AMSI
AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

National Collaboration in the Mathematical Sciences

AMSI Vacation Research Project
A comparison of the performances of several methods for finding Value at Risk

Rupert Kuveke
La Trobe University, Melbourne
Supervisor: Associate Professor Paul Kabaila

February 22, 2014

Abstract

The aim of this research project is threefold:

- Firstly, background theory on several common methods for finding the financial risk measure Value at Risk (VaR) will be briefly covered.
- Secondly, the VAR methods of Historical Simulation, Filtered Historical Simulation and fitting a GARCH model in combination with a standardised t distribution will be conducted, using S&P 500 Index data from 1988 to 2013.
- Finally, the performances of these VaR methods will be compared and discussed.

1 Value at Risk

Value at Risk is a financial risk measure used in the corporate world to provide a single figure which attempts to convey the total risk a specific portfolio of assets is exposed to, over a specific period of time (Hull 2011). As defined by Tsay (2005), Value at Risk estimates the amount by which a company's position may decline over a given period of time due to the movements of financial markets.

Financial institutions may utilise Value at Risk calculations to assess and manage their risk levels. Regulatory committees such as the Basel Committee on Banking Supervision may also utilise Value at Risk calculations in order to set the margin requirements for banks and curtail any excessive risk taking by financial institutions.

For the purposes of this project, we shall follow Christoffersen (2012), and define Value at Risk as the loss that a portfolio may incur over a specific period of trading days, such that only $p{\cdot}100\%$ of losses exceed the calculated amount. Due to the nature of this project, the percentage of the portfolio which is at risk of being lost, rather than the dollar value, will be reported. This will help to highlight the disparities between the efficacy of the different models assessed.

1.1 The S&P 500 Index

The calculations in this project are all based on data from the *S&P 500 Index*. The S&P 500 Index, or Standard & Poors 500 Index, is an American stock market index. It includes 500 leading American companies from various industries whose combined market capitalisation is approximately 80% of total American equities, according to S&P Dow Jones Indices McGraw Hill Financial (2013). As a result, the S&P 500 Index is considered to be the single best indicator of the American stock market and economy.

Investors can access index funds such as the Vanguard 500 Index Fund, which attempt to mimic the S&P 500 Index in composition. For the purposes of this project, if we assume that the 500 American companies comprising the S&P 500 Index do not change over the period under assessment, then we may treat the S&P 500 Index as a single stock, which is readily available to investors. This allows us to distil the characteristics of the stocks of the various companies comprising the index into the one equity instrument.

Therefore, while the VaR calculations in this project will be performed under the assumption of a single-stock portfolio, the resultant VaR figures may give a rough indication of the VaR of a diversified portfolio consisting of numerous stocks in various American companies.

1.2 Volatility

Financial stock markets often display high levels of *volatility*, which is reflected in the pricing of financial instruments. Volatile assets' returns have inherent uncertainty, and can be difficult to accurately predict (Hull 2011). When assessing financial markets, we may define volatility as 'the conditional standard deviation of the underlying asset return' (Tsay 2005, p. 97).

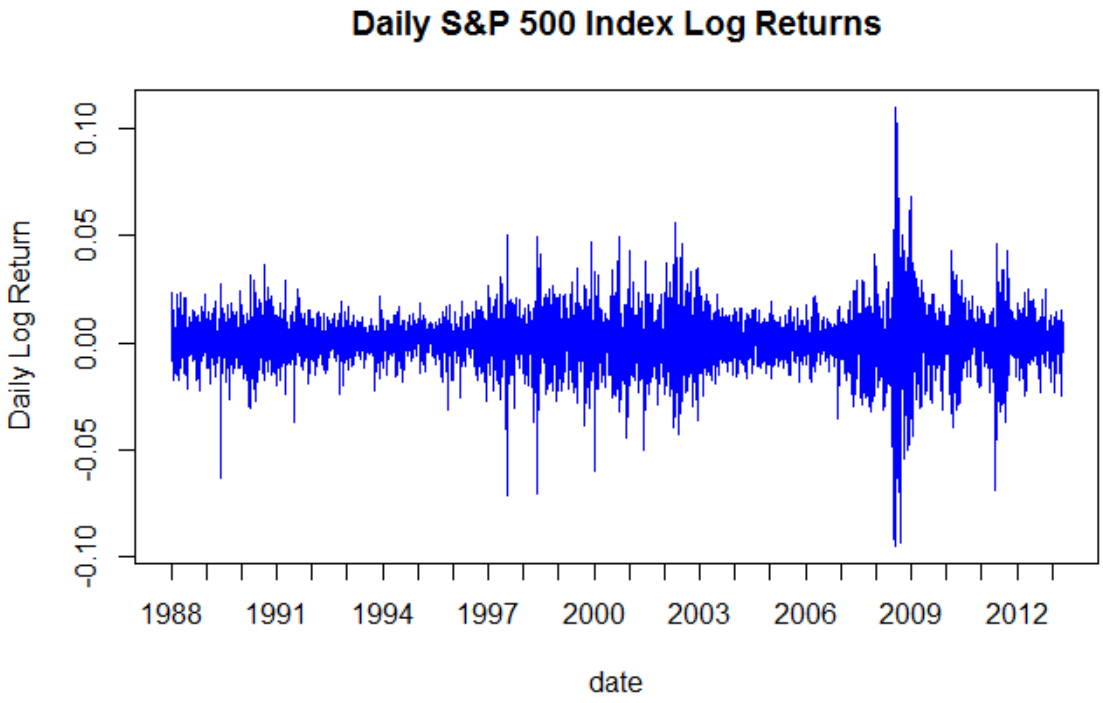

Figure 1: Daily log returns of the S&P 500 Index between June 1988 - June 2013. Here we can clearly see the phenomenon of volatility clustering occurring during the GFC.

1.3 Volatility Clustering

Stock markets can experience periods of *volatility clustering*. The phenomenon of volatility clustering, whereby 'large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes' (Mandelbrot 1963, p. 418) can result in periods of extreme market fluctuations, such as the Global Financial Crisis (GFC). The accuracy of VaR calculations is particularly important at such times, and hence the development of models which can accurately account for volatility clustering is vital.

2 Historical Simulation Value at Risk

The first method used in this project to calculate VaR is *Historical Simulation*. Historical Simulation (HS) is the 'most commonly used method for computing VaR' (Christoffersen 2012, p.21). By assessing historical data pertaining to a specific portfolio, HS attempts to answer the question: How much does the portfolio stand to lose tomorrow?

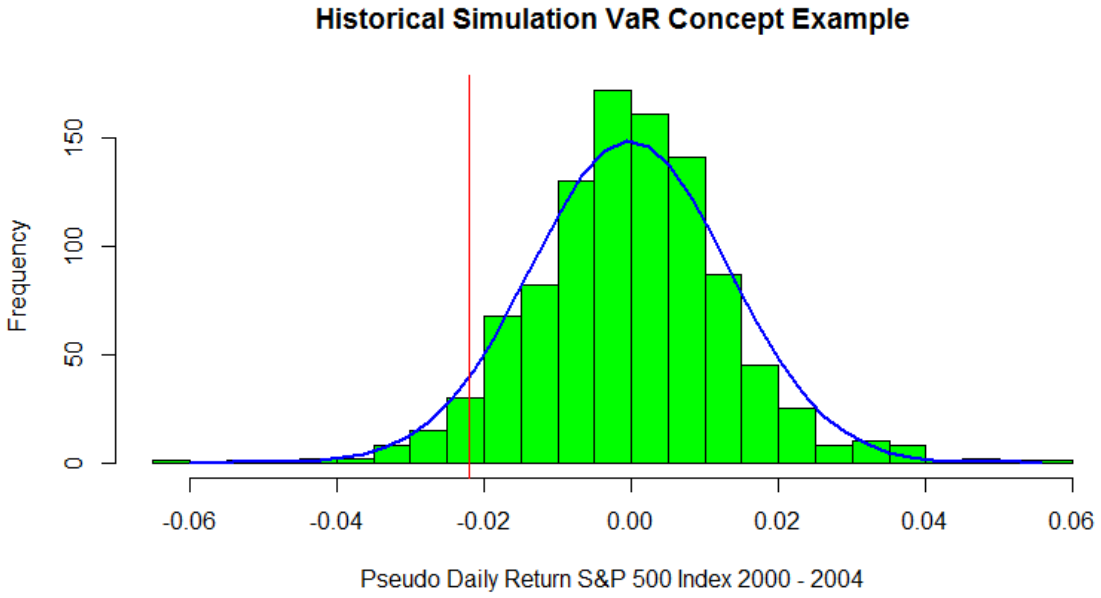

Figure 2: Visual example of the Historical Simulation method for finding VaR. 1000 S&P 500 Index daily log returns have been sorted in ascending order, such that the probability of returns being to the left of the red line is 5%. Therefore we can say that for $m = 1000$, the daily VaR of the portfolio is 2.169%, with 5% probability. Note that the histogram of returns does not follow a normal distribution curve (in blue).

It should be noted that the following notation and equations are based on the primary reference source, Christoffersen (2012). Many of these definitions are vital for the development of the concept of Historical Simulation as a method, and as such cannot be changed significantly.

To define Historical Simulation, using the structure and terminology of Christoffersen (2012), we begin by considering a portfolio of n assets.

Let the stock price of asset i be $S_{i,t}$, and let today be day t. Given that we own $N_{i,t}$ shares of asset i, the value of our portfolio today is

$$V_{PF,t} = \sum_{i=1}^{n} N_{i,t} S_{i,t} \tag{1}$$

Historical Simulation uses the present asset allocation of a portfolio, and applies these allocations to the historical asset prices of assets $i = 1, ..., n$.

This allows us to calculate a sequence of pseudo portfolio values, which would have occurred if today's portfolio allocations had remained constant throughout the period of time being assessed.

For instance, the pseudo portfolio value b days ago is

$$V_{PF,t-b} = \sum_{i=1}^{n} N_{i,t} S_{i,t-b} \tag{2}$$

This value for $V_{PF,t-b}$ is considered to be a hypothetical or pseudo value as the asset allocation of any portfolio will generally alter over time. We may now define the pseudo log return of the portfolio as

$$R_{PF,t} = \ln(V_{PF,t}/V_{PF,t-1}) \tag{3}$$

The HS method calculates a sequence of m daily hypothetical portfolio returns, using the current portfolio allocations in conjunction with m past prices of each of the portfolio's underlying assets.

Once this collection of hypothetical returns are calculated, using equation 3, they are sorted in ascending order. It is assumed that the distribution of $R_{PF,t+1}$ is well approximated by the histogram of the sequence of values

$$\{R_{PF,t+1-\tau}\}_{\tau=1}^{m} \tag{4}$$

If we consider the distribution of the pseudo log return sequence in histogram form, as has been done in Figure 2, then the VaR of the portfolio, for probability p, is the $p \cdot 100th$ percentile of $\{R_{PF,t+1-\tau}\}_{\tau=1}^{m}$.

Hence we may write

$$VaR_{t+1}^{p} = -Percentile(\{R_{PF,t+1-\tau}\}_{\tau=1}^{m}, 100p) \tag{5}$$

Therefore, the VaR_{t+1}^{p} is calculated as being the hypothetical return such that only $100p\%$ of the hypothetical returns are smaller than VaR_{t+1}^{p}.

2.1 Historical Simulation Daily Value at Risk applied to S&P 500 Index data

The following graphs show the Daily Value at Risk for a theoretical portfolio mirroring the S&P 500 Index, for different values of m and p, calculated using Historical Simulation. As expected, there is a clear increase in the Daily Value at Risk during the Global Financial Crisis, a period of high volatility.

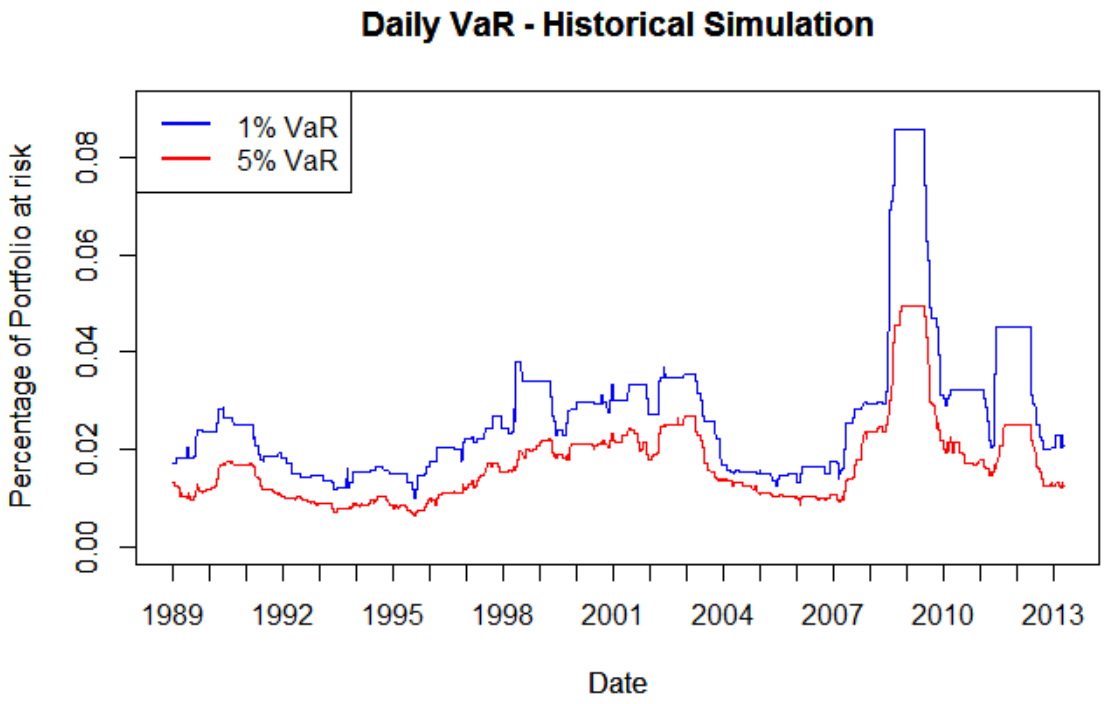

Figure 3: Daily VaR June 1989-June 2013 using $m = 250$. Here the blue line corresponds to a VaR for $p = 1\%$ and the red line to a VaR for $p = 5\%$.

As can be seen from Figures 3, 4 and 5, reducing the value of p, the probability that at least the reported VaR will be lost on a specific day, leads to consistently higher VaR values. This is a logical result. For example, if there is a 5% probability of experiencing a VaR of 2% on a given day, then within that 5% section of the normal distribution lies the section for the VaR which has a 1% probability of occurring, and therefore this latter VaR value will be greater than 2%. Mathematically, if we have probability w of having a VaR of $V\%$ or greater for a portfolio P, then for q, where $q < w$, the VaR of portfolio P will be $> V\%$.

