



LUNDS UNIVERSITET
Ekonomihögskolan

Department of Economics

Bachelor Thesis

January 2015

Volatility Forecasting In the Nordic Stock Market

Niklas Hummel

Supervisor

Professor Hans Byström

Abstract

This paper studies volatility prediction on OMX Stockholm 30, OMX Helsinki 25 and OMX Nordic 40. The models used are a historical variance model, an exponentially weighted moving average model and three models from the GARCH family. These are GARCH(1,1), EGARCH(1,1) and GJR(1,1), with normal and t-distribution respectively. The volatility for 2008-2013 is forecasted with a rolling window technique using historical data from 2002-2013. The models are ranked based on forecasting accuracy. The difference in accuracy is then translated into an average volatility forecasting error reduction. The financial crisis at the end of 2008 is studied separately, again ranking the models and comparing their relative forecasting ability. For the entire period I find that EGARCH and GJR are the superior models, followed by GARCH and EWMA, with the historical variance model yielding the greatest loss. Student's t-distribution compared to normal distribution yields varying results for the entire period. For the crisis, GJR is the best model in all of the markets, while standard GARCH is the worst performing of all the models. t-distribution yields a substantial forecasting improvement over normal distribution for the crisis. The relative volatility forecasting error between the best and worst of the GARCH models are several times larger during the crisis than during the entire period for Nordic and Stockholm.

Table of Contents

1. Introduction	5
1.1 Purpose	5
1.2 Delimitations and previous research	5
1.3 Disposition	6
2. Theory	7
2.1 Models for volatility forecasting.....	7
2.1.1 The historical variance	8
2.1.2 EWMA	8
2.1.3 GARCH	9
2.1.3.1 <i>GARCH(1,1)</i>	10
2.1.3.2 <i>GJR(1,1)</i>	11
2.1.3.3 <i>EGARCH(1,1)</i>	11
2.2 Measuring the loss-The QL loss function	12
2.3 Translating the loss into volatility forecasting error.....	13
2.4 From volatility forecasting error to pricing error.....	14
2.4.1 Option pricing	14
2.4.2 Value-at-Risk	15
3. Data & Method	16
3.1 Software.....	16
3.2 Data used and method of forecasting	16
4. Empirical Results	20
4.1 Forecasting results for the entire period	20
4.2 Forecasting results for the crisis	23
5. Conclusions	28
6. References	29

1. Introduction

The concept of randomness certainly constitutes a large part of finance. Even if one knew the price of all assets and all economic variables at a certain time one could not exactly say what they would be the day after. What can be done is to try and find a model that captures the essence of this randomness and allows an investor to make predictions about, for instance, the future price of an asset. The general idea for an investor is to maximize the returns while minimizing the risk. The easiest perceivable quantification of risk is the volatility or variance of an asset. For this reason the concept of volatility forecasting has been studied extensively. Not only is the volatility of an asset important for the price of that particular asset, the volatility is also used to price derivatives. In option pricing for example, the estimate of future volatility is used both for pricing, valuation and to derive the hedge ratios. The financial crisis in 2008 led to increasing volatilities across the board and with the added regulations of Basel III even more pressure are put on the institutions. These dictate the use of more rigorous risk calculations and introduce a risk of being grounded (Basel III, 2010). For these reasons, volatility forecasting may be more important than ever.

1.1 Purpose

The aim of this paper is to evaluate the volatility forecasting ability of five different models on several Nordic indices for the period 2008-2013. The loss of the different models and specifications will be analyzed and the models will be ranked. The results for the entire period will be compared to the loss of the different models during the 2008 financial crisis. To relate the statistical loss value to economic terms the statistical loss will be translated into an average error reduction in forecasting, when switching from the worst to the best fitting model for the different periods. By using a calibrated example I want to demonstrate how the difference in forecasting ability translate to relative errors in Value-at-Risk calculations and option pricing.

1.2. Delimitation and previous research

A great number of models have been developed to be used in volatility forecasting. Often they are designed to capture a certain observed feature of the asset studied, or a feature believed to exist. Models generated this way can usually be found to “fit” a certain set of data very well. However, the forecasting ability of these models may often be worse than vanilla models (Hansen & Lunde, 2001). For this reason this paper will focus on out-of sample forecasting and not the fit of the models to in sample data.

As stated in the introduction volatility forecasting has been studied extensively. Papers such as Hansen & Lunde (2001) perform studies on a wide variety of models. Several models, several loss functions and different forecast horizons are used. The results from different papers and for different markets and time periods are mixed, with different models yielding different losses for different markets. This paper focuses on the 1 day ahead volatility forecasting, utilizing only one loss function, forecasting the period of 2008-2013 with a rolling window technique.

1.3. Disposition

Chapter 2 outlines the theoretical framework relevant to the paper. Financial returns are defined and the volatility process for the different models are discussed. The QL-loss function is introduced and the statistical loss is related to an average volatility forecasting error. An example of Value-at-Risk calculation and option pricing is introduced to relate volatility forecasting errors to strictly financial applications. Chapter 3 shows the characteristics of the data used. Here the method and software used are also discussed. Chapter 4 gives the empirical results of the volatility forecasting. The QL-loss from the different models and markets are reported and discussed. Relative volatility forecasting error reductions are calculated. Chapter 5 hold the conclusions, comments and suggestions for further research.

