Stockholm School of Economics

Department of Finance

Master Thesis

Parametric Value-at-Risk in Leptokurtic Distributions

A Study of Currencies and Commodities

Patrik Essunger & Lulu Valevie

2016-12-07

Abstract

Value-at-risk offers a quick estimate of the market risk exposure inherent in an asset or portfolio. A wide range of value-at-risk methods exist, which differ slightly in the estimation procedures and their assumptions. The most widely used methods belong to the parametric value-at-risk family, which attempts to fit a probability distribution to the underlying data. In the majority of cases the assumed distribution is normal. However, most financial data have been shown to follow leptokurtic distributions, yielding the assumption of normality void. In fact, it would lead to severe underestimation of the tail risk. This study explores the usage of the leptokurtic Student's t-distribution as an alternative to the normal distribution. Five different volatility estimation techniques are used on monthly data from two currencies and two commodity futures contracts. The estimated conditional volatilities are then applied to both distributions in order to assess the performance of each distribution and volatility estimation technique. The performance is gauged by backtesting the computed VaR levels through a proportion of failures test, an independence test as well as a joint test. The results only imply minor improvements when using the t-distribution. These results are due to the poor estimations of the tail distributions stemming from skewness in the sample distributions. Moreover, the implied volatility could potentially be a strong candidate for accurate VaR estimations if the underlying options are liquid, thus reflecting an efficient market.

Key words: *Value-at-Risk, Leptokurtic, Student's t-distribution, Currencies, Commodities*

Supervisor: Michael Halling, Associate Professor, Department of Finance, SSE

Acknowledgements

We would like to extend our gratitude to Anders Åslund and his team at SEB - Corporate Risk Advisory and Analytics for being both inspirational and supportive from the beginning until the end of writing this thesis. Also, we would like to thank our supervisor Michael Halling, Associate professor of Finance at the Stockholm School of Economics, for his helpful insights and advise throughout the whole process.

Stockholm, 2016-12-07

Patrik Essunger & Lulu Valevie

Contents

List of Figures

List of Tables

1 Introduction

Value-at-Risk (VaR) is a popular risk measurement that is widely used by financial institutions, non-financial corporations and fund managers in order to monitor the market risk for financial assets or portfolios. (Hull & White, 1998). In 1999 the use of VaR became mandatory for banks through the Basel II regulations and established itself as the most widely used market risk exposure metric. VaR represents the maximum potential change in the value of a portfolio or asset over a targeted horizon with a certain level of probability (Manganelli & Engle, 2001). One of the main attractions with VaR is that the risk is presented through a single number, either in absolute or relative terms. It is more intuitive than many other risk measures with more advanced metrics. VaR is easy to understand by anyone throughout any industry and instantly provides the user with an overview of the current risk situation. With the help of VaR, managers can set overall risk targets to aid in the decisions on risk objectives and position limits. Also, a VaR measurement is helpful for determining the internal capital allocation as well as evaluating the performances of the business units after implemented projects.

During the last decade, the increased instability and unpredictability of the financial markets have led to continuous improvements of the VaR metric, resulting in more sophisticated and well designed variations of this risk tool. The numerous VaR methods differ in their assumptions and in the computation of the density function for the forecasted portfolio or asset movements. Inaccurate estimations of the underlying risk distribution could lead to suboptimal capital allocation with negative consequences on not only the short-term profitability but also on the long-run financial stability of a company. Nevertheless, many financial assets have been shown to display kurtosis in excess to the normal distribution. Some risk managers therefore argue that the historical simulation VaR method would be better to use than a parametric approach (Hull & White, 1998). The leptokurtic behavior of returns have been shown in, among others, equity (Mandelbrot, 1963; Fama, 1965; Praetz, 1972; Blattberg & Gonedes, 1974; Gray & French, 1990; Kim & Kon, 1994; Bekaert et al., 1998), futures contracts (Cornew et al., 1984; Huang & Lin, 2004) and currencies (Tucker & Pond, 1988; Aggarwal et al., 1990). Generally, a negative skewness is also present in equities and commodities (Brooks & Persand, 2002). The assumption of normally distributed variables in the parametric approach will lead to a substantial underestimation of the probability of a tail event if the data in reality is leptokurtic. An attempt to explain the leptokurtic behavior of returns has been done through the use of the heteroskedastic conditional variances, estimated by methods such as the generalized autoregressive conditional heteroskedasticity (GARCH) and the exponentially weighted moving average (EWMA) models. However, even these have shown tendencies to underestimate the VaR. The other possible solution is to characterize the distributions of the returns by non-normal distributions more suited for the data that is being analyzed (Manganelli & Engle, 2001). In this paper one such distribution is explored, namely the Student's t-distribution.