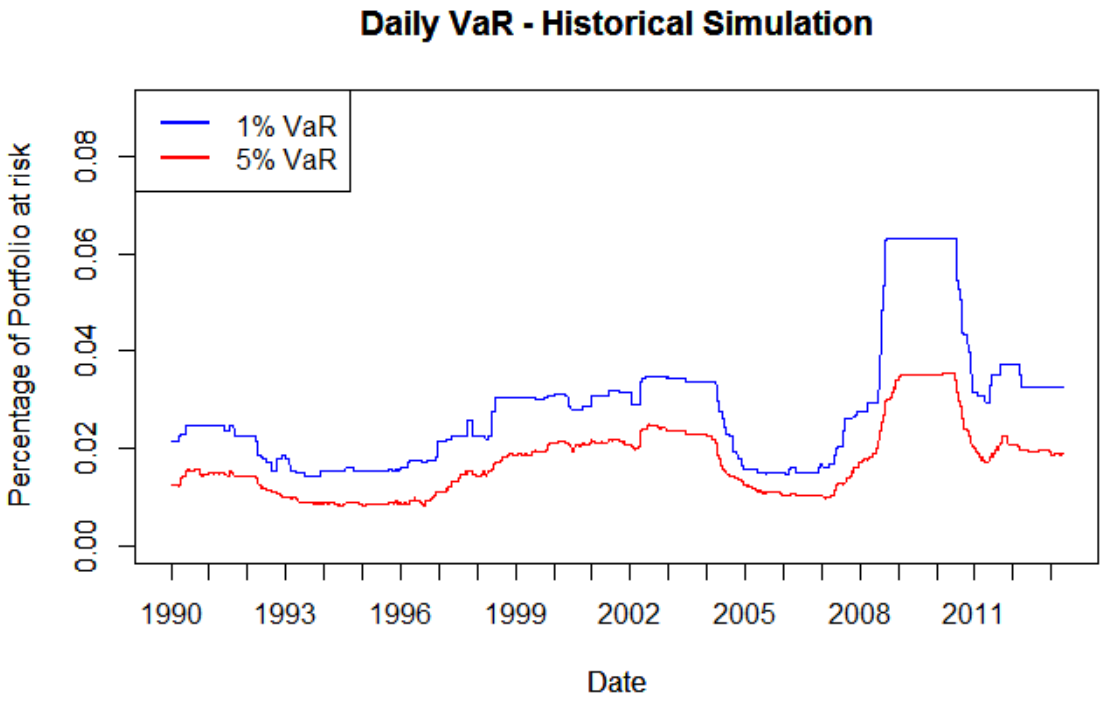

Figure 4: Daily VaR June 1990-June 2013 using $m = 500$. Here the blue line corresponds to a VaR for $p = 1\%$ and the red line to a VaR for $p = 5\%$.

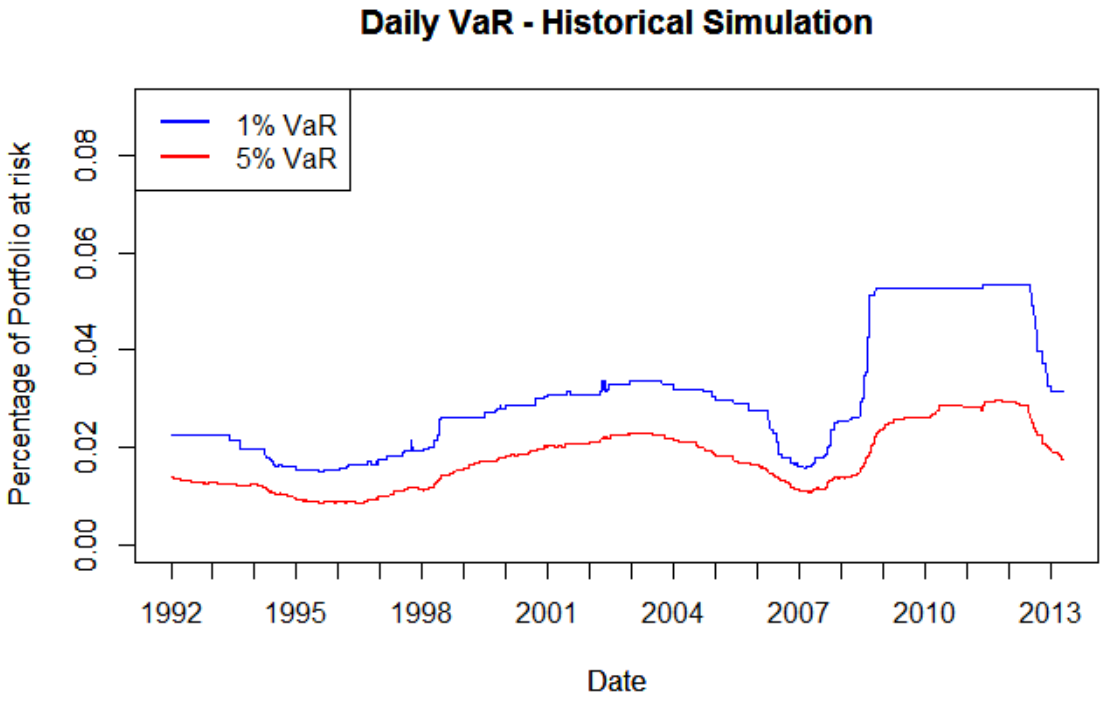

Figure 5: Daily VaR June 1992-June 2013 using $m = 1000$. Here the blue line corresponds to a VaR for $p = 1\%$ and the red line to a VaR for $p = 5\%$.

2.2 Historical Simulation Results

The Historical Simulation method for finding Value at Risk is relatively straightforward to understand and calculate. As it is a technique which does not rely on 'any particular parametric model' (Christoffersen 2012), it is superior to model-based approaches which rely on a poorly chosen model. However, this method does have several flaws. The number of days m of theoretical historical return which is used as a basis for determining the daily VaR is completely arbitrary. There is no global standard set in place, although values 'between 250 and 1000 days corresponding to approximately 1 to 4 years' (Christoffersen 2012, p.23), are typically chosen. While this does allow for tailored models to be created, it may be more beneficial from a corporate perspective to have an industry standard value to ensure accountability.

Historical Simulation appears to be based upon a logical foundation. However, it is apparent that this method overlooks the fact that it is inappropriate during a time of high volatility to be calculating a firm's VaR by using historical market data from times of normal market volatility. Clearly, the Historical Simulation method for finding VaR responds poorly to volatility fluctuations.

It is at times of high volatility that the Historical Simulation method reveals a serious flaw. Figures 3 and 5 in particular highlight this flaw. In the aftermath of the initial fallout from the Global Financial Crisis (GFC) in 2008, the HS model with $m = 250$ quickly returns to pre-GFC levels of VaR after a large spike in 2009. This is due to the model only using the past 250 days as a basis for how the current market will perform. As a result, by 2010, the majority of the volatility experienced in financial markets during the GFC is implicitly considered to be no longer relevant by the model.

Conversely, the HS model with $m = 1000$ remains almost constant with VaR for $p \cdot 100\% = 1\%$ at approximately 6% from January 2009 to January 2013. It is no coincidence that this period is equivalent to 1000 trading days. By utilising an m value of 1000, the high volatility experienced during the GFC continues to be taken into consideration for many years after the event. This results in VaR figures which are much higher than corresponding figures calculated using lower m values, as is the case in figures 3 and 4.

Ultimately, this means that using a low m value, such as $m = 250$ will lead to potential underestimation of VaR in the short term aftermath of periods of high volatility, and that using a high m, such as $m = 1000$ will lead to potential overestimation of VaR for extended periods of time following a period of high volatility.

For a statistical model, adding more information should lead to a more accurate and robust model. However the opposite is true for Historical Simulation. As the value of m increases, the HS VaR plot becomes less responsive to volatility fluctuations. This result has important implications for the corporate world, as Historical Simulation is considered to be widely used in practice. Therefore, analysts using this inaccurate method to calculate VaR may be unwittingly exposing their company to unnecessary levels of risk.

3 Filtered Historical Simulation

Filtered Historical Simulation (FHS) builds upon the Historical Simulation method for finding VaR. Filtered Historical Simulation combines the model-free HS method with 'model-based methods of dynamic variance, such as GARCH' (Christoffersen 2012, p.124).

As a result, this method attempts to better account for changing volatility levels, by being conditional on past information. Therefore, the definition of VaR for the FHS method, and indeed any method incorporating a GARCH model, is distinct from the definition of VaR used in HS. The HS VaR is a simple unconditional definition, whereas for FHS, the definition of VaR is conditional on the information utilised in estimating the conditional GARCH(1,1) dynamic volatility value σ_t.

3.1 GARCH Models

A *GARCH* model, or generalised autoregressive conditional heteroskedastic model, is a generalised ARCH model. ARCH models provide a 'systematic framework for volatility modeling' (Tsay 2005, p. 102). In comparison to ARCH models, GARCH models reduce the number of parameters required in order to 'adequately describe the volatility process of an asset return' (Tsay 2005, p. 113). As a result, a GARCH(1,1) model will yield similar information to an ARCH model of a higher order, while being significantly easier to calculate.

By using a GARCH model, we can account for volatility clustering to a greater degree than by using Historical Simulation. 'GARCH models are very strong in incorporating volatility persistence of returns. Without such models, spikes in volatility cannot be captured quickly and VaR predictions may not reflect sudden changes in volatility' (Supervision Regulation and Credit, Policy Analysis Unit, Federal Reserve Bank of Richmond 2010, p. 13).

A GARCH(1,1) model, which relies on 'one lag of returns squared and one lag of variance' (Christoffersen 2012, p. 78) is the simplest form of a GARCH model, and may be written as

$$\sigma_{t+1}^2 = \omega + \alpha R_t^2 + \beta \sigma_t^2, \text{where } \alpha + \beta < 1 \tag{6}$$

Here ω represents the weighted average of the long-run variance. The combined value of $\alpha + \beta$ is known as the model's persistance, the size of which reflects how long a shock to the long-run average variance will persist (Christoffersen 2012). Christoffersen (2012) notes that this value therefore cannot equal 1, as this would imply that shocks to variance would persist forever.

An important aspect of a GARCH(1,1) model is that if the R_{t-1}^2 or σ_{t-1}^2 term is large, then the GARCH(1,1) model will produce a large σ_t^2. This generates the behaviour of volatility clustering (Tsay 2005).

3.2 Filtered Historical Simulation with a GARCH(1,1) model

As stated in section 2, the following notation and equations are based on the primary reference source, Christoffersen (2012). Many of these definitions are vital for the development of the concept of Filtered Historical Simulation as a method, and as such cannot be changed significantly.

To define Filtered Historical Simulation, using the structure and terminology of Christoffersen (2012), we begin by considering a sequence of historical returns, $\{R_{PF,t+1-\tau}\}_{\tau=1}^{m}$, as developed in Section 2.

Using this sequence of values, the α, β and ω values for a GARCH(1,1) model can be estimated. Incorporating this GARCH(1,1) model in our calculation of VaR allows us to 'capture important features of returns data, while allowing for flexibility in accommodating specific aspects of individual assets' (Christoffersen 2012, p. 70).

As a result, we have a set of estimated standard deviations $\{\sigma_{PF,t+1-\tau}\}_{\tau=1}^{m}$, which represent the volatility of the portfolio on a specific day.

We can write

$$R_{PF,t+1} = \sigma_{PF,t+1} z_{t+1} \tag{7}$$

Here, z_t is defined as 'today's shock to return' (Christoffersen 2012, p. 77).

With this information, we can standardise the historical returns to adjust them for differing levels of volatility and account for their nonnormality (Christoffersen 2012).

Namely, we can calculate

$$\hat{z}_{t+1-\tau} = R_{PF,t+1-\tau}/\sigma_{PF,t+1-\tau}, \text{for } \tau = 1, 2, ..., m \tag{8}$$

This provides the set of standardised returns $\{\hat{z}_{t+1-\tau}\}_{\tau=1}^{m}$.

Hence Filtered Historical Simulation Value at Risk can be defined as

$$VaR_{t+1}^{p} = -\sigma_{PF,t+1} Percentile\{\{\hat{z}_{t+1-\tau}\}_{\tau=1}^{m}, 100p\} \tag{9}$$

By rescaling the standardised returns with the current $\sigma_{PF,t}$, the FHS VaR values attempt to account for current market volatility.

3.3 Filtered Historical Simulation Daily Value at Risk applied to S&P 500 Index data

The α, β and ω values of the GARCH(1,1) model were estimated by using the fGarch package in the statistical software package R. In order for these parameters to all be statistically significant at the 0.05 level of significant, it was necessary to use a minimum of 1250 days of pseudo daily log returns from the S&P 500 Index data. Using a lesser amount of days resulted in one or more of the parameter estimates being statistically insignificant.

In the process of fitting this GARCH(1,1) model to the S&P 500 Index data, sections of 250 days were used to ascertain the the number of days required to appropriately estimate the parameter values. That is, sections of 250 days, 500 days, 750 days and 1000 days were assessed before settling on 1250 days as the most appropriate selection. The 250 day skip was chosen as there are 252 trading days in a year; hence a section of 250 days corresponds to approximately one year.

The resultant values are listed in Table 1 below

Parameter	Estimated Value
α	0.011792
β	0.98223
ω	0.00000038683

Table 1: Estimated Parameter values for GARCH(1,1) model

Therefore the GARCH(1,1) model used was of the form

$$\sigma^2_{PF,t+1} = 0.00000038683 + 0.011792R_t^2 + 0.98223\sigma^2_{PF,t} \tag{10}$$

Before this model could be utilised, an initial $\sigma^2_{PF,1}$ value for the variance of the portfolio's returns was required. This value was necessary for the calculation of all the future variance values, which would be used to standardise the pseudo daily log returns of the portfolio. In order to calculate $\sigma^2_{PF,t}$, $\sigma^2_{PF,1}$ must be treated as a fixed, known value. To estimate a realistic approximation of $\sigma^2_{PF,1}$, the simple average of the first 1250 squared pseudo daily log return observations was used.

Namely

$$\sigma^2_{PF,1} = \frac{1}{n}\sum_{\tau=1}^{n} R_{t+1-\tau}, \text{where } n = 1250 \tag{11}$$

This yields the result $\sigma^2_{PF,1} = 0.00006964568$.

There are numerous methods with which to compute this value, but for simplicity and ease of calculation this approach was chosen. By using 1250 days of information, the calculated value is based upon the same information used to fit the GARCH(1,1) model, for continuity.