2. Theory

Certain features of financial markets have been well documented and should be incorporated in some form into a “good” model. These include volatility clustering, the fact that one period of high volatility makes the next period more likely to exhibit high volatility. Asymmetry, the fact that “bad” news has a higher impact on volatility than good news. Mean reversion, the price of an asset tending to move towards its average over time, is also observed in most markets (Poon & Granger, 2003). The different volatility models used in this paper range from basic, where none of these features are taken into account, to more sophisticated ones, that should in theory be able to capture these features to some extent, as will be explained when the models are introduced.

2.1. Models for volatility forecasting

The volatility of a financial asset is a statistical measure of the variation in the price of that asset over a certain period. The higher the volatility the greater the risk. Volatility and variance (squared volatility) will be used interchangeably in this paper. In the following chapters the different models that will be used in the forecasting are defined and discussed. Before introducing the models I define the logarithmic returns of the stock market, given as input to the different models.

From the historical inter-day prices the logarithmic returns are calculated according to

$$R_t = \ln\left(\frac{P_t}{P_{t-m}}\right) \quad (1)$$

Where P_t is the price of the asset today and P_{t-m} is the price of the asset m days back. This paper will focus on daily returns so m is set to 1.

2.1.1. The Historical Variance

A simple forecast of volatility is using the historical variance. The variance of returns for some days in the past can be calculated according to the formula

$$VAR = \frac{1}{n} \sum_{i=0}^{n-1} r_{t-i}^2 \quad (2)$$

Where

$$r_t = R_t - \bar{R} \quad (3)$$

All the historical observations receive equal weight. A question that arises is how many days of historical variance that should be taken into account. I do the forecasting with the variance calculated over a different number of days as reported in table 4.

The forecasted variance is then simply set to the historical variance.

$$\sigma_{t+1}^2 = VAR(t) \quad (4)$$

2.1.2 Exponentially Weighted Moving Average

A slightly more sophisticated method is the Exponentially Weighted Moving Average (EWMA). In this model the variance for the next day is forecasted as

$$\sigma_{t+1}^2 = \frac{1}{\Gamma} \sum_{i=0}^I B^i r_{t-i}^2 \quad (5)$$

Where

$$\Gamma = \sum_{i=0}^I B^i \quad (6)$$

Here the historical observations receive an exponentially declining weight B . This is perhaps intuitively an improved model, since an event yesterday is presumed to have greater impact on the

volatility tomorrow than an event two weeks ago. Here again arises the question of what weight should be used for the forecasting. Riskmetrics uses this model and sets its B to 0.94 for its Value-at-Risk calculations (Ederington & Guan, 2005). For this reason 0.94 will be one of the values evaluated.

2.1.3 Generalized Autoregressive Conditional Heteroskedasticity

The next method for forecasting variance is three different models from the GARCH family. The GARCH concept was originally developed by Robert Engle in 1982, who used his ARCH model for modeling the inflation rate of the United Kingdom (Engle, 1982). This concept was then “generalized” to a GARCH model by Bollerslev (1986). What makes this model appealing to finance is that it should be able to capture the volatility clustering observed for financial time series. Volatility clustering means that if the volatility has been high over the last few days the next day is likely to also exhibit high volatility. Or put in other words, the returns of an index are characterized by periods of high volatility followed by periods of calm.

In GARCH models it is not the historical standard deviation that is used in the forecasting. Instead a conditional variance is inferred from the historical data via maximum likelihood estimation or another procedure. In other words, the old in-sample variations are used to specify the parameters of the model, depending on the historical observed returns and the conditional variances of these. The models of the GARCH family has been specified in very many ways, tweaked to try and capture certain properties of the time series to be studied. I have chosen to work with three often studied, relatively basic models, namely GARCH(1,1) EGARCH(1,1) and GJR(1,1).

For the GARCH models the return for a certain period is presumed to be able to model according to

$$r_t = \sigma_t z_t \tag{7}$$

The returns depend on the volatility as well as z_t , an independent and identically distributed (iid) process. These numbers are the innovations, the randomness that influence the returns over a certain period. Two different specifications for this process is studied. The first is the normal distribution. Due to the fact that financial time series often are showed to have fatter tails than a normal distribution (Richardson & Smith, 1993, for example), Student’s t-distribution will also be studied. The way that these numbers are distributed is shown in figure 1. The t-distribution has a more well defined peak and fatter tails. The t-distribution varies with the degrees of freedom and as this measure increases the t-distribution converges to a normal distribution. For the model

estimation the degrees of freedom will not specified, instead the best suited value will be evaluated by the estimation software.




Figure 1 Normal distribution (dotted line) and t-distribution with 5 degrees of freedom (solid line).

