squared returns

¹⁸ A hedging scheme that is designed to make the price of a portfolio of derivatives insensitive to small changes in the price of the underlying asset, Hull (2007).

for S&P 500 clearly contains (G)ARCH effects and volatility clustering. The paper found that the ARCH model does poorly in precisely pinning down the shape of the memory structure and therefore has a problem with too many lags. For this reason, I expanded the analysis to a GARCH(1,1) model. The GARCH(1,1) test found that the (G)ARCH effects in the S&P 500 return were clearly significant with t-ratios for α and β of 6.89 and 119, respectively. The test also showed that we almost have a unit root in the ARMA(1,1) process for squared error terms; implying that shocks to the squared residuals have persistent effects and the financial markets will only slowly revert to a center of gravity. In accordance with Andersen and Bollerslev (1997), this paper finds that the GARCH(1,1) model provides good volatility forecasts and it fits data very well. The forecasts were easily applied by programming in OxMetrics. The VIX seems to be a rather robust measure of future volatility in the market and a precise measure of investor sentiment. However, it tends to overestimate the actual volatility in the market. Presumably this can partly be explained by the risk premium built into equity options. Furthermore, the paper found that investors can use VIX options as a hedging tool because it has a strong negative correlation to the S&P 500. This inverse correlation can partly be explained by the increase in demand for out-of-the-money put options, when S&P 500 declines. This drives up implied volatility and thus the value of these options. Effectively, this increase in implied volatility is reflected in an increase in the level of the VIX. VIX options have several beneficial hedging features, because they are pure trades on volatility and the underlying is the highly liquid S&P 500 index. Furthermore, the VIX provides investors with the ability to take a position on implied volatility independent of the level and directions of stock prices. In conclusion, investors can gain valuable insights on financial volatility by applying the methods and instruments covered throughout this paper.

8. References

- AHONIEMI, K. (2006): *Modeling and Forecasting Implied Volatility – an Econometric Analysis of the VIX Index*. Discussion paper No. 129, Helsinki School of Economics, FDPE, and HECER.
- ANDERSEN, T. G. & BOLLERSLEV, T. (1997): *Answering the Critics: Yes, ARCH Models Do Provide Good Volatility Forecasts*. Working paper 6023, NBER.
- BOLLERSLEV, T. (1986): *Generalized Autoregressive Conditional Heteroskedasticity*. *Journal of Econometrics*, 31(3), p. 307-327
- CBOE, (2003): *VIX – CBOE Volatility Index*. Chicago Board Option Exchange.
- CBOE, (2009): CBOE-TV Spotlight at <http://www.cboe.com/>.
- ENGLE, R. F. (1982): *Autoregressive Conditional Heteroskedasticity with Estimates of the variance of United Kingdom Inflation*. *Econometrica*, 50(4), 987-1008.
- GLOSTEN, L, et al (1993): *Relationship Between the Expected Value and the Volatility of the Nominal Excess Return of stocks*. *Journal of Finance*, 48(5), 1779-1802.
- HULL, J. C. (2008): *Option, futures and other derivatives*. 7th edition, Prentice Hall.
- HULL, J. C. (2007): *Risk Management and Financial Institutions*. Prentice Hall.
- MANDELBROT, B (1963): *The Variation of Certain Speculative Prices*. *The Journal of Business*, 36(4), 394-419.
- MORAN, MT., and DASH, S. (2007): *VIX Futures and Options: Pricing and Using Volatility Products to Manage Downside Risk and Improve Efficiency in Equity Portfolios*. *The Journal of Trading*: 96-105.
- NELSON, D. B. (1990): *Stationarity and Persistence in the GARCH(1,1) Model*. *Econometric Theory*, 6(3), 318-334.
- NELSON, D.B, (1991): *Conditional Heteroskedasticity in Asset Returns: A New Approach*. *Econometrica*, 59, 347-370.

- NIELSEN, H. B. (2007): *Modeling Volatility in Financial Time Series: An Introduction to ARCH*. University of Copenhagen, Lecture Note.
- SIEGEL, J. J. (2008): *Stocks for the Long Run*. 4th edition, New York: McGraw-Hill.
- SLOYER, M. and TOLKIN, R.: *The VIX as a Fix: Equity Volatility as a Lifelong Investment Enhancer*. Honors thesis, Duke University, Durham, North Carolina.
- TAYLOR, S. J. (1986): *Modeling Financial Time Series*. Wiley, New York.
- TSAY, R. S.(2005): *Analysis of Financial Time Series*. 2nd edition, Wiley-Interscience.
- VERBEEK, M. (2004): *A Guide to Modern Econometrics*. 2nd edition, John Wiley & Sons Ltd.
- WOOLDRIDGE, J. M. (2006): *Introductory Econometrics – A Modern Approach*. 3rd edition, South-Western College Pub.