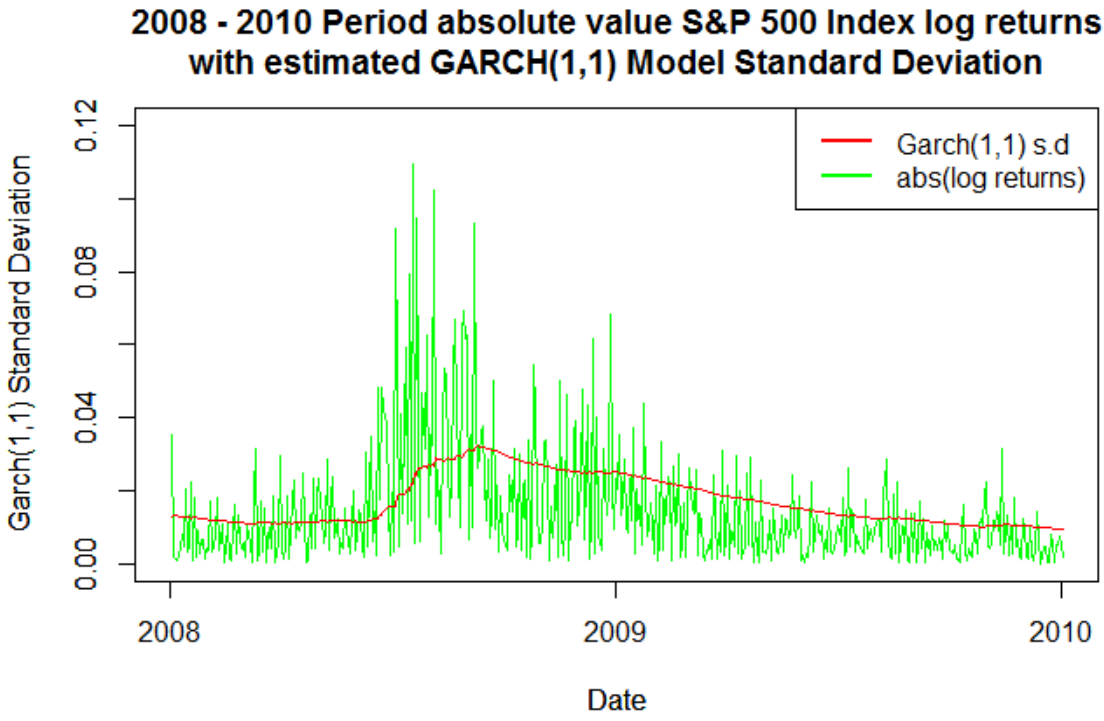

Figure 6: The period 2008 - 2010 is shown here to highlight the ability of the GARCH(1,1) model to respond to fluctuating levels of volatility.

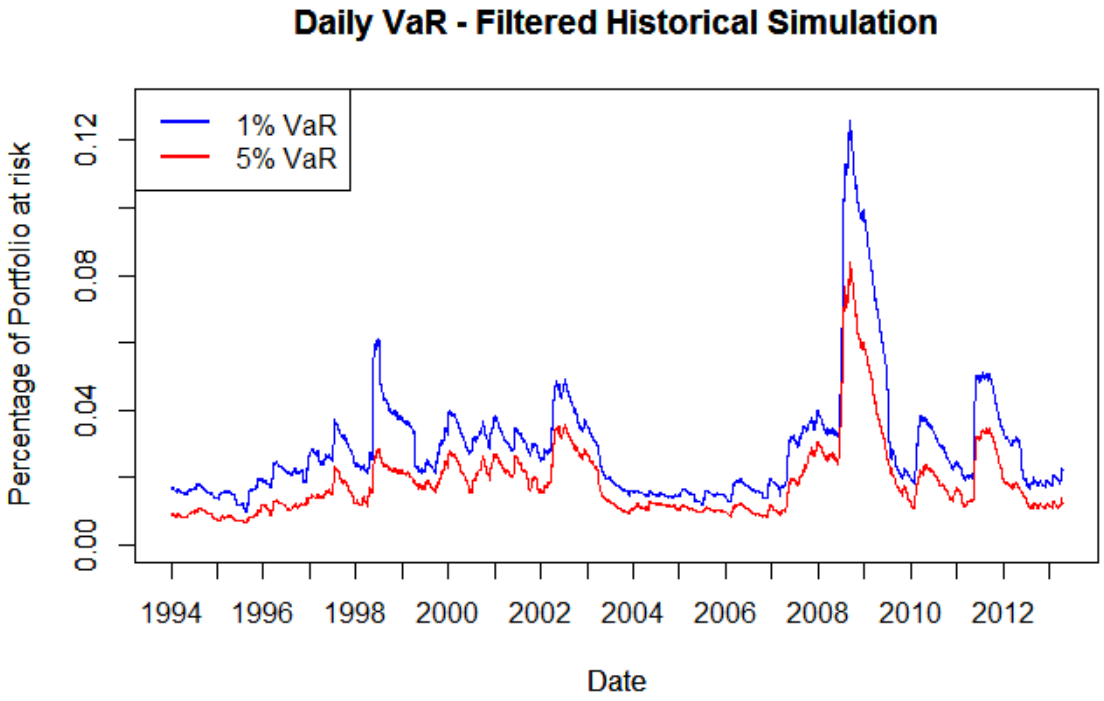

Figure 7: Daily FHS VaR June 1994-June 2013 using $m = 250$. Here the blue line corresponds to a VaR for $p = 1\%$ and the red line to a VaR for $p = 5\%$.

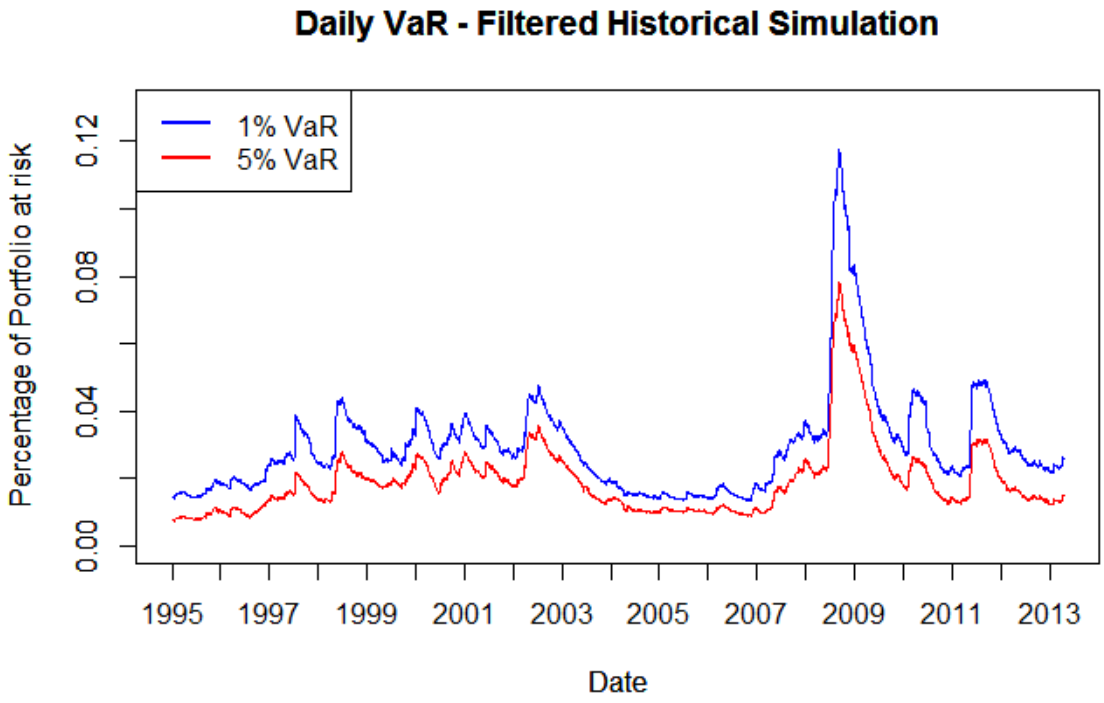

Figure 8: Daily FHS VaR June 1995-June 2013 using $m = 500$. Here the blue line corresponds to a VaR for $p = 1\%$ and the red line to a VaR for $p = 5\%$.

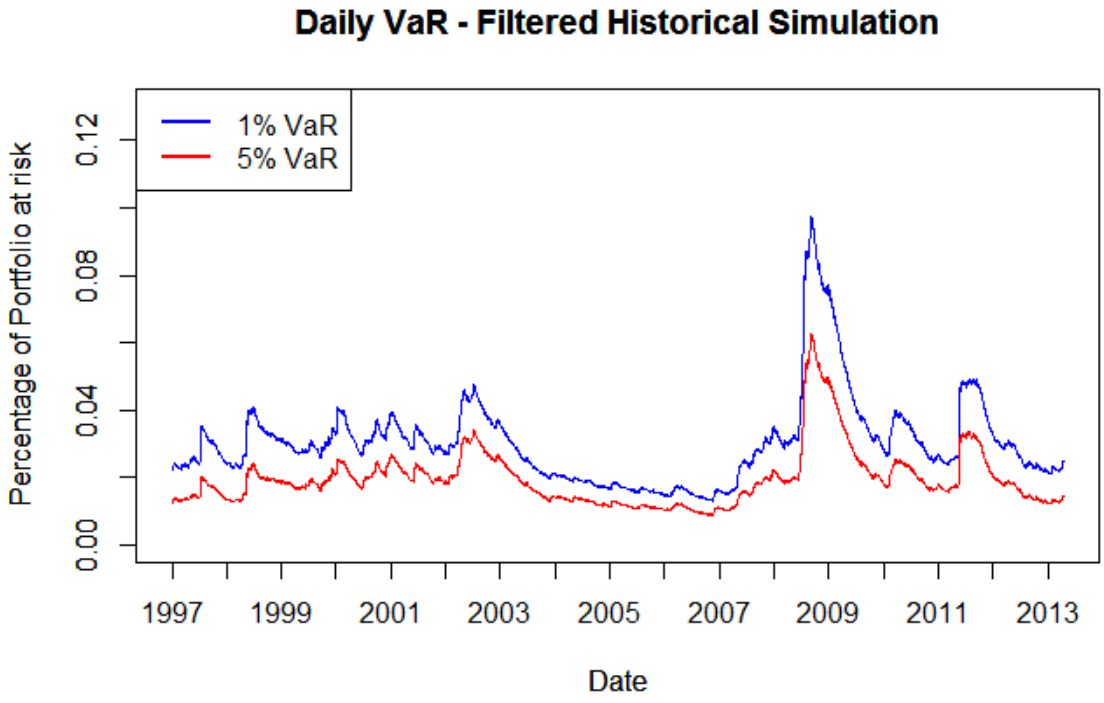

Figure 9: Daily FHS VaR June 1997-June 2013 using $m = 1000$. Here the blue line corresponds to a VaR for $p = 1\%$ and the red line to a VaR for $p = 5\%$.

3.4 Filtered Historical Simulation Results

FHS consistently responds well to volatility fluctuations, predicting much higher VaR levels during periods of high volatility than during periods of normal market volatility. Increasing m, the number of days of historical information used to calculate VaR, leads to lower VaR values for periods of high volatility. However, increasing m has little impact on VaR calculations during periods of low volatility, such as 2003-2007.

Regardless of the choice of m, the FHS method does not produce periods of constant VaR values. This is due in part to the incorporation of dynamic variance, estimated through a GARCH(1,1) model. As a result the accuracy of the VaR estimates can be considered to increase due to the filtering effect of standardising the asset returns by the estimated volatility.

However, increasing the value of m does lead to lower VaR estimates for periods of high volatility. Therefore, we may consider the possibility that the FHS model over-estimates VaR values when a low m value is used.

The FHS approach shares some flaws with its HS counterpart. Just as in the HS method for finding VaR, the FHS method utilises arbitrary m and p values. However, as mentioned above, the introduction of a conditional model in the form of a GARCH(1,1) model reduces the impact of changing the m value. Therefore, while the correct choice of m is still at the individual user's discretion, it is of less import in comparison to the choice of m for Historical Simulation.

It should be noted that the GARCH(1,1) parameters were estimated using S&P 500 Index data from 1988 - 1993. These parameter values were not re-estimated at any future date. In practice, it would be more likely that Filtered Historical Simulation VaR calculations would work on a rolling basis, with the GARCH(1,1) parameters re-estimated every trading day. That is, under the assumption that five years of data is required for statistically significant parameter values to be estimated, the VaR on January 1 1994 would require a GARCH(1,1) model fitted using information from January 1 1988 - December 31 1993, while the VaR on January 2 1994 would require a new GARCH(1,1) model fitted using information from January 2 1988 - January 1 1994.

Therefore, the results displayed here are considered to give a rough indication of how Filtered Historical Simulation would perform in practice.

Hence we could consider that, if the GARCH(1,1) model's parameters had been estimated using data from a period of time including the GFC, then the dynamic variance estimates would be much larger than those reported here. As a result, practitioners currently utilising Filtered Historical Simulation with such a GARCH(1,1) model may in fact be producing over-estimates of future dynamic variance. Therefore it is clear that the range of data used to fit a GARCH(1,1) model's parameters for Filtered Historical Simulation is of great importance, and should be carefully considered.

It is notable that even using a simplified version of the Filtered Historical Simulation method, the VaR values estimated consistently outperform Historical Simulation method VaR values which were calculated using the same m and p values.

4 Calculating Value at Risk using a GARCH(1,1) model combined with the standardised t distribution

As shown in Figure 2, asset returns are not normally distributed. GARCH(1,1) models can in part account for this nonnormality by incorporating distributions which allow for nonnormal shocks. However, GARCH models with normal shocks, such as the GARCH model estimated for the Filtered Historical Simulation VaR method in section 3.2 'by definition do not capture what we call conditional nonnormality in the returns' (Christoffersen 2012, p.121).

Therefore, to address this issue, a GARCH(1,1) model can be combined with a standardised t distribution. Using an exact distribution such as the standardised t distribution 'allows us to compute VaR...for extreme probabilities' (Christoffersen 2012, p. 128).

Student's t distribution has fatter tails and a more pronounced peak than a normal distribution (Christoffersen 2012). Hence this method for finding VaR should be able to account for large potential losses as the returns will not be assumed to be conditionally normal.

Using the structure and terminology of Christoffersen (2012), the Student's t Distribution is defined by

$$f_{t(d)}(x;d) = \frac{\Gamma((d+1)/2)}{\Gamma(d/2)\sqrt{(d\pi)}}(1+x^2/d)^{-(1+d)/2} \text{ for } d>0 \tag{12}$$

Here d is the degrees of freedom of the Student's t distribution. This value has an important relationship with *kurtosis*. Kurtosis, $K(x)$, is the fourth moment of a continuous random variable X, and 'the quantity $K(x)-3$ is called the *excess kurtosis* because $K(x)=3$ for a normal distribution' (Tsay 2005, p. 9).

Therefore, if a distribution has a positive excess kurtosis, it will have fatter tails, as is the case for the standardised t distribution (Tsay 2005). Christoffersen (2012) notes that in order for the kurtosis of the standardised t distribution to be well defined, the d value must be higher than 4.