2.1.3.1 GARCH(1,1)

Bollerslev's original GARCH(p,q) has the following specification for the conditional variance.

$$\sigma_{t+1} = \alpha_0 + \sum_{i=1}^q \alpha_1 r_t^2 + \sum_{i=1}^p \alpha_2 \sigma_t \quad (8)$$

Many studies show that the number of lags, p and q, can be set to one without much loss of forecast accuracy (Hansen & Lunde, 2001, for example). In that case the GARCH model can be simplified to

$$\sigma_{t+1} = \alpha_0 + \alpha_1 r_t^2 + \alpha_2 \sigma_t \quad (9)$$

This model has a mean reverting element in the parameter restriction $\alpha_1 + \alpha_2 < 1$. Studying the factor r_t^2 one also notes that the future variance is not influenced by the signs of past returns, but solely on their magnitude. Since we believe that the return depends on the past innovations this means that "bad" and "good" news are taken in this model to impact the future conditional variance in the same way. However it has been shown showed that bad news (negative returns) tend to give larger future volatilities than good news (positive returns) (Andersen et al, 2006). This is called the

leverage effect in financial literature. Therefore two widely used models that take this asymmetric effect into account will also be studied and compared to the standard GARCH(1,1).

2.1.3.2 GJR(1,1)

Glosten, Jagannathan and Runkle (1993) introduced the GJR-GARCH, with specification as below.

$$\sigma_{t+1} = \alpha_0 + \alpha_1 r_t^2 + \alpha_2 \sigma_t + \Phi I r_t^2 \quad (10)$$

This model is sometimes called a threshold GARCH, due to the indicator I which equals one if the return in the period before is less than zero (and zero otherwise). This model takes the leverage effect into account by specifying Φ as another ARCH parameter that only comes into play when the return (or innovation) the period before was negative.

2.1.3.3 EGARCH(1,1)

The EGARCH was introduced by Daniel B. Nelson (1991). The EGARCH(1,1) has the following specification for the conditional variance.

$$\ln(\sigma_{t+1}) = \alpha_0 + \alpha_1 (|\epsilon_t| - E[|\epsilon_t|]) + \gamma \epsilon_t + \alpha_2 \ln(\sigma_t) \quad (11)$$

Where $\epsilon_t = \frac{r_t}{\sigma_t}$

This specification takes the aforementioned leverage effect into account as well, by letting the magnitude and the sign of $(|\epsilon_t| - E[|\epsilon_t|]) + \gamma \epsilon_t$ have different effects on the variance. By specifying σ logarithmically, the constraints on the parameters can be relaxed since the forecasted conditional variance will always be positive.

2.2 Measuring the loss – The QL loss function

To measure the forecasting accuracy of the different models it is necessary to specify a way to model the average error or loss. This is done by comparing the forecasted variance for a certain period with the realized variance. The different functional forms that relates the forecasted value to the realized variance are called loss functions. Of course the volatility is unobservable and an ex-ante proxy has to be utilized. I have chosen to work with the following specification for the daily realized variance.

$$\sigma_t = r_t^2 \quad (12)$$

Using squared returns as volatility has shown to be an unbiased but quite a noisy measure of volatility (Andersen et al, 2000). However due to data availability this specification is used nonetheless.

The loss function itself can be specified in almost as many ways as the model used for the forecasting. Patton shows that only two of the several models he tested were “robust” ones (Patton, 2006). The robust loss functions rank different model specifications in the same way regardless of the volatility proxy being used. I will only make use of one proxy for volatility, however it is still desirable to use a robust model, to not have the loss function interfere with the results. I have chosen to work with the following loss function:

$$QL_t: L(\sigma_t^2, h_t) = \frac{\sigma_t^2}{h_t} - \log \frac{\sigma_t^2}{h_t} - 1 \quad (13)$$

Where σ_t^2 is my ex post proxy of the conditional variance and h_t is the forecasted conditional variance for the time period. The mean loss for a period is then simply

$$\overline{QL} = \frac{1}{N} \sum_{t=1}^N QL_t \quad (14)$$

A useful quality of the QL loss function is that it asses the losses multiplicatively. Often used additive loss functions such as MSE or RMFSE depend on the realized variance. Since I’m going to study quite a long period I want a loss function that allows me to compare the loss over different periods.

Whereas MSE for example scales with the squared variance, the QL does not depend on the actual level of volatility and therefore the difference in forecast accuracy during the financial crisis compared to the entire sample can easily be calculated.

2.3 Translating the loss into volatility forecasting error

The loss from the different models are interesting in themselves, since a lower QL-value will mean a better forecasting accuracy. However, to be able to relate these values to financially applicable ones they need to be translated into a volatility forecasting error. To do this a concept from Brownlees et al (2011) will be discussed.

If the standard deviation of returns is taken to be the mean over the entire period, the average loss reported from the QL function can be translated into an average error in volatility prediction.

$$\frac{\sigma_{mean}^2}{x^2} - \log\left(\frac{\sigma_{mean}^2}{x^2}\right) - 1 = \text{average loss} \quad (15)$$

The loss function is symmetrical as can be seen in figure 2. This means that solving for x we will get two values, which is the average underestimate and overestimate of volatility.