However, consensus in the risk management field seems to be that several VaR methods are needed in order to receive a broader scope of the risk sources in a portfolio. Moreover, it is necessary to develop further statistical tools to verify the adequacy of the different VaR methods. Two of the most commonly known backtesting models that are used to assess the accuracy of the different VaR approaches are the proportion of failures test (Kupiec, 1995) and the independence test (Christoffersen, 1998) which together evaluate the most important characteristics of a VaR model (Lopez, 1998).

The purpose of this paper is to evaluate the accuracy of the parametric approach in computing VaR. Five different methods are used for estimating the conditional volatility: a simple moving average (SMA), an EWMA, a GARCH process, an asymmetrical variation of the GARCH called GJR-GARCH, and finally the implied volatility (IV). Furthermore, the VaR is calculated through the application of two different distributions, a normal distribution and a t-distribution that has been fitted to the data through its kurtosis. The five volatility estimation methods are examined for monthly data of two currency pairs, EUR/SEK and GBP/USD; as well of two commodity futures contracts, Brent crude oil (CO1) and copper (LMCADS03). In order to gauge the accuracy of these methods, backtesting is performed targeting four different horizons; one, three, six and twelve months; as well as three different probability levels; 1%, 5% and 32%. The backtesting is accomplished by the use of a proportion of failures test, an independence test and a joint test. Through the use of the two currency pairs and the two commodities, this paper hopes to illuminate possible differences both between and within the asset classes. Up to today, the larger part of the existing research has focused on daily VaR measures in common asset classes such as equity and options. Since this paper focus on the combination of commodities and currencies examined at longer target horizons and within more probability levels, it hopes to add to the prevailing literature within the parametric VaR approach.

2 Literature Review

2.1 Value-at-Risk

Value-at-Risk is a quantitative tool that measures the market exposure of an asset or portfolio. It can be interpreted as the loss threshold the asset will exceed with a specific probability over the course of the specified horizon. Given for example a 3-month 5%-VaR, the loss would be expected to exceed the VaR in 5% of all 3-month periods. Another interpretation of the VaR is that it is the maximum potential loss of an asset or a portfolio in absolute terms or in relative terms during a target horizon and based on a preselected confidence interval (Jorion, 1997). The risk is measured by estimating the volatility embedded in the asset or the portfolio through the use of statistical or simulation models. 'In the context of market risk, VaR measures the market value exposure of a financial instrument in case tomorrow is a statistically defined bad day' (Saunders & Allen, 1999). The VaR metric contains three parameters:

1. Denomination of VaR
2. Target Horizon
3. Confidence Level

The denomination of VaR determines whether the computed VaR is stated in either absolute terms or relative terms. Both the absolute and relative denomination should be easily understood by anyone with basic mathematical skills, which is one of many reasons as to why the VaR metric have made a major impact in risk management. The target horizon (Δt) is the planned holding period of the asset. The typical period used is one day, but there are no limitations to use one week, one month or any other horizon. The confidence level $(1-\alpha)$ refers to the specified probability for the estimation, with regular used levels at 95 percentage and the 99 percentage confidence intervals, which responds to a 5%-VaR and a 1%-VaR, respectively. A Δt-horizon α%-VaR$_t$ (or $VaR^{\alpha}_{\Delta t,t}$) can be explained mathematically as:

$$P(L_{\Delta t,t} < \mathrm{VaR}^{\alpha}_{\Delta t,t-1}) = \alpha \tag{2.1}$$

Where L_t is the loss (in either absolute or relative terms) during the target horizon and α is the probability corresponding to a $1-\alpha$ confidence level. VaR can be computed through a variety of methods that are either non-parametric or parametric driven. The three most commonly known methods are the historical simulation, Monte Carlo simulation and the variance-covariance model. The historical simulation assumes that the future movements of the underlying asset will have the same distribution as the past movements. The Monte Carlo simulation assumes that the future movements of the underlying asset will randomly follow a specified distribution. By simulating sufficiently many observations, an approximated distribution can be created. The variance-covariance model also known as the delta-normal model, assumes that the log-returns of the underlying asset are normally distributed and that changes in the underlying asset are linearly dependent. The delta-normal VaR can be computed as:

$$\mathrm{VaR}^{\alpha}_{\Delta t,t|\Phi} = \Phi^{-1}(\alpha)\ \sigma_t\ \sqrt{\Delta t} \tag{2.2}$$

Where $\Phi^{-1}(\alpha)$ is the inverse of the cumulative standard normal distribution with probability α. The Δt is, as discussed above, the chosen horizon and σ_t^2 is the conditional variance. Generally, the VaR concept can be illustrated as in Figure 1.