The following notation and equations are based on the primary reference source, Christoffersen (2012). Many of these definitions are vital for the development of this method, and as such cannot be changed significantly.

Let $\tilde{t}(d)$ be the standardised t distribution, which we derive from the Student's t distribution in order to obtain a distribution with a variance of 1.

Once we have estimated a value for d, we may calculate the portfolio return as

$$R_{PF,t+1} = \sigma_{PF,t+1} z_{t+1}, \text{with } z_{t+1} \overset{i.i.d.}{\sim} \tilde{t}(d) \tag{13}$$

Hence we may compute the Value at Risk as

$$VaR^p_{t+1} = -\sigma_{PF,t+1} \tilde{t}^{-1}_p(d) \tag{14}$$

Here $\tilde{t}^{-1}_p(d)$ is the pth quantile of the $\tilde{t}(d)$ distribution.
Therefore

$$VaR^p_{t+1} = -\sigma_{PF,t+1} \sqrt{\frac{d-2}{d}} t^{-1}_p(d) \tag{15}$$

4.1 Calculating Daily Value at Risk using a GARCH(1,1) model combined with the standardised *t* distribution applied to S&P 500 Index Data

As in section 3.3, the α, β and ω values of the GARCH(1,1) model were estimated by using the fGarch package in the statistical software package R. The degrees of freedom of the Student's t distribution, d, were also estimated within the same calculation.

A two-step approach to fitting a standardised t distribution with a GARCH(1,1) model consists of first calculating the GARCH(1,1) parameters and then inferring the degrees of freedom from the resultant dynamic variance model. One potential problem with this approach is that if the parameter estimates are poor, the following estimate for d will also be poor, leading to an inaccurate model.

Therefore it was deemed more appropriate to estimate all the unknown values simultaneously. Both the GARCH(1,1) parameters and the degrees of freedom for the Student's t distribution were re-estimated each year, using the previous 5 years' historical returns.

The resultant values are listed in Table 2 below. The initial variance, $\sigma^2_{PF,1}$ was taken to be the same as calculated in section 3.3. An asterisk next to the value means that the estimate was statistically significant at the 5% level of significance.

Period Used	Estimated Parameter Value			
	ω	α	β	d
1988-1993	0.00000010360	0.016484*	0.98168*	5*
1989-1994	0.00000013070	0.020680*	0.97716*	5*
1990-1995	0.00000013220	0.027466*	0.96983*	6*
1991-1996	0.00000069999	0.027765*	0.95465*	5*
1992-1997	0.00000072087	0.042335*	0.94151*	6*
1993-1998	0.00000058714*	0.050524*	0.93957*	6*
1994-1999	0.00000065005	0.069698*	0.92660*	6*
1995-2000	0.00000082679	0.065580*	0.93120*	6*
1996-2001	0.0000029901*	0.085199*	0.89908*	7*
1997-2002	0.0000093542*	0.079034*	0.862821*	8*
1998-2003	0.0000081416*	0.083565*	0.87644*	5*
1999-2004	0.0000025766	0.073737*	0.91341*	5*
2000-2005	0.0000011317	0.070295*	0.92343*	5*
2001-2006	0.00000044432	0.050510*	0.94627*	5*
2002-2007	0.00000042791	0.050997*	0.94433*	5*
2003-2008	0.00000052852	0.052370*	0.94178*	9*
2004-2009	0.00000065275	0.081150*	0.91873*	7*
2005-2010	0.00000074699	0.088280*	0.91120*	6*
2006-2011	0.00000086080*	0.10728*	0.90300*	5*
2007-2012	0.0000022184*	0.12220*	0.88009*	5*
2008-2013	0.0000018524*	0.11087*	0.88843*	5*

Table 2: Estimated Parameter values for GARCH(1,1) model and Student's t Distribution

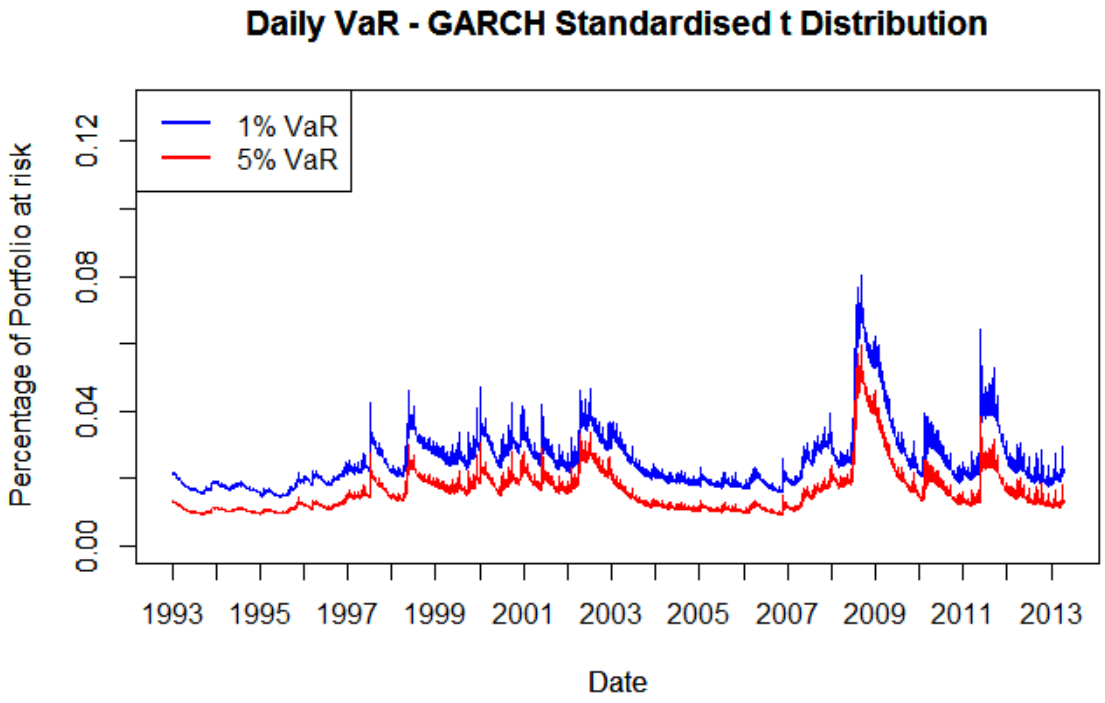

Figure 10: Daily VaR June 1993-June 2013 using a GARCH(1,1) model with nonnormal shocks. The blue line denotes VaR for $p = 1\%$ and the red line VaR for $p = 5\%$.

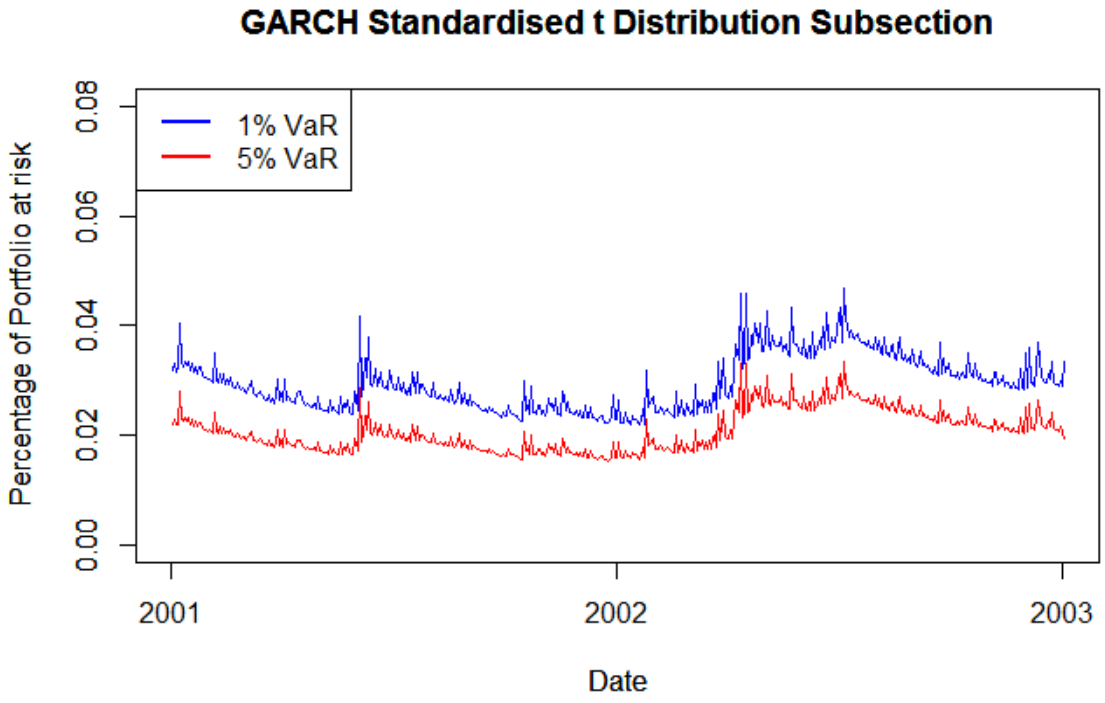

Figure 11: Daily VaR 2001-2003 using a GARCH(1,1) model with nonnormal shocks. Here the blue line denotes VaR for $p = 1\%$ and the red line VaR for $p = 5\%$.

4.2 GARCH(1,1) model combined with the standardised t distribution Results

Computing Value at Risk by using a GARCH(1,1) model combined with a standardised t distribution results in VaR estimates which account for fluctuating levels of volatility. This model may be considered to be theoretically robust, as the combination of a nonnormal distribution of z_t with a GARCH(1,1) model allows us to take into consideration the fatter tails of an asset return distribution.

The model reacts quickly to shocks, and is able to incorporate daily changes in volatility, producing a jagged plot with no areas of constant VaR estimation. The model is capable of distinguishing between minor short term changes in the market, which briefly elicit higher VaR estimates, and more important periods of high volatility like the GFC, which produce extended periods of higher VaR estimation.

Figure 11 highlights the GARCH(1,1) model's ability to produce distinct daily VaR estimates, which is in part due to its combination with a nonnormal distribution, allowing it to account for nonnormal shocks. As mentioned in section 4.1, the GARCH(1,1) parameters and the degrees of freedom for the Student's t distribution were re-estimated each year. This allows the GARCH(1,1) model to remain relevant, while the 5 year period of returns assessed ensures that the integrity of the resultant parameter estimates is high.

However, in practice, it would be more likely that such VaR calculations would work on a rolling basis, with the parameters re-estimated every trading day. Therefore, the results of this model are considered to provide a rough indication of the actual application in a business environment of fitting a GARCH(1,1) model with nonnormal shocks for the purposes of VaR computation.

All of the α and β estimates for the GARCH(1,1) model were statistically significant at the 5% level of significance. However, two thirds of the ω estimates were not. Despite the statistical insignificance of most of the ω estimates, their low values meant that their overall impact on the resultant variance values was minor. Notably, the α and β values estimated using the periods 2006-2011 and 2007-2012 are such that $\alpha + \beta > 1$, which is in violation of the concept of persistence, as covered in equation 6. This may be due to the GFC period and the corresponding extreme spikes in volatility. Theoretically, these parameter estimates suggest that the shocks to the variance are of such magnitude that the variance will never revert to its long run average. As a result, the VaR values estimated for the period 2012-2013 may be unreliable.

The amount of data required to estimate the GARCH(1,1) parameters for this method is arbitrary, as is the choice of p. However, utilising statistical significance as an assessment criterion allows us to accurately assess the validity of parameter estimates from a statistical viewpoint. Having a model with statistically significant parameters lends credence to any further estimates calculated by the model. Therefore, while several years of returns data may be required for such a model, the resultant VaR estimates can be considered more sophisticated and technically accurate than those estimated by a simpler model.

5 Comparison between the three assessed methods for finding Daily Value at Risk

5.1 Comparison between Historical Simulation and Filtered Historical Simulation Results

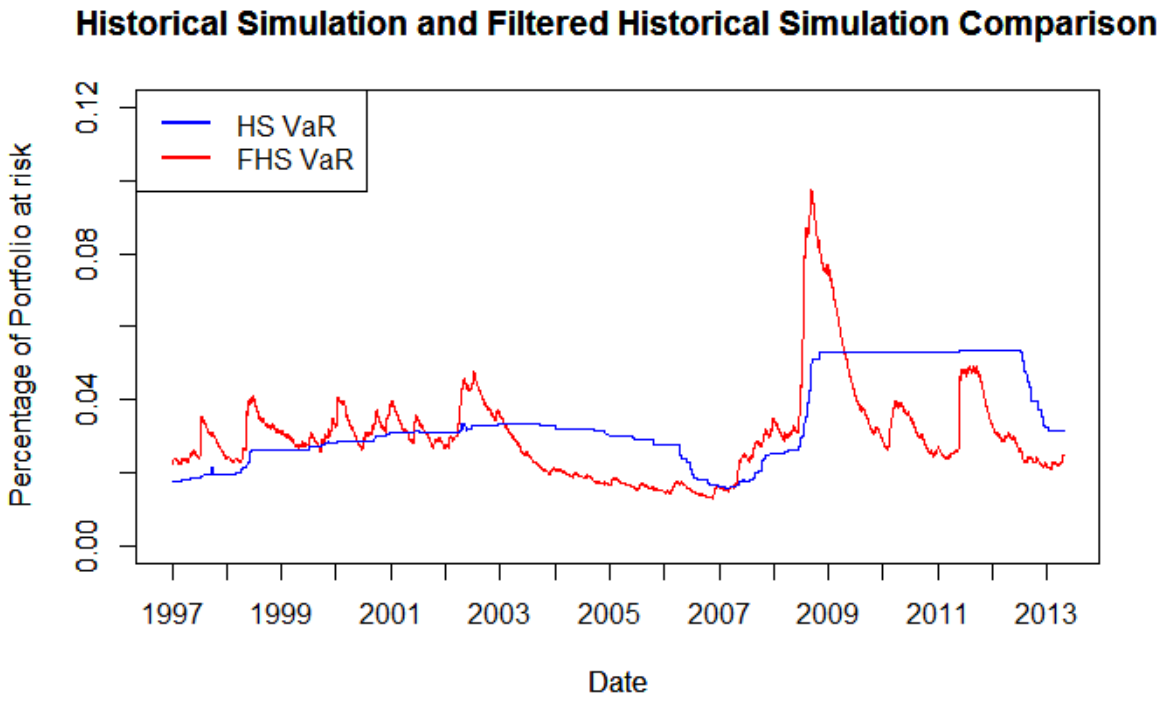

Figure 12: Comparison plot of HS and FHS 1% daily VaR with $m = 1000$.

Overall, the Filtered Historical Simulation method for finding VaR is clearly superior to the Historical Simulation method. The FHS method is more responsive to shocks, and is able to accommodate new market information more quickly than HS. For the initial GFC period (2008-2010), the FHS model predicts a VaR value almost double that of the HS VaR counterpart, showing that the method is more capable of accounting for volatility clustering.