Figure 2 The QL loss function for a hypothetical example.

Furthermore, by comparing the respective values of x for two different models an average volatility forecasting error reduction can be calculated. Say that model 1 has a lower QL loss than model 2. To compare the improvement in the average loss when switching from model 2 to model 1, I can use the function for the underestimate and overestimate respectively to find a span of the average volatility error reduction.

$$\text{Volatility error reduction} = 1 - \left| \frac{x_1 - \sigma_{mean}}{x_2 - \sigma_{mean}} \right| \quad (16)$$

2.4 From volatility forecasting error to pricing error

To calculate how using an inferior model affects the error in financial calculations the next two chapters demonstrate an example of how the volatility error relates to errors in option pricing and Value-at-Risk calculations.

2.4.1 Option Pricing

As stated earlier volatility forecasting is essential in derivatives pricing. For this reason I will discuss how the average forecasting loss can be translated into an average loss in option pricing. The well know Black-Scholes model for a call option reads

$$C(S, t) = N(d_1)S - N(d_2)Ke^{-r(T-t)} \quad (17)$$

Where $N(x)$ is the standard normal cumulative distribution.

The only unobservable quantity in this formula is the estimate of future volatility, used to calculate d_1 and d_2 . Consider an at-the-money call option with $S=K=1$ with one day left to maturity and a risk free interest rate of $r_f=1\%$. Even though the Black-Scholes model is far from perfect, for one day the assumption that the spot price has a lognormal distribution should hold (Heston & Nandi, 2000). So, keeping the other values fixed in the BS-formula the error reduction in pricing by switching from model 2 to model 1 is

$$BS \text{ error reduction} = 1 - \left| \frac{BS(x_1) - BS(\sigma)}{BS(x_2) - BS(\sigma)} \right| \quad (18)$$

This is calculated both for the under- and overestimates respectively, yielding a resulting error reduction in pricing of the same range as the reduction in volatility forecasting error (Bronwnlees et al, 2011).

2.4.2 Value-at-risk

The estimates of future volatility is crucial in Value-at-Risk calculations as well. A similar calibrated example is demonstrated by (Brownlees et al, 2011). As stated earlier the ability to predict the risk in an asset or a portfolio is crucial for the financial actor. The VaR concept is a way to say that with a certain confidence level $0 < \alpha < 1$ the loss shouldn't exceed the VaR for a certain horizon. It is a tool that attempts to quantify the amount of risk one is taking on by investing in a certain stock or portfolio. If the horizon is one day and the confidence level is α , then the loss is expected to exceed the VaR no more than $(1 - \alpha)$ percent of the days. This concept is illustrated by the probability density function in figure 3, where the VaR in this case represents the fifth quantile of the return distribution. This concept tries in no way a way to put a maximum value on the loss, rather it is a tool quantify the maximum loss within a certain confidence level.




Figure 3 Probability density function for Value-at-Risk. The light area to the left represents corresponds to 5 % of the total area.

Using the deduced over- and underestimates the reduction in the daily 99 % VaR can be calculated according to:

$$VaR \text{ error reduction} = 1 - \left| \frac{\phi^{-1}(0.01, 0, x_1) - \phi^{-1}(0.01, 0, \sigma)}{\phi^{-1}(0.01, 0, x_2) - \phi^{-1}(0.01, 0, \sigma)} \right| \quad (19)$$

Where ϕ^{-1} is the inverse cumulative Gaussian probability density function, assuming a mean of zero.

Just as for the Option price, the average error reduction in VaR has the same range as the error reduction in volatility forecasting.

3. Data & Method

3.1 Software

The estimation and forecasting was done in Matlab. It was chosen because it is rigorous, flexible and much used. Matlab estimates the GARCH models through maximum likelihood estimation. Using the historical data of the log-returns and the different specifications of the GARCH models, the software finds the most likely value of the parameters to have produced such a series of returns. The forecasting is trivial since I only use one day ahead forecasts of the conditional variance. As can be seen in the specifications of the respective models the conditional variance at time $t+1$ is known at time t .

3.2 Data used and method of forecasting

The data used consists of daily closing prices of the Stockholm 30, Helsinki 25 and Nordic 40 indices for the entire years of 2002 to 2013. The period of 2008-2013 is forecasted using a rolling window technique. All the data up to 2008 is used to forecast the first day of 2008. Then the oldest observation (the first day of 2002) is dropped and the first day of 2008 is added to the in-sample data to forecast the variance of the second day of 2008. The data was downloaded from Nasdaq OMX. The Stockholm and Helsinki indices represent the stocks of the 30 and 25 most traded companies in the respective countries. The Nordic 40 index represents the 40 most traded Nordic companies, there amongst companies from the Stockholm and Helsinki markets. Table 1 shows the characteristics of the different markets. Figure 4 to figure 6 below shows the price and daily returns of the respective indices for the entire 2002-2013 period.

Table 1 Features of the log-return for the different indices.