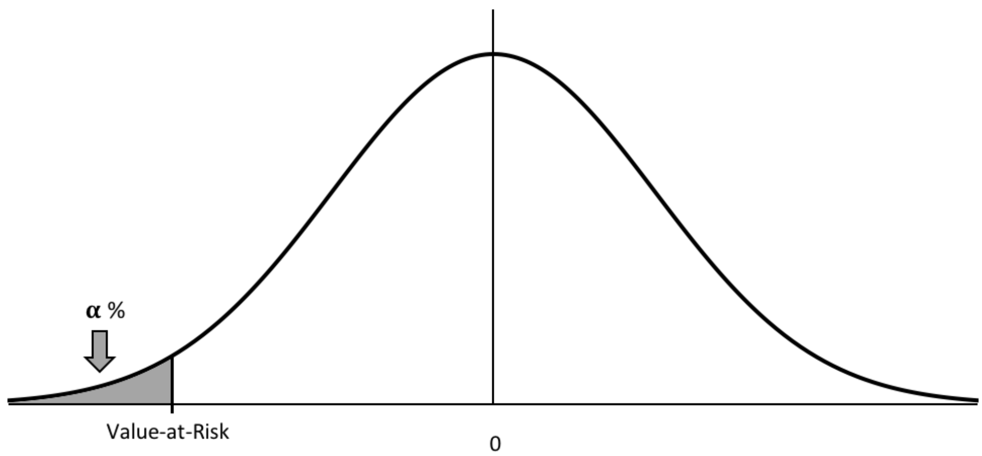

Figure 1: Graphical illustration of the value-at-risk concept. Here shown as a normal distribution, but the choice of distribution may differ between value-at-risk methodologies.

The first two methods, historical simulation and Monte Carlo simulation, belong to the non-parametric family of VaR. The third method, the variance-covariance method, instead belongs to the parametric family, which will be the primary focus in this thesis (Jorion, 1997). However, there are also methods that have characteristics from both the parametric and non-parametric families and are thus said to belong to the hybrid family. Examples of hybrid methods include the Extreme Value Theory (EVT) and Conditional Autoregressive Value at Risk by Regression Quantiles (CAViaR) (Manganelli & Engle, 2001).

2.1.1 Non-Parametric VaR

One of the most well known non-parametric approach, previously mentioned, is the historical simulation. It divides the return series of an asset into quantiles based on the specified α. The historical return in the series that is the cutpoint between the two lowest quantiles will be the estimated VaR. The advantages of this method is its simplicity in implementation, the readily available data and efficient calculations. The historical simulation is based on actual returns and includes all correlations that is embedded in the market rate changes, hence is not exposed to model risks. The lack of parametric assumptions make it possible to accommodate fat tails, skewness and other non-normal features which could otherwise cause problems for parametric methods. A main disadvantage of this technique is the complete dependency on the quality and length of the data sets. The pattern of past returns may not reflect the pattern of future returns, making the model slow to adapt to new market conditions. Furthermore, the use of small sample data could lead to the risk of having insufficient information about the distribution tails (Jorion, 1997). Other non-parametric estimations are the bootstrapping methods, density estimation, principal components and factor analysis methods.

2.1.2 Parametric VaR

The most common parametric approach is the variance-covariance method. It is also known as the delta-normal method when the return series is assumed to follow a normal distribution. A normal distribution indicates that a substantial part of the asset's returns will be close to the mean value, the rest tailing off symmetrically from the mean with a magnitude determined by the standard deviation. By assuming normally distributed variables, the

VaR can be computed by using the estimated mean and standard deviation of the return series. A standard assumption when using daily log-returns is to set the mean to zero. The standard deviation of the asset's returns can be calculated through different volatility estimation approaches such as the GARCH, implied volatility, the simple moving average and the EWMA. The same volatility estimation techniques can also be used to estimate the covariances within a portfolio. When the variances and covariance are know for the assets in the portfolio, the total variance can be calculated and used as input in the parametric VaR approach. However, the log-returns of financial data often exhibit leptokurtosis which begs to question whether the assumption of normally distributed variables is sound. In the cases where the log-returns are leptokurtic it would be better to apply other probability distributions with more kurtosis such as Student's t-distribution (Dowd, 2007).

The leptokurtic distribution looks similar to the curvature of a normal distribution, but with a more narrow and higher peak around the mean and fatter tails. The fat tails appear as a result of the higher frequency of outlier events and extreme observations. The kurtosis value, which helps to gauge an asset's level of risk is positive and greater than 3 in a leptokurtic distribution. This means that small changes occur less frequently since the historical values are clustered around the mean. However, there is a higher probability for extreme outcomes compared to a normal distribution within the fat tails. Data that is assumed to be normally distributed, but in reality is leptokurtically distributed, will generate a VaR that will underestimate the risk at higher levels of significance due to the fatter tails.