Following periods of volatility clustering, HS will continue to predict high VaR values which, depending on the choice of m, will remain constant for extended periods of time. In contrast, due to the incorporation of the GARCH(1,1) model and its dynamic variance estimation, the FHS method does not display periods of constant VaR estimation.

HS is a simpler method which requires less historical information in order to be calculated. For figure 12, the HS VaR values required only the past 1000 days of historical returns. In contrast, in order to utilise a GARCH(1,1) model with statistically

significant parameter values, the FHS VaR values required the past 2250 days of historical returns. Hence it is clear that either accuracy or speed must be compromised when choosing between HS and FHS.

Another aspect to consider is that in emerging markets where there is only a few years of historical returns data available, FHS may be unuseable. However, in established markets, with decades of historical returns data, it would appear to be inappropriate to use HS when a better alternative in the form of FHS exists.

5.2 Comparison between Historical Simulation and a GARCH(1,1) model combined with a standardised t distribution

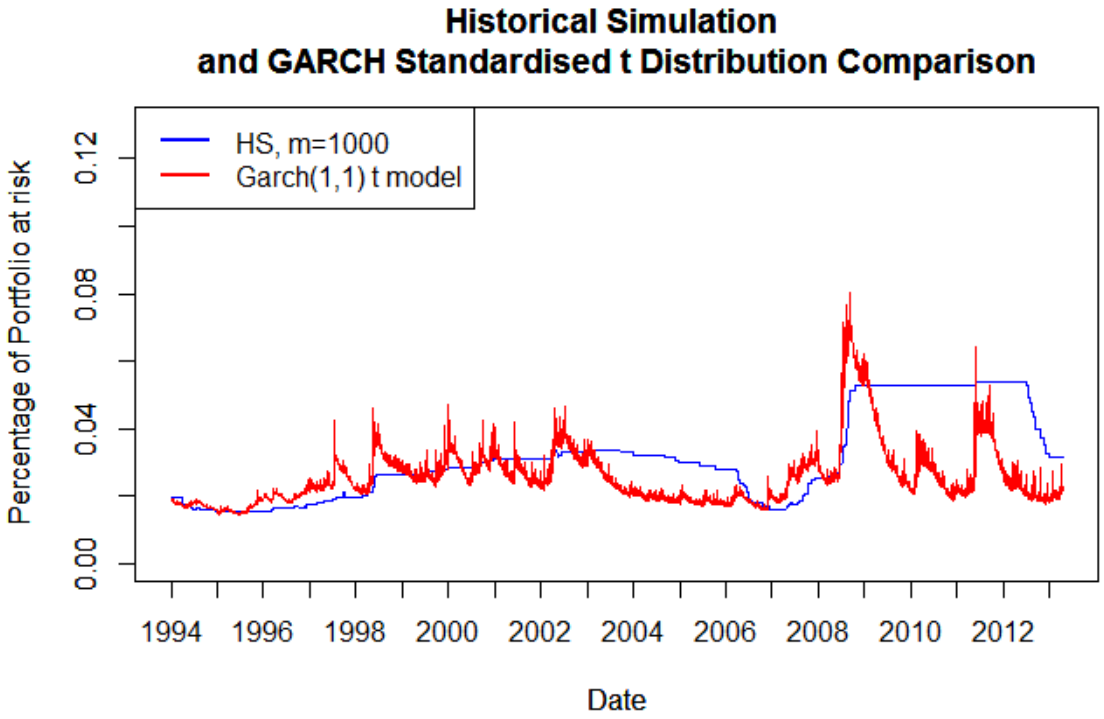

Figure 13: Comparison of Daily 1% VaR June 1994 - June 2013 between HS (in blue) with $m = 1000$, and a GARCH(1,1) model with non-normal shocks (in red).

The method for finding VaR by utilising a GARCH(1,1) model combined with a standardised t distribution is vastly superior to the Historical Simulation method for finding VaR. Not only does the GARCH(1,1) model with nonnormal shocks more accurately account for the nonnormal distribution of asset returns, and hence more accurately predict VaR, it also responds far more quickly to fluctuating levels of volatility. This may be due to the inclusion of dynamic volatility, which is estimated using five years of historical data. As a result the VaR estimates are sophisticated and responsive to volatility clustering.

In comparison Historical Simulation is a simplistic and ponderous method, estimating entire years of constant VaR. As noted in section 2.2, the choice of m for HS results in either potential overestimation or underestimation of VaR following periods of extreme volatility levels, such as the GFC.

Historical Simulation does have the advantage of being simple and easy to calculate. However, if the resources are available to conduct VaR estimation by using a GARCH(1,1) model which accounts for nonnormality of returns, then Historical Simulation should not be considered as a viable method for computing VaR.

5.3 Comparison between Filtered Historical Simulation and a GARCH(1,1) model combined with a standardised t distribution

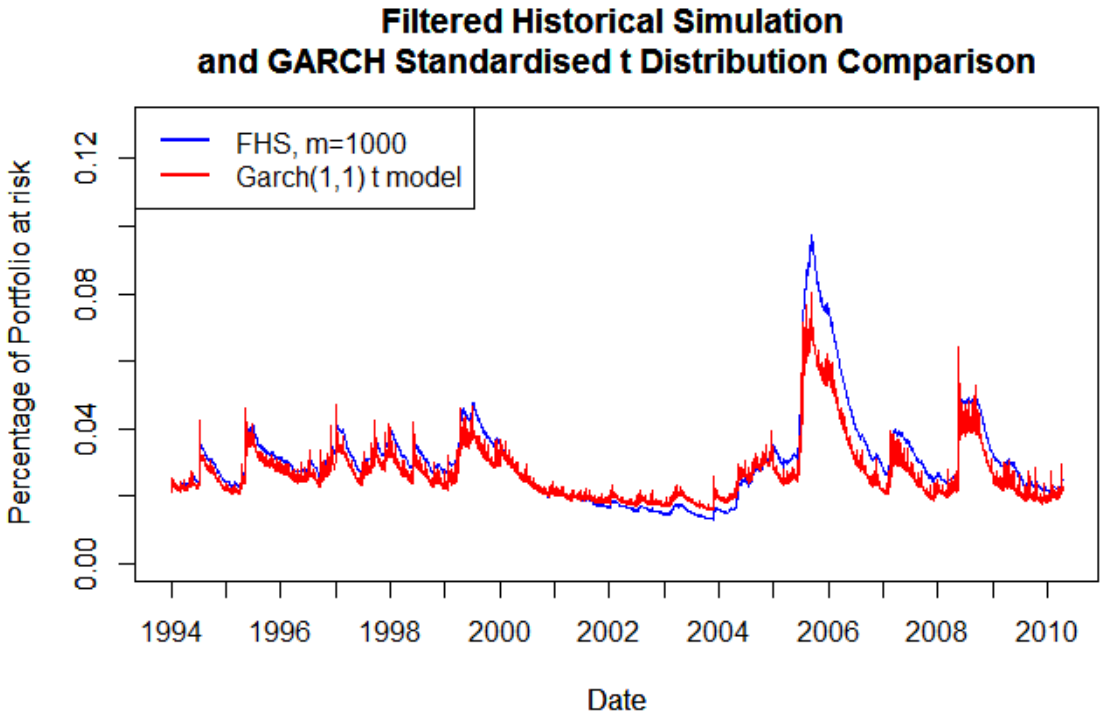

Figure 14: Comparison of Daily 1% VaR June 1994 - June 2013 between FHS (in blue) with $m = 1000$, and a GARCH(1,1) model with non-normal shocks (in red).

Here we can see that the Filtered Historical Simulation method and the method of fitting a GARCH(1,1) model with a standardised t distribution result in similar VaR calculations. The GARCH(1,1) model with nonnormal shocks responds quickly to daily changes in volatility, producing a more jagged plot with faster fluctuations than the FHS model. However, FHS predicts slightly higher VaR values for the GFC period, where there is a large amount of volatility clustering.

The common factor in both these methods is the inclusion of a GARCH(1,1) model to estimate dynamic volatility. While combining the GARCH(1,1) model with a standardised t distribution allows us to tailor the model more specifically to asset returns, the resultant effect on VaR estimates is small. Although the daily VaR estimates become more responsive, from a long term perspective figure 14 above demonstrates that the two methods' VaR estimates remain similar.

Therefore, utilising a GARCH(1,1) model within VaR calculations is a clear improvement over the more simplistic Historical Simulation method for finding VaR.

It should be noted that these methods are not directly comparable as the FHS model relies on a GARCH(1,1) with parameters estimated using only one period of

historical returns data, whereas the parameter values of the GARCH(1,1) combined with the standardised t distribution were annually re-estimated. As a result, the latter FHS VaR values may not be as accurate.

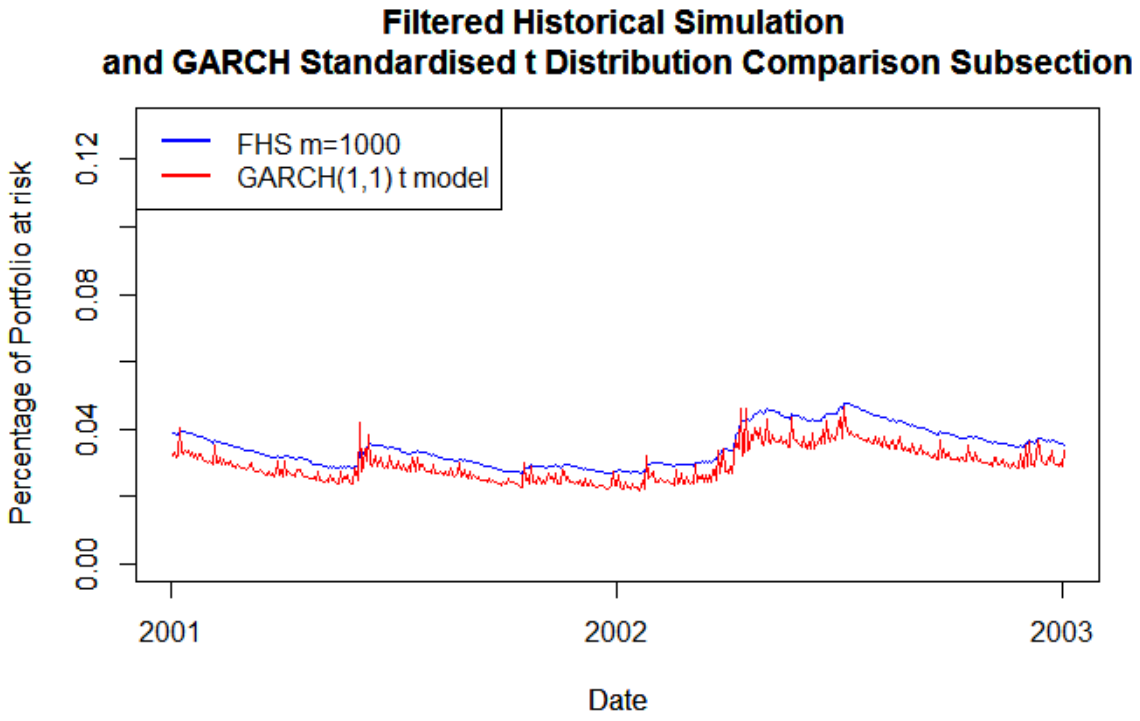

Figure 15: Comparison of Daily 1% VaR January 2001- January 2003 between Filtered Historical Simulation (in blue) with $m = 1000$, and a GARCH(1,1) model combined with non-normal shocks (in red).

From the subsection plot above, we can clearly see that the GARCH(1,1) model with nonnormal shocks is more responsive to daily fluctuations in volatility. This result is expected, as by incorporating a nonnormal distribution in the form of the standardised Student's t distribution, the model should more accurately follow the distribution of returns. This method appears to respond well to shocks, with immediate changes in the VaR value estimated. These responses may be more modest than for the FHS method, but this could suggest that the FHS method is overestimating the VaR.

The Filtered Historical Simulation method, which incorporates a GARCH(1,1) model with normal shocks results in a considerably smoother graph of VaR. This could be a result of the FHS method using a sequence of m past returns in the calculation of a VaR value. As these returns are sorted into ascending order and the $p \cdot 100th$ percentile selected, the return selected may be the same value over a short period of time. While the VaR values calculated over such a period of time may differ due to different levels of estimated volatility, these VaR values would be similar.

In comparison to the GARCH(1,1) model with nonnormal shocks, the FHS plot often results in higher VaR percentages. This may be due to the fact that the parameters of the FHS Garch(1,1) model were estimated using returns from the period 1988-1993. These parameters were not recalculated each year, in contrast to the GARCH(1,1) model with nonnormal shocks. Hence the resultant FHS VaR values for latter periods such as the 2001-2003 period shown above may be less accurate than VaR values calculated for earlier periods.

6 Conclusion

This project assessed three distinct methods for finding the financial risk measure Value at Risk. These were Historical Simulation, Filtered Historical Simulation, and fitting a GARCH(1,1) model combined with a standardised t distribution to a data set.

Of the methods assessed, Historical Simulation was clearly the worst performing method for calculating VaR. HS responds poorly to volatility fluctuations, and cannot easily accommodate new information following a period of high volatility clustering. HS uses the arbitrary value m to determine the range of historical information to be assessed. Increasing the value of m does not improve the performance of the method, but rather reduces it, an outcome which is counter-intuitive and highlights the inadequacies of the method. Ultimately HS should not be used if an alternative method for finding VaR, such as the two other methods assessed in this paper, can be utilised.

Filtered Historical Simulation outperforms HS as it is more responsive to volatility clustering. As the FHS method incorporates dynamic volatility, it can react more appropriately to fluctuating levels of volatility than HS, and more accurately match financial markets' reactions. FHS is however more dificult to calculate than HS, and requires a larger amount of past information.

Using a GARCH(1,1) model combined with a standardised t distribution also outperforms Historical Simulation as a method for finding VaR. The GARCH(1,1) model with a nonnormal distribution responds far more quickly to daily volatility fluctuations than HS, and can account for the nonnormality of asset returns. It requires more complex calculations than HS, but this increase in complexity is matched by an increase in the sophistication of VaR estimates.

In comparison to a GARCH(1,1) model with nonnormal shocks, the Filtered Historical Simulation method reacts less quickly to daily changes in volatility, producing a much smoother plot of daily VaR. However, during periods of high volatility, the FHS method actually appears to respond to a greater degree, resulting in higher VaR calculations. FHS does not account for the nonnormality of asset returns as well as a GARCH(1,1) model with nonnormal shocks, and as such could be considered to produce less accurate VaR estimates. Ultimately it is unclear which method is superior, as they both have favourable results, although the GARCH(1,1) model incorporating nonnormal shocks may be more theoretically accurate.