	Observations	Mean	Stdv	Kurtosis	Skewness
Stockholm	3016	1.5616e-4	0.0152	6.9843	0.1119
Helsinki	3016	1.8907e-4	0.0147	6.5887	-0.0243
Nordic	3053	8.3891e-4	0.0157	6.7606	-0.0090




Figure 4 Daily closing price (top) and daily log-returns (bottom) for the OMX Stockholm 30 index 2002-2013.




Figure 5 Daily closing price (top) and daily log-returns (bottom) for the OMX Helsinki 25 index 2002-2013.




Figure 6 Daily closing price (top) and daily log-returns (bottom) for the OMX Nordic 40 index 2002-2013.

The Stockholm, Helsinki and Nordic indices exhibit great similarities. This is natural since the OMX Nordic 40 is made up of stocks from among others the Stockholm and Helsinki indices. Moreover both Sweden and Finland are members of the EU and are presumed to have quite highly correlated business cycles.

Table 2 shows properties of the log-returns of the indices during the crisis period. The number of days studied correspond to a period of high volatility and negative returns for the respective indices. The Stockholm and Nordic indices are studied during 45 and 65 observations respectively. The Helsinki index is studied for 130 observations. These time periods have been chosen to correspond to a time period representing the crisis for the respective index. The Helsinki index had a longer turbulent period with a smaller average daily loss than the Stockholm and Nordic indices.

Table 2 Properties of the log-returns of the indices for the crisis forecasting.

Index	Observations	Mean	Standard deviation	Skewness	Kurtosis	Start date
Stockholm	45	-0.076	0.0361	0.5810	3.6133	27-Aug-08
Helsinki	130	-0.0042	0.0311	0.4082	3.6225	20-Aug-08
Nordic	65	-0.071	0.0380	0.4676	3.2730	25-Aug-08

4. Empirical Results

4.1 Forecasting results for the entire period

In table 3 below the average QL losses for the entire 2008-2013 forecasting period are reported. The specifications of EWMA and the historical variance with the smallest loss are also reported. For comparison table 4 shows all the estimated specifications of these models.

Table 3 Average QL losses during the 2008-2013 forecasting period. The model with the smallest loss for the respective index is marked in blue.

	GARCH	GARCH t	EGARCH	EGARCH t	GJR	GJR t	Historical	EWMA
Stockholm	1.4750	1.4769	1.4456	1.4448	1.4371	1.4347	1.5108	1.4787
Helsinki	1.4033	1.4035	1.3618	1.3611	1.3669	1.3696	1.4330	1.4033
Nordic	1.4875	1.4899	1.4553	1.4546	1.4535	1.4535	1.5256	1.4917

Table 4 Average QL losses during the 2008-2013 forecasting period. The model specification for the historical variance and EWMA with the smallest loss for the respective index are marked in blue.

Historical	10	20	50	100	EWMA	0.7	0.8	0.9	0.94
Stockholm	1.5832	1.5110	1.5108	1.5961		1.6836	1.5739	1.4926	1.4787
Helsinki	1.5052	1.4371	1.4330	1.5159		1.6015	1.4914	1.4149	1.4033
Nordic	1.6032	1.5385	1.5256	1.6137		1.7281	1.5996	1.5075	1.4917

The results are consistent in that 50 days of historical data seems to be optimal for all of the indices. 20 days are second best, 10 days third best and 100 days is the worst specification for all of the markets. These results are interesting as they tell us something about the pattern of volatility clustering for these markets. For the EWMA model $B=0.94$ is the superior model specification for all of the indices. Declining B results in declining forecast accuracy in all of the markets. These results correspond well with the performance of the different specifications of the historical variance models, a larger Beta is similar to including more days of historical data. Since the value of Beta is just

0.045 by the 50th day for B=0.94 this can be thought of as a historical variance model with 20-50 days of significant historical data and with a declining impact on the forecasted variance. Concerning the GARCH models the results are a little more diverse. Generally, assuming normal or t-distribution has a small impact on the forecasting error over the entire period. t-distribution and normal distribution performs better in 4 cases each with GJR for the Nordic index yielding equally large losses with normal and t-distribution. EGARCH with t-distribution is superior for Helsinki whereas GJR with t-distribution is superior for Stockholm. The standard GARCH is the worst performing of the GARCH models, and have a loss very close to the EWMA(B=0.94) in all of the markets. To summarize, the more sophisticated the model, the better the result in this case.

To translate the QL-losses into average volatility error reduction I proceed as outlined in chapter 2.5. As stated earlier this is done to translate the effect of differences in forecasting ability into error reductions in financial calculations. In table 4 the reduction in forecasting error by switching from the worst performing to the best forming model of the GARCH models are reported.

Table 5 Average Error reduction in volatility forecasting by switching from the worst to the best of the GARCH models for the respective markets during the 2008-2013 forecasting period.