Appendix 1

Table 2 – Modeling S&P 500 stock index by OLS.

| | Coefficient | Std. Error | t-value | t-prob | Part.R ² |
|---------------------|------------------------|--------------------|----------------------|--------|---------------------|
| Return_1 | -0.0561806 | 0.01404 | -4.00 | 0.0001 | 0.0031 |
| Return_2 | -0.0657559 | 0.01406 | -4.68 | 0.0000 | 0.0043 |
| Return_3 | 0.0147422 | 0.01410 | 1.05 | 0.2958 | 0.0002 |
| Return_4 | -0.0273965 | 0.01408 | -1.95 | 0.0518 | 0.0007 |
| Return_5 | -0.0345660 | 0.01407 | -2.46 | 0.0140 | 0.0012 |
| Constant | 0.000228555 | 0.0001593 | 1.43 | 0.1515 | 0.0004 |
| R ² | 0.00931166 | F(5,5069) | 9.52889 [0.0000] *** | | |
| Log-Likelihood | 155.346.385 | DW | 2.11 | | |
| No. of observations | 5075 | No. of parameters | 6 | | |
| Mean(Return) | 0.00019462 | Var(Return) | 0.000129666 | | |
| AR 1-2 test: | F(2,5077) = | 18.246 [0.0000]*** | | | |
| ARCH 1-1 test: | F(1,5077) = | 257.00 [0.0000]*** | | | |
| Normality test: | Chi ² (2) = | 4925.8 [0.0000]*** | | | |

***: Significantly different from 0 on a 1% significance level.

Autoregressive equation for the conditional mean:

$$AR(5): y_t = \delta + \theta_1 y_{t-1} + \theta_2 y_{t-2} + \theta_3 y_{t-3} + \theta_4 y_{t-4} + \theta_5 y_{t-5} + \varepsilon_t, \quad (4)$$

where ε_t is a white noise process.

Table 3 – Normality test for residuals.

| | |
|--------------------|-----------|
| Observations | 5075 |
| Mean | 0.00000 |
| Standard deviation | 0.011334 |
| Skewness | -0.46201 |
| Excess Kurtosis | 9.893 |
| Minimum | -0.092524 |
| Maximum | 0.10067 |

Table 4 – Test for no autocorrelation in the residuals from lags 1 to 7.

| | | |
|----------------|-------------|------------------|
| ARCH 1-7 test: | F(7,5062) = | 0.30299 [0.9528] |
|----------------|-------------|------------------|

Table 5 – Test for no ARCH from lags 1 to 7.

| ARCH coefficients: | | |
|--------------------|-------------|--------------------|
| Lag | Coefficient | Std.Error |
| 1 | 0.014674 | 0.01397 |
| 2 | 0.22029 | 0.01386 |
| 3 | -0.034432 | 0.0139 |
| 4 | 0.1168 | 0.01383 |
| 5 | 0.20886 | 0.01392 |
| 6 | 0.12622 | 0.01388 |
| 7 | 0.11925 | 0.01399 |
| ARCH 1-7 test: | F(7,5055) = | 256.32 [0.0000]*** |

***: Significantly different from 0 on a 1% significance level.

The ARCH(7) model for the conditional variance: $\sigma_t^2 = \bar{\omega} + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 \varepsilon_{t-2}^2 + \dots + \alpha_7 \varepsilon_{t-7}^2. \quad (5)$

To ensure a consistent model that generates a positive variance, one need to constrain the parameters, $\bar{\omega} > 0, \alpha_i \geq 0$.

Appendix 2

The GARCH(1,1) model implies that the squared residuals follow an ARMA(1,1) process

The GARCH model implies an ARMA structure for the squared residuals. Define the surprise innovation¹⁹ as $v_t = \varepsilon_t^2 - \sigma_t^2$. Inserting this in equation (3) yields:

$$\begin{aligned}\sigma_t^2 &= \bar{\omega} + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \\ \varepsilon_t^2 - v_t &= \bar{\omega} + \alpha\varepsilon_{t-1}^2 + \beta(\varepsilon_{t-1}^2 - v_{t-1}) \\ \varepsilon_t^2 &= \bar{\omega} + (\alpha + \beta)\varepsilon_{t-1}^2 + v_t - \beta v_{t-1},\end{aligned}\tag{6}$$

which is an ARMA (1,1) process for the squared residuals, Tsay (2005). Stationarity of the ARMA process requires that the autoregressive root is smaller than unity, i.e. $\alpha + \beta < 1$. Provided that this is fulfilled one can write the *unconditional* variance as:

$$\sigma^2 = E[\varepsilon_t^2] = \frac{\bar{\omega}}{1 - \alpha - \beta}.\tag{7}$$