All of these methods do however share a common flaw. By definition, Value at Risk yields a figure which is the minimum loss expected for a given day, with a specific probability. However it gives no insight into the distribution of the losses within that probability coverage. Therefore, while the VaR on a given day may be reported to be at least $V\%$ of a portfolio P, the actual value of the portfolio at risk could be much larger than $V\%$. None of the methods for finding VaR assessed in this project address this issue. Should further work be undertaken on this subject, a more advanced method capable of conveying this information, known as Expected Shortfall, would be utilised (Christoffersen 2012).

7 Bibliography

Christoffersen, PF 2012, *Elements of Financial Risk Management, 2nd edn*, Academic Press, Massachusetts.

Ergen, I 2010, *VaR Prediction for Emerging Stock Markets: GARCH Filtered Skewed t Distribution and GARCH Filtered EVT Method*, Supervision Regulation and Credit, Policy Analysis Unit, Federal Reserve Bank of Richmond, Baltimore.

Hull, JC 2011, *Fundamentals of Futures and Options Markets, 7th edn*, Pearson Education Limited, Essex.

Kabaila, P 1999, 'The Relevance Property for Prediction Intervals', *Journal of Time Series Analysis*, vol. 20, no. 6, pp. 655-662.

Mandelbrot, BB 1963, 'The variation of certain speculative prices', *Journal of Business*, vol. 36, no. 4 , pp. 392-419.

Roy, I 2011, 'Estimating Value at Risk (VaR) using Filtered Historical Simulation in the Indian Capital market', Reserve Bank of India Occasional Papers, vol. 32, no. 2, pp. 81-98.

S&P Dow Jones Indices McGraw Hill Financial 2013, *S&P 500*, S&P Dow Jones Indices McGraw Hill Financial, viewed 8 January 2014, <http://au.spindices.com/indices/equity/sp-500>.

Tsay, RS 2005, *Analysis of Financial Time Series, 2nd edn*, John Wiley & Sons, Inc., New Jersey.

8 Appendix - R coding

This Appendix includes the coding used in the software package R for this project. The S&P 500 Index data is from Bloomberg and the Chicago Board Options Exchange (CBOE).

8.1 Historical Simulation Coding

```
## Importing data and manipulating into useable format ##

SP <- read.csv("sp500tir.fixed.dates.csv", header=TRUE)
# import S&P 500 data
head(SP)
plot(SP, main="S&P500 Index Daily Closing Price")
# plots daily closing price over period 1988 to present
Number.of.stocks.held <- 100
# this is an arbitrary value, won't change the return rates
SP["Portfolio.Value"] <-
SP$Closing.Price*Number.of.stocks.held
# make new column for portfolio value

dim(SP) #6324 observations

i <- c(1:6324)
j <- c(NA,i-1)

SP["Pseudo.Return"] <- SP$Portfolio.Value[i]/
SP$Portfolio.Value[j]
# following Christoffersen's equation

SP["Log.Pseudo.Return"] <- log(SP$Pseudo.Return)
# taking the log to yield estimated returns

SP.date.and.lr <-
data.frame(SP$Date,SP$Log.Pseudo.Return)
# create dataframe with date and log return only
head(SP.date.and.lr)

## To plot log returns ##

returns.dates <- c(1988, 1989, 1990, 1991,
1992, 1993, 1994, 1995, 1996, 1997, 1998,
1999, 2000, 2001, 2002, 2003, 2004, 2005,
```

```
2006, 2007, 2008, 2009, 2010, 2011, 2012,
          2013)
returns.numbers <- c(0, 250, 500, 750, 1000,
1250, 1500, 1750, 2000, 2250, 2500, 2750,
3000, 3250, 3500, 3750, 4000, 4250, 4500,
4750, 5000, 5250, 5500,  5750, 6000, 6250)

plot(SP$Log.Pseudo.Return, type='l', col='blue',
     main='Daily S&P 500 Index Log Returns',
ylab='Daily Log Return', xlab='date',
xaxt = "n")
par(new=TRUE)
plot(SP$Log.Pseudo.Return, type='l', col='blue',
     main='Daily S&P 500 Index Log Returns',
 ylab='Daily Log Return', xlab='date',
axis(1,
at=c(returns.numbers),
       labels=c(returns.dates)) )

## Finding Historical Simulation VaR ##

library(PerformanceAnalytics)
# need this to compute historical simulation VaR

## HS with m=250 ##

m <- 250
x <- 0.01
z <- 0.05

daily.VaR <- numeric(6324-m)

for (t in 1:(6324-m)) {
  daily.VaR[t] <- -VaR(SP.date.and.lr[t:(m-1+t),2],
p=1-x, method="historical")
# negative for graph purposes
} # for loop to calculate daily VaR using
# previous m days data,
#from day m+1 onwards
head(daily.VaR)
# So daily.VaR[1] corresponds to the
# daily VaR for the 251st day,
# ie 1989-05-26
# and the daily.VaR[6074] corresponds to the
```

```
# daily VaR for the
# 6324th or final day

daily.VaRz <- numeric(6324-m)

for (t in 1:(6324-m)) {
  daily.VaRz[t] <- -VaR(SP.date.and.lr[t:(m-1+t),2],
p=1-z, method="historical")
  # negative for graph purposes
} # for loop to calculate daily VaR using
# previous m days data, from day
#  m+1 onwards with different probability

# to plot the graph#

SP.date.m.plus.one.onwards <- SP$Date[(m+1):6324]
SP.date.m.plus.one.onwards[1]
SP.date.m.plus.one.onwards[1000]
# To check  figures are correct
# test 2000, 3000 etc

dates <- c(1989, 1990, 1991, 1992, 1993, 1994,
1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002,
2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010,
2011, 2012, 2013)
numbers <- c(0, 250, 500, 750, 1000, 1250, 1500,
1750, 2000, 2250, 2500, 2750, 3000, 3250, 3500,
3750, 4000, 4250, 4500, 4750, 5000, 5250, 5500,
5750, 6000)


plot(daily.VaR, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="blue", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n"
)

par(new=TRUE)
plot(daily.VaRz, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n")
```

```
# don't worry this is not a typo, due to the coding
# the graph needs to be plotted again in order to
# include the dates
plot(daily.VaRz, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n"
     , axis(1, at=c(numbers),
       labels=c(dates)) )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

## HS with m=500 ##

m <- 500
x <- 0.01
z <- 0.05


daily.VaR.500 <- numeric(6324-m)

for (t in 1:(6324-m)) {
  daily.VaR.500[t] <- -VaR(SP.date.and.lr[t:(m-1+t),2],
 p=1-x, method="historical")
  # negative for graph purposes
} # for loop to calculate daily VaR using
# previous m days data,
# from day m+1 onwards
head(daily.VaR.500)

daily.VaRz.500 <- numeric(6324-m)

for (t in 1:(6324-m)) {
  daily.VaRz.500[t] <- -VaR(SP.date.and.lr[t:(m-1+t),2],
p=1-z,
method="historical")
  # negative for graph purposes
} # for loop to calculate daily VaR using
# previous m days data,
#from day m+1 onwards
```

```
# to plot the graph #

SP.date.m.plus.one.onwards <- SP$Date[(m+1):6324]
500.dates <- c(1990, 1991, 1992, 1993, 1994, 1995,
1996, 1997, 1998, 1999, 2000, 2001, 2002, 2003,
2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011,
2012, 2013)
500.numbers <- c(0, 250, 500, 750, 1000, 1250, 1500,
1750, 2000, 2250, 2500, 2750, 3000, 3250, 3500,
3750, 4000, 4250, 4500, 4750, 5000, 5250, 5500, 5750)

plot(daily.VaR.500, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="blue", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n"
)

par(new=TRUE) # plot on same graph as before

plot(daily.VaRz.500, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n")
# don't worry this is not a typo, due to the coding
# the graph must be plotted again in order to include
# correct dates
plot(daily.VaRz.500, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n"
     , axis(1, at=c(500.numbers),
       labels=c(500.dates))
)
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

## HS with m=1000 ##
```

```
m <- 1000
x <- 0.01
z <- 0.05

daily.VaR.1000 <- numeric(6324-m)

for (t in 1:(6324-m)) {
  daily.VaR.1000[t] <- -VaR(SP.date.and.lr[t:(m-1+t),2],
p=1-x, method="historical")
# negative for graph purposes
} # for loop to calculate daily VaR using previous
# m days data, from day m+1 onwards
head(daily.VaR.1000)

daily.VaRz.1000 <- numeric(6324-m)

for (t in 1:(6324-m)) {
  daily.VaRz.1000[t] <- -VaR(SP.date.and.lr[t:(m-1+t),2], p=1-z,
method="historical")
  # negative for graph purposes
} # for loop to calculate daily VaR using previous m days data,
# from day m+1 onwards

# to plot the graph #

SP.date.m.plus.one.onwards <- SP$Date[(m+1):6324]
1000.dates <- c(1992, 1993, 1994, 1995, 1996, 1997,
1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005,
2006, 2007,  2008, 2009, 2010, 2011, 2012, 2013)
1000.numbers <- c(0, 250, 500, 750, 1000, 1250, 1500,
1750, 2000, 2250, 2500, 2750, 3000, 3250, 3500,
3750, 4000, 4250, 4500, 4750, 5000, 5250)

plot(daily.VaR.1000, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
col="blue",
     cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n"
    )

par(new=TRUE)

plot(daily.VaRz.1000, type = 'l', xlab = 'Date',
```

```
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n")
# don't worry this is not a typo, due to the coding
# the graph must be plotted again in order to include
# correct dates
plot(daily.VaRz.1000, type = 'l', xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.09)
     , xaxt = "n"
     , axis(1, at=c(1000.numbers),
            labels=c(1000.dates))
     )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

## to draw the histogram ##

daily.VaRz[3000] # 0.02169296
# this is for the 3000th observation,
# which would be the 4000th day after 1988,
# ie 16 years later, 2004
# So the VaR for this date would be calculated
# using the previous
# 1000 days pseudo daily returns,
# hence we have the histogram
# with returns from 2000-2004

SP.date.m.plus.one.onwards[3000] # 2004

L <- SP.date.and.lr[3000:4000,2]
#this is between 2000 and 2004

SP.date.and.lr[3000,1]

h<-hist(L, breaks=20, col="green",
        xlab=
"Pseudo Daily Return S&P 500 Index 2000 - 2004",
        main=
"Historical Simulation VaR Concept Example")
```

```
# To plot the histogram of index returns
xfit<-seq(min(L),max(L),length=40)
yfit<-dnorm(xfit,mean=mean(L),sd=sd(L))
yfit <- yfit*diff(h$mids[1:2])*length(L)
lines(xfit, yfit, col="blue", lwd=2)
# To include a normal distribution curve

abline(v=-0.022,col="red", cex=2)
# To include the red 5% line
```