Index	Volatility error reduction
Stockholm	0.80-3.31%
Helsinki	0.85-3.41%
Nordic	0.68-2.84%

As can be seen the reduction in forecasting error is somewhere around 1-3 percent for all of the markets, when switching from the best performing of the GARCH models to the worst (standard GARCH in all cases). The calculated reduction in option pricing and VaR calculations also match these number down to three decimals, as expected since it was a calibrated example. The aim is to compare these numbers to the error reduction from the crisis forecasting, to translate the models' relative performances into relatable economic magnitudes.

Figures 7-9 below shows the one day ahead annualized volatility forecast versus the “realized volatility” (also annualized) for the model with the smallest loss for the respective index.

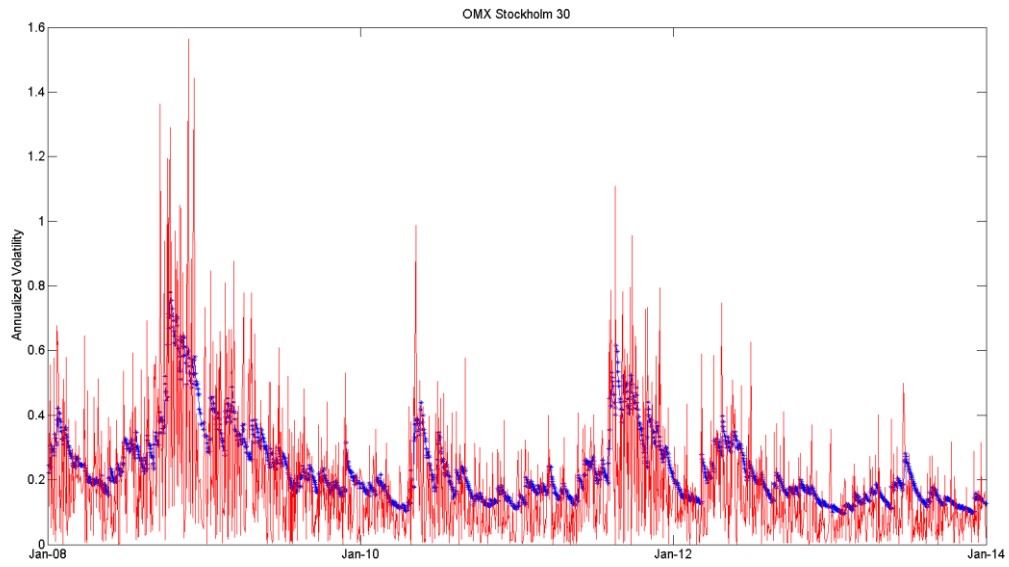


Figure 7 Forecasted volatility (thick blue line) versus realized volatility (red line) for OMX Stockholm 30 for 2008-2013.




Figure 8 Forecasted volatility (thick blue line) versus realized volatility (red line) for OMX Helsinki 25 for 2008-2013.




Figure 9 Forecasted volatility (thick blue line) versus realized volatility (red line) for OMX Nordic 40 for 2008-2013.

By comparing the thick dotted line with the thinner lines of realized volatility one notes that while the models seem to follow the general trend of the volatility clustering well, they fail to account for extreme values. The forecasted variance is smoothed and does not reach the extreme high or low values of volatility observed. This effect is an innate quality of the stock market and a consequence of the fundamental randomness of the innovations.

4.2 Forecasting results for the crisis

Table 6 shows the average QL-losses for the forecasting of the crisis.

Table 6 Average QL-losses during the 2008 crisis. The model with the smallest loss for the respective index is marked in blue.

	GARCH	GARCH t	EGARCH	EGARCH t	GJR	GJR t	EWMA	Historical
Stockholm	1.6573	1.6492	1.5514	1.5232	1.5070	1.4577	1.5668(0.90)	1.5520(20)
Helsinki	1.5159	1.5053	1.5085	1.5034	1.4828	1.4833	1.4923(0.90)	1.4955(20)
Nordic	1.6370	1.6342	1.5623	1.5476	1.5434	1.5217	1.5892(0.90)	1.5875(20)

As expected all the models performed worse during the crisis. As stated in the theory section the QL loss function is not sensitive to the level of volatility. Therefore these losses solely represent the

increased loss in forecasting accuracy, not the fact that the volatility was actually heightened during this period. GJR seems to be the dominating model for forecasting, providing smaller losses for all the indices. The standard GARCH is the worst performing model in all of the markets. For Stockholm and Nordic the GARCH yields a substantially worse forecasting than EGARCH and GJR. The forecast of these markets uses fewer observations than the Helsinki index and it might be the case that the GARCH is too slow in adapting to a crisis. The crisis period studied for the Stockholm index only has 45 observations. Here the GARCH model performs the worst compared to the GJR and EGARCH models. During the longer observation of the crisis for the Helsinki index the GARCH model still forecasts worse, however not as badly as for the other two indices. For the 3 models and 3 markets, the t-distribution is the better model in 8 of the 9 instances. The improvement ranges from quite small, to quite substantial, as for the GJR model in the case of the Stockholm and Nordic indices. The fatter tails of the t-distribution curve were apparently able to better capture the distribution of the innovations during the crisis. However it is still the case that the distribution overall has a much smaller impact on the forecasting ability than has the model itself. Concerning the historical variance model, it is noteworthy that now 20 days of historical data yields the better forecast for all of the indices, compared to 50 days for the entire period. The EWMA model now yields the best result for $B=0.9$ for all of the indices. However it is outperformed by the historical variance model for Stockholm and Helsinki and just slightly better than the historical model for the Nordic index. It should be noted that the results for the historical variance model were significantly worse when more days were used to calculate the variance.