The generalized VIX formula

The VIX is calculated using the two nearest expiration months of the S&P 500 options in order to achieve a 30-calender-day period. The value of the VIX index is derived from the prices of at-the-money and out-of-the money calls and puts. The closer the strike price to the at-the-money value, the higher the weight its price receives in the calculation. The generalized formula for calculating the VIX index is:

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} * e^{RT} Q(K_i) - \frac{1}{T} \left[\frac{F}{K_0} - 1 \right]^2,\tag{8}$$

where σ is the value of the VIX divided by 100, T is the time to expiration of the option contract, F is the forward index level derived from option prices, K_i is the strike price of the i^{th} out-of-the-money option (call if $K_i > F$ and put if $K_i < F$), ΔK_i is the interval between strike prices, or $(K_{i+1} - K_{i-1})/2$, K_0 is the first strike below F , r is the risk-free interest rate up to expiration of the option contract, and $Q(K_i)$ is the midpoint of the bid-ask spread for each option with strike K_i . Calls and puts are included up to the point where there are two consecutive strike prices with a bid price equal to zero, CBOE (2003) and Ahoniemi (2006).

¹⁹ V_t denotes the error in the forecast. The error arises, because of chocks to the variance and these chocks are white noise processes (mean zero and uncorrelated over time – however, heteroskedastic). For this reason, V_t becomes a white noise process, which means that a GARCH(1,1) model for the conditional variance can be rewritten as a ARMA(1,1) model for the squared residuals.

Appendix 3

Theory of volatility forecasts

This paper uses $\sigma_{T+h|T}^2 = E[\varepsilon_{T+h}^2 | I_T]$ to describe the volatility forecast at time $T+h$ given the information set at time T . By using the unconditional variance from the GARCH(1,1) model and then rewriting we obtain:

$$\sigma^2 = \frac{\bar{\omega}}{1 - \alpha - \beta} \Leftrightarrow \bar{\omega} = \sigma^2(1 - \alpha - \beta). \quad (9)$$

By substituting this into the variance equation for the GARCH(1,1) model one can find an expression for the deviations from the unconditional variance

$$\begin{aligned} \sigma_t^2 &= \bar{\omega} + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \Leftrightarrow \\ \sigma_t^2 &= \sigma^2(1 - \alpha - \beta) + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \Leftrightarrow \\ \sigma_t^2 - \sigma^2 &= \alpha(\varepsilon_{t-1}^2 - \sigma^2) + \beta(\sigma_{t-1}^2 - \sigma^2) \Leftrightarrow \\ \sigma_t^2 &= \sigma^2 + \alpha(\varepsilon_{t-1}^2 - \sigma^2) + \beta(\sigma_{t-1}^2 - \sigma^2). \end{aligned} \quad (10)$$

To forecast volatility in the financial markets one period ahead, $T+1$, we use the insights obtained from the above recursion, so the best prediction for the variance one period ahead have to follow the process outlined above. Doing so the first period forecast will be given by:

$$\begin{aligned} \sigma_{T+1|T}^2 &= E[\varepsilon_{T+1}^2 | I_T] \Leftrightarrow \\ \sigma_{T+1|T}^2 &= E[\sigma^2 + \alpha(\varepsilon_T^2 - \sigma^2) + \beta(\sigma_T^2 - \sigma^2) | I_T] \Leftrightarrow \\ \sigma_{T+1|T}^2 &= \sigma^2 + \alpha(\varepsilon_T^2 - \sigma^2) + \beta(\sigma_T^2 - \sigma^2), \end{aligned} \quad (11)$$

due to the fact that ε_T is in the information set at time T , i.e. $E[\varepsilon_{T+1}^2 | I_T] = \varepsilon_T^2$, and σ_T^2 can be derived from the information set, Nielsen (2007). This tells us that the best prediction for the unconditional variance given the information set, is the expression for the deviations from the unconditional variance.

One can proceed this way and forecast up to $T+h$, for which the forecast will look this way:

$$\sigma_{T+2|H}^2 = \sigma^2 + (\alpha + \beta)(\sigma_{T+h-1|T}^2 - \sigma^2). \quad (12)$$

Note that the forecasts will converge exponentially with speed $\alpha + \beta$ towards the unconditional variance, σ^2 , Nielsen (2007).