8.2 Filtered Historical Simulation Coding

```
library(fGarch)
# Need this package to compute the Garch(1,1) model

spdata <-na.omit(SP.date.and.lr)
head(spdata)

# the problem with using a model which uses all 6324
# observations of return to estimate the garch(1,1)
# model parameters omega, alpha and beta is that it
# has incorporated future' values into its
# parameter estimation.

# This is incorrect, as I should be using a sequence
# of past returns
# to estimate a GARCH(1,1) model which will give
# insight into future observations.
# Hence, I will use

m = 1250
garchFit(~garch(1,1), data=spdata[(1:m),2])

# R output #

# Title:
#   GARCH Modelling

# Call:
#   garchFit(formula = ~garch(1, 1),
data = spdata[(1:m), 2])
```

```
# Mean and Variance Equation:
#   data ~ garch(1, 1)
# <environment: 0x11ac8ed0>
#   [data = spdata[(1:m), 2]]

#Conditional Distribution:
#  norm

# Coefficient(s):
#   mu       omega      alpha1       beta1
# 5.2581e-04  3.8683e-07  1.1792e-02  9.8223e-01

# Std. Errors:
#   based on Hessian

# Error Analysis:
#   Estimate  Std. Error  t value Pr(>|t|)
# mu     5.258e-04   2.260e-04    2.326   0.0200 *
#   omega  3.868e-07   1.770e-07    2.185   0.0289 *
#   alpha1 1.179e-02   2.625e-03    4.491 7.08e-06 ***
#   beta1  9.822e-01   3.885e-03  252.850  < 2e-16 ***
#   ---
#   Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

# Log Likelihood:
#   4234.546    normalized:  3.387637

# Description:
#   Thu Jan 09 15:20:19 2014 by user: Rupert

# So we have a GARCH(1,1) model where alpha,
# beta and omega are
# all statistically significant
# This suggests the model
# sigma_{pf,t+1}^{2} = omega + alpha*R_{PF,t}^{2} +
# beta*sigma_{pf,t}^{2}
# where the optimal parameters imply
# the following variance dynamics
# so sigma_{pf,t+1}^{2}
# = 3.8683e-07 +1.1792e-02 *R_{PF,t}^{2}
# + 9.8223e-01 * sigma_{pf,t+1}^{2}
# here the persistence of variance is 0.994022
# which is close to 1 therefore good
```

```
head(spdata)
head(SP$Squared.Daily.Log.Return)
length(sp.return.squared)
sp.return.squared <-
na.omit(SP$Squared.Daily.Log.Return)
simple.m.variance <-
(1/m)*sum(sp.return.squared[1:m])
# to calculate initial variance estimate
head(sp.return.squared)

garch.variance <- numeric(6321)
garch.variance[1:(m-1)] <- 0
garch.variance[m] <- simple.m.variance
length(garch.variance)
head(garch.variance)
for (t in (m):(6321)) { #so starts at t=1250
  garch.variance[t+1] <- (3.8683e-07
+ (1.1792e-02*sp.return.squared[t])
+(9.8223e-01*garch.variance[t]) )
}
garch.variance[m]

# To plot graph of garch variance #

gv.dates <- c(2008, 2009, 2010)
gv.numbers <- c(0, 250, 500)

plot(sqrt(sp.return.squared[5000:5500]), type='l',
col="green", ylim=c(0,0.12)
     ,main =
'2008 - 2010 Period absolute value S&P 500
Index log returns
     with estimated GARCH(1,1) Model
Standard Deviation',
     xaxt = 'n',  xlab = 'Date',
ylab = 'Garch(1,1) Standard Deviation')
par(new=TRUE)
plot(sqrt(garch.variance[5000:5500]), type='l',
col="red", ylim=c(0,0.12)
     , xaxt = 'n', xlab= 'Date',
ylab = 'Garch(1,1) Standard Deviation')
par(new=TRUE)
plot(sqrt(garch.variance[5000:5500]),
type='l', col="red",
```

```
ylab='Garch(1,1) Standard Deviation',
     ylim=c(0,0.12)
     , xaxt = "n", xlab = 'Date'
     , axis(1, at=c(gv.numbers),
            labels=c(gv.dates)))
legend("topright",9.5, c("Garch(1,1) s.d",
                         "abs(log returns)")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("red","green"))

# To standardise returns #

zhat <- numeric(6321)
zhat[1:1250] = 0
returns <- na.omit(SP$Log.Pseudo.Return)
length(returns)

for (t in (m+1):(6321)) {
  zhat[t] <- returns[t]/(sqrt(garch.variance[t]))
}
head(returns)
head(SP)
returns[1251]
garch.variance[1251]
sqrt(garch.variance[1251])
returns[1251]/sqrt(garch.variance[1251])
zhat[1251] # just checking coding is working as desired
zhat[6314]
head(zhat)
zhat <- zhat/100
# this is to reduce all observations below 1,
# ie abs.value for all observations <1, so that
# the VaR function will compute meaningful
# values, then I just adjust the result by
# multiplying  by 100 to bring values back
# to original
# simply multiplying by a constant doesn't
# change the order
# of the sequence so
# no information is lost.

library(PerformanceAnalytics)
# Need this package to compute
```

```
# the following calculations

#### TO find VaR with 250 day lag #####
n <- 250
x <- 0.01
z <- 0.05

daily.FHS.VaR.250 <- numeric(6324)
daily.FHS.VaR.250[1:(m+n)] = 0
zhat[1240:1500]

# here we have t from m+n onwards, so that from
# the start the VaR being calculated is using
# hypthetical zhat values which are not zero
for (t in (m+n):(6324)) {
  daily.FHS.VaR.250[t+1] <-
-100*sqrt(garch.variance[t+1])
*VaR(zhat[(t-n+1):(t)],
   p=1-x, method="historical")
  # negative for graph purposes
}

head(daily.FHS.VaR) # 1st 1500 should all equal 0
daily.FHS.VaR.250[1501]

daily.FHS.VaRz.250 <- numeric(6324)
daily.FHS.VaRz.250[1:(m+n)] = 0

for (t in (m+n):(6324)) {
  daily.FHS.VaRz.250[t+1] <-
-100*sqrt(garch.variance[t+1])
*VaR(zhat[(t-n+1):(t)],
     p=1-z, method="historical")
  # negative for graph purposes
}

head(daily.FHS.VaRz.250)
# 1st 1500 should all equal 0
daily.FHS.VaRz.250[1501]

plot(daily.FHS.VaR.250[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main =
```

```
'Daily VaR - Filtered Historical Simulation',
     col="blue", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)

# we start 1500 days in, ie 6 years in, at 1994
fhs.250.dates <- c(1994, 1995, 1996, 1997, 1998,
1999, 2000, 2001, 2002,
 2003, 2004, 2005, 2006, 2007,
2008, 2009, 2010, 2011, 2012, 2013)
fhs.250.numbers <- c(0, 250, 500, 750,
1000, 1250, 1500, 1750, 2000,
 2250, 2500, 2750, 3000,
 3250, 3500, 3750, 4000, 4250, 4500, 4750)

par(new=TRUE)
plot(daily.FHS.VaRz.250[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main =
'Daily VaR - Filtered Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)
par(new=TRUE)
plot(daily.FHS.VaRz.250[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main =
 'Daily VaR - Filtered Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
     , axis(1, at=c(fhs.250.numbers),
            labels=c(fhs.250.dates)) )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

#### To find VaR with 500 day lag #####

n <- 500
x <- 0.01
z <- 0.05
```

```
daily.FHS.VaR.500 <- numeric(6324)
daily.FHS.VaR.500[1:(m+n)] =0
zhat[1240:1500]

# here we have t from m+n onwards, so that
# from the start the VaR being calculated is using
# hypthetical zhat values which are
# not zero/undefined

for (t in (m+n):(6324)) {
  daily.FHS.VaR.500[t+1] <-
-100*sqrt(garch.variance[t+1])
*VaR(zhat[(t-n+1):(t)],
p=1-x, method="historical")
  # negative for graph purposes
}

head(daily.FHS.VaR) # 1st 1250 should all equal 0
daily.FHS.VaR.500[1751]

daily.FHS.VaRz.500 <- numeric(6324)
daily.FHS.VaRz.500[1:(m+n)] =0

for (t in (m+n):(6324)) {
  daily.FHS.VaRz.500[t+1] <-
-100*sqrt(garch.variance[t+1])
*VaR(zhat[(t-n+1):(t)]
 , p=1-z, method="historical")
  # negative for graph purposes
}

head(daily.FHS.VaR) # 1st 1250 should all equal 0
daily.FHS.VaRz.500[1751]

plot(daily.FHS.VaR.500[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main =
 'Daily VaR - Filtered Historical Simulation',
     col="blue", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)
```

```
# starting one year after 1994, ie 1995
fhs.500.dates <- c( 1995, 1996, 1997, 1998,
1999, 2000, 2001, 2002,
 2003, 2004, 2005, 2006,
 2007, 2008, 2009, 2010,
2011, 2012, 2013)
fhs.500.numbers <- c(0, 250, 500, 750, 1000,
1250, 1500, 1750, 2000, 2250, 2500, 2750,
3000, 3250, 3500, 3750, 4000, 4250, 4500)

par(new=TRUE)
plot(daily.FHS.VaRz.500[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Filtered Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)
par(new=TRUE)
plot(daily.FHS.VaRz.500[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main =
 'Daily VaR - Filtered Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
     , axis(1, at=c(fhs.500.numbers),
            labels=c(fhs.500.dates)) )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

#### To find VaR with 1000 day lag #####

n <- 1000
x <- 0.01
z <- 0.05

daily.FHS.VaR.1000 <- numeric(6324)
daily.FHS.VaR.1000[1:(m+n)] =0
zhat[1240:1500]

# here we have t from m+n onwards, so that from
```

```
# the start the VaR being calculated is using
# hypthetical zhat values which are not
# zero/undefined
for (t in (m+n):(6324)) {
  daily.FHS.VaR.1000[t+1] <-
-100*sqrt(garch.variance[t+1])
*VaR(zhat[(t-n+1):(t)],
p=1-x, method="historical")
  # negative for graph purposes
}
head(daily.FHS.VaR) # 1st 1250 should all equal 0
daily.FHS.VaR.1000[2250]

daily.FHS.VaRz.1000 <- numeric(6324)
daily.FHS.VaRz.1000[1:(m+n)] =0

for (t in (m+n):(6324)) {
  daily.FHS.VaRz.1000[t+1] <-
-100*sqrt(garch.variance[t+1])
*VaR(zhat[(t-n+1):(t)],
 p=1-z, method="historical")
  # negative for graph purposes
}
head(daily.FHS.VaR) # 1st 1250 should all equal 0
daily.FHS.VaRz.1000[2251]

plot(daily.FHS.VaR.1000[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Filtered Historical Simulation',
     col="blue", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)

# starting 2 years after 1995, ie 1997
fhs.1000.dates <- c( 1997, 1998, 1999, 2000,
2001, 2002, 2003, 2004, 2005, 2006, 2007,
2008, 2009, 2010, 2011, 2012, 2013)
fhs.1000.numbers <- c(0, 250, 500, 750, 1000,
1250, 1500, 1750, 2000, 2250,
 2500, 2750, 3000, 3250, 3500, 3750, 4000)

par(new=TRUE)
plot(daily.FHS.VaRz.1000[(m+n+1):6324], type = 'l',
```

```
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Filtered Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)
par(new=TRUE)
plot(daily.FHS.VaRz.1000[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - Filtered Historical Simulation',
     col="red", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
     , axis(1, at=c(fhs.1000.numbers),
            labels=c(fhs.1000.dates)) )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

## To plot comparison with FHS ##

plot(daily.FHS.VaR.1000[(m+n+1):6324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Historical Simulation and
Filtered Historical Simulation Comparison',
     col="red", cex=0.75, ylim=c(0, 0.12)
     , xaxt = "n"
)
par(new=TRUE)
plot(daily.VaR[1250:5324], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Historical Simulation and
 Filtered Historical Simulation Comparison',
 col="blue",
     cex=0.75, ylim=c(0, 0.12)
     , xaxt = "n"
)

par(new=TRUE)
plot(daily.VaR[1250:5324], type = 'l',
xlab = 'Date',
```

```
     ylab= 'Percentage of Portfolio at risk',
     main = 'Historical Simulation and
Filtered Historical Simulation Comparison',
     col="blue", cex=0.75,
 ylim=c(0, 0.12)
     , xaxt = "n"
     , axis(1, at=c(fhs.1000.numbers),
           labels=c(fhs.1000.dates))
)
legend("topleft",9.5, c("HS VaR", "FHS VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))
```