Table 7 shows the average error reduction, again switching from the worst performing of the GARCH models to the best performing one for every index.

Table 7 Average Error reduction in volatility forecasting by switching from the worst to the best of the GARCH models for every index during the forecasting of the crisis.

Index	Average error reduction
Stockholm	3.42-14.25
Helsinki	0.6-2.56
Nordic	01.96-8.43

As can be seen the error reduction from the crisis is much more significant than for the entire period on the Stockholm and Nordic markets, due to the poor performance of the standard GARCH. The span of average error reduction in volatility forecasting, option pricing and VaR is about five times greater than for the entire period for Stockholm and three times greater for Nordic. These numbers

are telling, indicating that there is a lot more at stake when choosing the forecasting model in very turbulent periods. For the Helsinki index the average error reduction during the crisis is about the same as for the entire period, due to the fact that the standard GARCH did not perform as badly compared to the other models as on the other markets.

Figure 10 to figure 12 below shows the one day ahead volatility forecasts for the best performing model on every index. The number of observations are different for the different markets, so the x-axis has different scaling in each of the figures.




Figure 10 Forecasted volatility (blue line) versus realized volatility (black dots) for OMX Stockholm 30 from 2008-08-27.

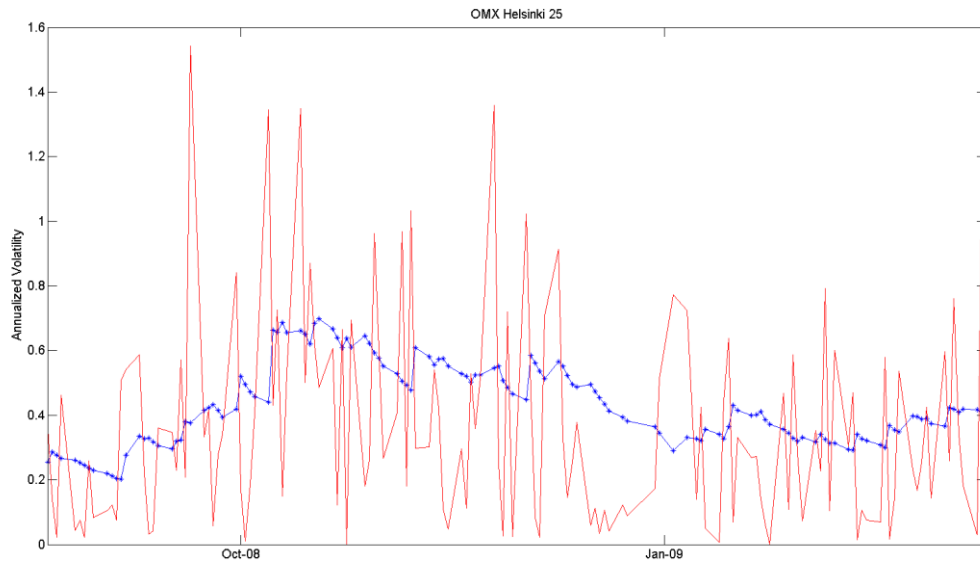


Figure 11 Forecasted volatility (blue line) versus realized volatility (black dots) for OMX Helsinki 25 from 2008-08-20.

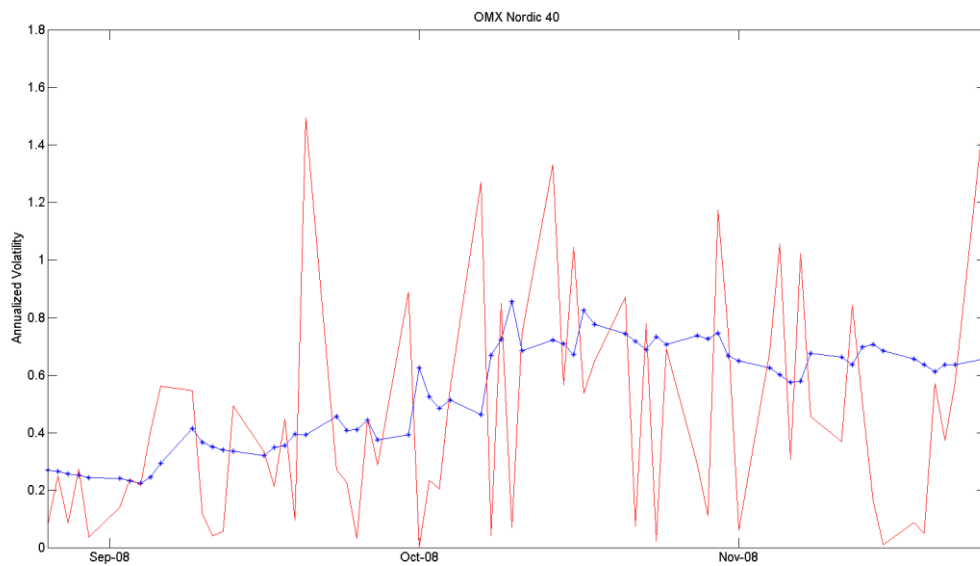


Figure 12 Forecasted volatility (blue line) versus realized volatility (black dots) for OMX Nordic 40 from 2008-08-25.