8.3 GARCH(1,1) with standardised t Distribution Coding

```
## GARCH(1,1) with standardised t distribution ##
# please note, the GarchFit command used here will
# result in a printout of information as listed in
# the filtered historical section
# For brevity the printouts obtained below are not
# included, as they would needlessly extend
# this document and are not significantly necessary
# for comprehension. The estimated GARCH(1,1)
# parameters can be seen from the test.gv loops
# 'shape' refers to the estimated degrees of freedom

# to estimate the parameters, working on a reset each year #

test <- garchFit(~garch(1,1), data=spdata[(1:1250),2],
cond.dist = "std", include.shape = TRUE)
test2 <- garchFit(~garch(1,1), data=spdata[(251:1500),2],
cond.dist = "std", include.shape = TRUE)
test3 <- garchFit(~garch(1,1), data=spdata[(501:1750),2],
cond.dist = "std", include.shape = TRUE)
test4 <- garchFit(~garch(1,1), data=spdata[(751:2000),2],
cond.dist = "std", include.shape = TRUE)
test5 <- garchFit(~garch(1,1), data=spdata[(1001:2250),2],
cond.dist = "std", include.shape = TRUE)
test6 <- garchFit(~garch(1,1), data=spdata[(1251:2500),2],
cond.dist = "std", include.shape = TRUE)
test7 <- garchFit(~garch(1,1), data=spdata[(1501:2750),2],
cond.dist = "std", include.shape = TRUE)
test8 <- garchFit(~garch(1,1), data=spdata[(1751:3000),2],
```

```
cond.dist = "std", include.shape = TRUE)
test9 <- garchFit(~garch(1,1), data=spdata[(2001:3250),2],
cond.dist = "std", include.shape = TRUE)
test10 <- garchFit(~garch(1,1), data=spdata[(2251:3500),2],
 cond.dist = "std", include.shape = TRUE)
test11 <- garchFit(~garch(1,1), data=spdata[(2501:3750),2],
cond.dist = "std", include.shape = TRUE)
test12 <- garchFit(~garch(1,1), data=spdata[(2751:4000),2],
cond.dist = "std", include.shape = TRUE)
test13 <- garchFit(~garch(1,1), data=spdata[(3001:4250),2],
 cond.dist = "std", include.shape = TRUE)
test14 <- garchFit(~garch(1,1), data=spdata[(3251:4500),2],
cond.dist = "std", include.shape = TRUE)
test15 <- garchFit(~garch(1,1), data=spdata[(3501:4750),2],
cond.dist = "std", include.shape = TRUE)
test16 <- garchFit(~garch(1,1), data=spdata[(3751:5000),2],
 cond.dist = "std", include.shape = TRUE)
test17 <- garchFit(~garch(1,1), data=spdata[(4001:5250),2],
cond.dist = "std", include.shape = TRUE)
test18 <- garchFit(~garch(1,1), data=spdata[(4251:5500),2],
 cond.dist = "std", include.shape = TRUE)
test19 <- garchFit(~garch(1,1), data=spdata[(4501:5750),2],
cond.dist = "std", include.shape = TRUE)
test20 <- garchFit(~garch(1,1), data=spdata[(4751:6000),2],
cond.dist = "std", include.shape = TRUE)
test21 <- garchFit(~garch(1,1), data=spdata[(5001:6250),2],
 cond.dist = "std", include.shape = TRUE)

m = 1250
test.gv <- numeric(6321)
test.gv[1:(m-1)] <- 0
test.gv[m] <- simple.m.variance

# to calculate variance #
#1
for (t in (m):(m+249)) {
  test.gv[t+1] <- (1.0360e-07 +
(1.6484e-02*sp.return.squared[t])
 +(9.8168e-01*garch.variance[t]) )}
# shape approx 5 here
#2
for (t in (m+250):(m+499)) {
  test.gv[t+1] <- (1.3070e-07 +
(2.0680e-02*sp.return.squared[t])
```

```
 +(9.7716e-01*garch.variance[t]) )}
# shape approx 5 here
#3
for (t in (m+500):(m+749)) {
  test.gv[t+1] <- (1.3220e-07 +
(2.7466e-02*sp.return.squared[t])
+(9.6983e-01*garch.variance[t]) )}
# shape approx 6 here
#4
for (t in (m+750):(m+999)) {
  test.gv[t+1] <- (6.9999e-07 +
(2.7765e-02*sp.return.squared[t])
 +(9.5465e-01*garch.variance[t]) )}
# shape approx 5 here
#5
for (t in (m+1000):(m+1249)) {
  test.gv[t+1] <- (7.2087e-07 +
 (4.2335e-02*sp.return.squared[t])
+(9.4151e-01*garch.variance[t]) )}
# shape approx 6 here
#6
for (t in (m+1250):(m+1499)) {
  test.gv[t+1] <- (5.8714e-07 +
(5.0524e-02*sp.return.squared[t])
 +9.3957e-01*garch.variance[t]) }
# shape approx 6 here
#7
for (t in (m+1500):(m+1749)) {
  test.gv[t+1] <- (6.5005e-07 +
(6.9698e-02*sp.return.squared[t])
 +(9.2660e-01*garch.variance[t]) )}
# shape approx 6 here
#8
for (t in (m+1750):(m+1999)) {
  test.gv[t+1] <- (8.2679e-07 +
(6.5580e-02*sp.return.squared[t])
+(9.3120e-01*garch.variance[t]) )}
# shape approx 6 here
#9
for (t in (m+2000):(m+2249)) {
  test.gv[t+1] <- (2.9901e-06 +
(8.5199e-02*sp.return.squared[t])
 +(8.9908e-01*garch.variance[t]) )}
# shape approx 7 here
```

```
#10
for (t in (m+2250):(m+2499)) {
  test.gv[t+1] <- (9.3542e-06 +
(7.9034e-02*sp.return.squared[t])
+(8.6282e-01*garch.variance[t]) )}
# shape approx 8 here
#11
for (t in (m+2500):(m+2749)) {
  test.gv[t+1] <- (8.1416e-06 +
 (8.3565e-02*sp.return.squared[t])
+(8.7644e-01*garch.variance[t]) )}
# go with base shape of 5
#12
for (t in (m+2750):(m+2999)) {
  test.gv[t+1] <- (2.5766e-06 +
(7.3737e-02*sp.return.squared[t])
+(9.1341e-01*garch.variance[t]) )}
# go with base shape of 5
#13
for (t in (m+3000):(m+3249)) {
  test.gv[t+1] <- (1.1317e-06 +
(7.0295e-02*sp.return.squared[t])
 +(9.2343e-01*garch.variance[t]) )}
# go with base shape of 5
#14
for (t in (m+3250):(m+3499)) {
  test.gv[t+1] <- (4.4432e-07 +
(5.0510e-02*sp.return.squared[t])
+(9.4627e-01*garch.variance[t]) )}
# go with base shape of 5
#15
for (t in (m+3500):(m+3749)) {
  test.gv[t+1] <- (4.2791e-07 +
(5.0997e-02*sp.return.squared[t])
+(9.4433e-01*garch.variance[t]) )}
# go with base shape of 5
test16
#16
for (t in (m+3750):(m+3999)) {
  test.gv[t+1] <- (5.2852e-07 +
(5.2370e-02*sp.return.squared[t])
  +(9.4178e-01*garch.variance[t]) )}
# 9 here
#17
```

```
test17
for (t in (m+4000):(m+4249)) {
  test.gv[t+1] <- (6.5275e-07+
(8.1150e-02 *sp.return.squared[t])
+(9.1873e-01*garch.variance[t]) )}
# 7 here
#18
for (t in (m+4250):(m+4499)) {
  test.gv[t+1] <- (7.4699e-07+
(8.8280e-02*sp.return.squared[t])
 +(9.1120e-01*garch.variance[t]) )}
# 6 here
#19
for (t in (m+4500):(m+4749)) {
  test.gv[t+1] <- (8.6080e-07+
(1.0728e-01*sp.return.squared[t])
+(9.0300e-01*garch.variance[t]) )}
# 5 here
#20
test20
for (t in (m+4750):(m+4999)) {
  test.gv[t+1] <- (2.2184e-06 +
(1.2220e-01*sp.return.squared[t])
+(8.8009e-01*garch.variance[t]) )}
# 6 here
#21
test21
for (t in (m+5000):(m+5071))  {
test.gv[t+1] <- (1.8524e-06+
(1.1087e-01*sp.return.squared[t])
+(8.8843e-01*garch.variance[t]) )}
# 5 here

test.gv
# to calculate the pth quantile of
# the standardised t distribution
x = 0.01
d.values <-c(5,5,6,5,6,6,6,6,7,8,5,5,5,5,5,9,7,6,5,6,5)
for (t in (1:21)) {
  t.d[t] <- qt(x, df= d.values[t])
}
for (t in (1:21)) {
  standardising.effect[t] <- sqrt((d.values[t]-2)/d.values[t])
}
```

```
standardising.effect <- sqrt((d-2)/d)
standardising.effect*t.d
multiplier <- standardising.effect*t.d

# To calculate 1% Daily VaR #
n = 250
std.gv.VaR <- numeric(6321)
std.gv.VaR[1:(1250)] = 0
for (t in (m):(m+n)) {
  std.gv.VaR[t+1] <- -multiplier[1]*sqrt(test.gv[t+1])
}
for (t in (m+n):(m+2*n)) {
  std.gv.VaR[t+1] <- -multiplier[2]*sqrt(test.gv[t+1])
}
for (t in (m+2*n):(m+3*n)) {
  std.gv.VaR[t+1] <- -multiplier[3]*sqrt(test.gv[t+1])
}
for (t in (m+3*n):(m+4*n)) {
  std.gv.VaR[t+1] <- -multiplier[4]*sqrt(test.gv[t+1])
}
for (t in (m+4*n):(m+5*n)) {
  std.gv.VaR[t+1] <- -multiplier[5]*sqrt(test.gv[t+1])
}
for (t in (m+5*n):(m+6*n)) {
  std.gv.VaR[t+1] <- -multiplier[6]*sqrt(test.gv[t+1])
}
for (t in (m+6*n):(m+7*n)) {
  std.gv.VaR[t+1] <- -multiplier[7]*sqrt(test.gv[t+1])
}
for (t in (m+7*n):(m+8*n)) {
  std.gv.VaR[t+1] <- -multiplier[8]*sqrt(test.gv[t+1])
}
for (t in (m+8*n):(m+9*n)) {
  std.gv.VaR[t+1] <- -multiplier[9]*sqrt(test.gv[t+1])
}
for (t in (m+9*n):(m+10*n)) {
  std.gv.VaR[t+1] <- -multiplier[10]*sqrt(test.gv[t+1])
}
for (t in (m+10*n):(m+11*n)) {
  std.gv.VaR[t+1] <- -multiplier[11]*sqrt(test.gv[t+1])
}
for (t in (m+11*n):(m+12*n)) {
  std.gv.VaR[t+1] <- -multiplier[12]*sqrt(test.gv[t+1])
}
```

```
for (t in (m+12*n):(m+13*n)) {
  std.gv.VaR[t+1] <- -multiplier[13]*sqrt(test.gv[t+1])
}
for (t in (m+13*n):(m+14*n)) {
  std.gv.VaR[t+1] <- -multiplier[14]*sqrt(test.gv[t+1])
}
for (t in (m+14*n):(m+15*n)) {
  std.gv.VaR[t+1] <- -multiplier[15]*sqrt(test.gv[t+1])
}
for (t in (m+15*n):(m+16*n)) {
  std.gv.VaR[t+1] <- -multiplier[16]*sqrt(test.gv[t+1])
}
for (t in (m+16*n):(m+17*n)) {
  std.gv.VaR[t+1] <- -multiplier[17]*sqrt(test.gv[t+1])
}
for (t in (m+17*n):(m+18*n)) {
  std.gv.VaR[t+1] <- -multiplier[18]*sqrt(test.gv[t+1])
}
for (t in (m+18*n):(m+19*n)) {
  std.gv.VaR[t+1] <- -multiplier[19]*sqrt(test.gv[t+1])
}
for (t in (m+19*n):(m+20*n)) {
  std.gv.VaR[t+1] <- -multiplier[20]*sqrt(test.gv[t+1])
}
for (t in (m+20*n):(6321)) {
  std.gv.VaR[t+1] <- -multiplier[21]*sqrt(test.gv[t+1])
}

## for 5% Daily VaR ##

z = 0.05
d.values <-c(5,5,6,5,6,6,6,6,7,8,5,5,5,5,5,9,7,6,5,6,5)
t.dz <-numeric(21)
for (t in (1:21)) {
  t.dz[t] <- qt(z, df= d.values[t])
}
for (t in (1:21)) {
  standardising.effect[t] <- sqrt((d.values[t]-2)/d.values[t])
}
multiplierz <- standardising.effect*t.dz
n = 250
std.gv.VaRz <- numeric(6321)
std.gv.VaRz[1:(1250)] = 0
for (t in (m):(m+n)) {
```

```
  std.gv.VaRz[t+1] <- -multiplierz[1]*sqrt(test.gv[t+1])
}
for (t in (m+n):(m+2*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[2]*sqrt(test.gv[t+1])
}
for (t in (m+2*n):(m+3*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[3]*sqrt(test.gv[t+1])
}
for (t in (m+3*n):(m+4*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[4]*sqrt(test.gv[t+1])
}
for (t in (m+4*n):(m+5*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[5]*sqrt(test.gv[t+1])
}
for (t in (m+5*n):(m+6*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[6]*sqrt(test.gv[t+1])
}
for (t in (m+6*n):(m+7*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[7]*sqrt(test.gv[t+1])
}
for (t in (m+7*n):(m+8*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[8]*sqrt(test.gv[t+1])
}
for (t in (m+8*n):(m+9*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[9]*sqrt(test.gv[t+1])
}
for (t in (m+9*n):(m+10*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[10]*sqrt(test.gv[t+1])
}
for (t in (m+10*n):(m+11*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[11]*sqrt(test.gv[t+1])
}
for (t in (m+11*n):(m+12*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[12]*sqrt(test.gv[t+1])
}
for (t in (m+12*n):(m+13*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[13]*sqrt(test.gv[t+1])
}
for (t in (m+13*n):(m+14*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[14]*sqrt(test.gv[t+1])
}
for (t in (m+14*n):(m+15*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[15]*sqrt(test.gv[t+1])
}
```

```
for (t in (m+15*n):(m+16*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[16]*sqrt(test.gv[t+1])
}
for (t in (m+16*n):(m+17*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[17]*sqrt(test.gv[t+1])
}
for (t in (m+17*n):(m+18*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[18]*sqrt(test.gv[t+1])
}
for (t in (m+18*n):(m+19*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[19]*sqrt(test.gv[t+1])
}
for (t in (m+19*n):(m+20*n)) {
  std.gv.VaRz[t+1] <- -multiplierz[20]*sqrt(test.gv[t+1])
}
for (t in (m+20*n):(6321)) {
  std.gv.VaRz[t+1] <- -multiplierz[21]*sqrt(test.gv[t+1])
}

# To plot #

std.gv.dates <- c(1993, 1994, 1995, 1996, 1997, 1998,
1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007,
2008, 2009, 2010, 2011, 2012, 2013)
std.gv.numbers <- c(0, 250, 500, 750, 1000, 1250, 1500,
1750, 2000, 2250, 2500, 2750, 3000, 3250, 3500, 3750,
4000, 4250, 4500, 4750, 5000)

# we start 5 years in, ie 1993

plot(std.gv.VaR[1251:6321], type = 'l', col="blue",
xlab = 'Date',  ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - GARCH Standardised t Distribution',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaRz[1251:6321], type = 'l', col="red",
xlab = 'Date',  ylab= 'Percentage of Portfolio at risk',
     main = 'Daily VaR - GARCH Standardised t Distribution',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaRz[1251:6321], type = 'l', col="red",
xlab = 'Date',  ylab= 'Percentage of Portfolio at risk',
```

```
     main = 'Daily VaR - GARCH Standardised t Distribution',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
    , axis(1, at=c(std.gv.numbers),
       labels=c(std.gv.dates)) )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

# to plot subsection #

close.up.dates <- c( 2001, 2002, 2003)
close.up.numbers <- c(0, 250, 500)

plot(std.gv.VaR[3251:3751], type = 'l', col="blue",
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'GARCH Standardised t Distribution
 Subsection',
     cex=0.75, ylim=c(0, 0.08)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaRz[3251:3751], type = 'l', col="red",
 xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'GARCH Standardised t Distribution
 Subsection',
     cex=0.75, ylim=c(0, 0.08)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaRz[3251:3751], type = 'l', col="red",
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'GARCH Standardised t Distribution
Subsection',
     cex=0.75, ylim=c(0, 0.08)
     , xaxt = "n", axis(1, at=c(close.up.numbers),
                        labels=c(close.up.dates)) )
legend("topleft",9.5, c("1% VaR", "5% VaR")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))
```

```
## To plot comparison plots subsection ##
# HS vs Standardised t GARCH m=1000 #
comp.std.gv.dates <- c(1994, 1995, 1996, 1997, 1998,
1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007,
                  2008, 2009, 2010, 2011, 2012, 2013)
comp.std.gv.numbers <- c(0, 250, 500, 750, 1000, 1250,
1500, 1750, 2000, 2250, 2500, 2750, 3000, 3250, 3500,
3750, 4000, 4250, 4500, 4750)

plot(daily.VaR[504:5324], type = 'l',
xlab = 'Date', ylab= 'Percentage of Portfolio at risk',
     main = 'Historical Simulation
\nand GARCH Standardised t Distribution Comparison',
     col="blue", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)
par(new=TRUE)
plot(std.gv.VaR[1501:6321], type = 'l', col="red",
xlab = 'Date', ylab= 'Percentage of Portfolio at risk',
     main = 'Historical Simulation
\nand GARCH Standardised t Distribution Comparison',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n")
# this starts at 1993
par(new=TRUE)
plot(std.gv.VaR[1501:6321], type = 'l', col="red",
xlab = 'Date', ylab= 'Percentage of Portfolio at risk',
     main = 'Historical Simulation
\nand GARCH Standardised t Distribution Comparison',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n", axis(1, at=c(comp.std.gv.numbers),
       labels=c(comp.std.gv.dates)) )
legend("topleft",9.5, c("HS, m=1000", "Garch(1,1) t model")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))

## To plot FHS against standardised t GARCH ##

fhs.250.dates <- c(1994, 1995, 1996, 1997, 1998,
1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006,
2007, 2008, 2009, 2010, 2011, 2012, 2013)
fhs.250.numbers <- c(0, 250, 500, 750, 1000, 1250,
1500, 1750, 2000, 2250, 2500, 2750, 3000, 3250,
```

```
3500, 3750, 4000, 4250, 4500, 4750)

plot(daily.FHS.VaR.250[1501:6321], type = 'l',
xlab = 'Date',  ylab= 'Percentage of Portfolio at risk',
     main = 'Filtered Historical Simulation
\nand GARCH Standardised t Distribution Comparison',
     col="blue", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n"
)
par(new=TRUE)
plot(std.gv.VaR[1501:6321], type = 'l', col="red",
xlab = 'Date', ylab= 'Percentage of Portfolio at risk',
     main = 'Filtered Historical Simulation
\nand GARCH Standardised t Distribution Comparison',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaR[1501:6321], type = 'l', col="red",
xlab = 'Date', ylab= 'Percentage of Portfolio at risk',
     main = 'Filtered Historical Simulation
\nand GARCH Standardised t Distribution Comparison',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n", axis(1, at=c(fhs.250.numbers),
                        labels=c(fhs.250.dates)) )
legend("topleft",9.5, c("FHS, m=250", "Garch(1,1) t model")
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))
## to plot subsection comparison ##

close.up.dates <- c( 2001, 2002, 2003)
close.up.numbers <- c(0, 250, 500)

plot(daily.FHS.VaR.250[3251:3751], type = 'l',
xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Filtered Historical Simulation
\nand GARCH Standardised t Distribution
Comparison Subsection',
     col="blue", cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaR[3251:3751], type = 'l', col="red",
xlab = 'Date',
```

```
     ylab= 'Percentage of Portfolio at risk',
     main = 'Filtered Historical Simulation
\nand GARCH Standardised t Distribution
Comparison Subsection',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n")
par(new=TRUE)
plot(std.gv.VaR[3251:3751], type = 'l', col="red",
 xlab = 'Date',
     ylab= 'Percentage of Portfolio at risk',
     main = 'Filtered Historical Simulation
\nand GARCH Standardised t Distribution
Comparison Subsection',
     cex=0.75, ylim=c(0, 0.13)
     , xaxt = "n", axis(1, at=c(close.up.numbers),
                        labels=c(close.up.dates)) )
legend("topleft",9.5, c("FHS", "GARCH(1,1) t model" )
       ,lty=c(1,1)
       ,lwd=c(2.5,2.5)
       , col=c("blue","red"))
```

Postal Address: 111 Barry Street
c/- The University of Melbourne
Victoria 3010 Australia

Email: enquiries@amsi.org.au
Phone: +61 3 8344 1777
Fax: +61 3 9349 4106
Web: www.amsi.org.au