From these figures it is apparent that the even the best fitted model has a hard time adapting to quick swings in volatility. For the crisis, the pattern is the same as for the entire period, the models fail to capture the extreme highs and lows of volatility. However the models seem to adapt after a certain time of heightened volatility, as is expected.

5. Conclusions

I have compared the volatility forecasting ability of several models on the Nordic stock market. For the entire 2008-2013 period an EGARCH with t-distribution was the best model for the Helsinki index, a GJR with t-distribution for the Stockholm index and a GJR (with normal and t-distribution) for the Nordic index. In general the most sophisticated models, GJR and EGARCH, performed best. The slightly less sophisticated GARCH performed about as well as EWMA, while the historical variance yielded the greatest loss. When studying the entire period assuming a normal or t-distribution had quite a small impact on the forecasting error. Moreover the difference in forecasting accuracy of the different models were much lower than during the crisis period. The difference in volatility error reduction for the worst and best model ranged from about 1-3%.

During the 2008 crisis the forecasting losses increased for all of the markets and forecasting methods. GJR was the model with the smallest loss for all of the indices and GARCH was the one with the greatest loss. For this period t-distribution yielded a substantially lower loss in several of the model specifications. This apparently means that assuming a t-distribution allows the model to better capture the more extreme values of the innovations. The relative forecasting error reduction between the worst and best of the models were shown to be about five and three times higher during the crisis for the Stockholm and Nordic markets respectively. It was shown that the VaR and option pricing errors could be substantially decreased by switching from a GARCH to a GJR-GARCH. During periods of financial turmoil it is of even greater importance to select the best model as the average error reduction between different models greatly increased.

In further studies it would be interesting to perform an extensive study on these markets. One could use significantly more models and forecast over different time periods, varying the amount of historical data used for the model estimation. It would also be interesting to fit the GARCH models via another method. Instead of maximum likelihood one could for example minimize the squares of the "forecasted" variances. It is also possible to imagine a method where the forecasts from different models are combined into a single forecast. The number of possible studies are endless.

It would also be interesting to undertake a similar study of the 2008 financial crisis on other markets to see if the results are consistent. To conclude, it must be noted that the fact that a certain model performed better for a certain period is no guarantee that this will be the case for other periods.

6. Reference List

- Andersen T, Bollerslev T, Christoffersen P, Diebold F, (2006), "Volatility and correlation forecasting". *Handbook of Economic Forecasting*. Amsterdam: North-Holland.
- Andersen T, Bollerslev T, Diebold F, Ebens H, (2000), "The Distribution of Stock Return Volatility". *PIER Working Paper 01-001*.
- Basel Committee on Banking Supervision, (2010), "Basel III: A global regulatory framework for more resilient banks and banking systems".
- Bollerslev T, (1986), "Generalized Autoregressive Conditional Heteroskedasticity". *Journal of Econometrics* Vol. 31, No. 3.
- Brownlees C, Engle R, Kelly B, (2011), "A practical guide to volatility forecasting through calm and storm". *The Journal of Risk* Vol. 14, No. 2.
- Ederington LH, Guan W, (2005), "Forecasting Volatility". *Journal of Futures Markets* Vol. 25, No. 5
- Engle F, (1982), "Autoregressive Conditional Heteroscedasticity With Estimates of the Variance of United Kingdom Inflation". *Econometrica* Vol. 50, No. 4.
- Glosten L, Jagannathan R, Runkle D, (1993), "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks". *Journal of Finance*, Vol. 48 No. 5.
- Hansen P, Lunde A, (2001), "A comparison of volatility models: Does anything beat a GARCH(1,1)?" *Working Papers series, No. 84, University of Aarhus*.
- Heston S, Nandi S, (2000), "A Closed-Form GARCH Option Valuation Model". *The Review of Financial Studies*, Vol. 13 No. 3.
- Nelson D, (1993), "Conditional Heteroskedasticity in Asset Returns: A New Approach". *Econometrica*, Vol. 59, No. 2.
- Patton A, (2011), "Volatility forecast comparison using imperfect volatility proxies". *Journal of Econometrics*, Vol. 160, No. 1.
- Poon S, Granger C, (2003), "Forecasting Volatility in financial Markets: A Review". *Journal of Economic Literature*, Vol. 41, No. 2.
- Richardson M, Smith T, (1993), "A Test for Multivariate Normality in Stock Returns". *Journal of Business* Vol.66, No. 2.