The strengths of the parametric approach are the simple calculations and the easily accessed data needed as input. However, the biggest weakness is the assumption of normality which is by far the most impactful parameter in the VaR computation. Another pitfall is the time-varying nature of the volatility. Oftentimes the return series exhibit heteroskedasticity, with periods when the volatility is substantially larger than usual. These periods coincide with extreme market events such as stock market crashes and defaults(Danielsson & De Vries, 2000).The volatility parameter in the parametric VaR models must be adjusted in the same manner, with larger volatility values in these periods. Unfortunately the most widely used volatility estimation techniques are too slow to react to quickly changing market conditions. This is where using a GARCH process to estimate the volatility offers its main advantage. The GARCH method quickly responds to and keeps the volatility elevated in periods of adverse markets. Also, if the volatility is varying over time it would be expected to find biased VaR estimates at longer horizons. An assumption when scaling up the VaR with the square root of the horizon is that the variance is constant over time. If then the volatility changes during that period, the correspond VaR level will be inaccurate. Another assumption of the delta-normal approach is a linear dependence of the returns in the underlying asset. This assumption is most of the times not a problem since the majority of the assets are linear such as stocks, futures, forwards and currencies. However, assets such as options have a non-linear relationship which makes the delta-normal useless. A delta-gamma method must then be used which takes into account the non-linearity of the underlying asset(Enders, 2010).

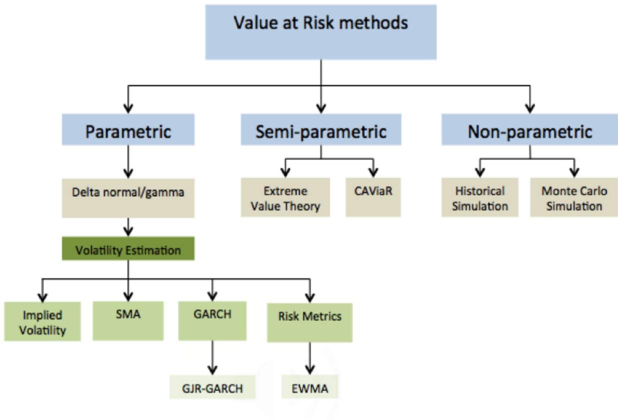

Figure 2: Graphical illustration of different value-at-risk approaches.

2.2 Critics of VaR

Even if VaR is an attractive metric there are also weaknesses that cannot be ignored, which have been pointed out by numerous researches and practitioners. The many approaches of the tool, from parametric to non-parametric methods, are known to give varying results that differ from each other, making it hard to actually compare the VaR measures. Also, even if similar theoretical methods are used, the implementation will affect the outcome; hence the impreciseness of VaR has been argued among researchers within the risk management field (Dowd, 2007). "False certainty is more dangerous than acknowledged ignorance" (Hoppe, 1998) and "you are worse off relying on misleading information than on not having any information at all" (Taleb, 1997), are some of the outspoken critics within the field. Furthermore the VaR can be seen as a much naive method coming from physical sciences that is not perfectly applicable on social systems. Important features of social systems are the non-stationary and dynamic interdependence of the market processes that are difficult to incorporate in the models, hence making the VaR estimates open for major errors. Taleb (1997) also pointed out that if VaR is used by everyone - the similar hedging behaviors will make uncorrelated risk factors become correlated, which in turn will lead to firms bearing much greater risk than what their VaR models suggest. This in turn could destabilize the whole financial system. Moreover, VaR does not describe the worst-case scenario in the tails, but only specifies where the tail begins. Firms that would like to know the maximum loss with 100% confidence must set other restrictions outside the VaR limits, such as operational limits, nominal orders or stop loss orders, in order to reach the highest possible coverage (Papaioannou & Gatzonas, 2002). There are metrics such as the conditional VaR that attempts to quantify the average loss within the tails but still struggles to deal with extreme events and other statistical anomalies, ofter referred to as black swan events (Rockafellar & Uryasev, 2002).

2.3 Model Risk

An important risk factor that is unavoidable is the model risk. Despite constantly improving risk measurements the model risk cannot completely be ignored. Every model used within risk management is to some degree exposed to model risk since the gap between what we assume to know and what we actually know makes us vulnerable to inaccuracies. A model is a simplified structure of the reality, hence will give various grades of error in the output. It is important to understand the caveats of a model and its components in order to make it suitable for the task at hand. Backtesting and evaluation of the performance of the risk model is continuously needed to ensure a minimal amount of model risk. Especially the estimation of the volatility parameter in VaR is exposed to this kind of risk. The model risk is present in the entire process of the VaR measure, from the specification of the volatility model but also in its implementation and application (Dowd, 2007).

2.4 Volatility Estimations

A key component for calculating the value-at-risk with a parametric approach is the volatility parameter. Volatility is normally estimated by calculating the variance, which is the average squared deviation from the mean as seen in Equation 2.3 below:

$$\mathrm{E}(\sigma_t^2) = \frac{\sum\limits_{t=1}^{T}(r_t - \mu_r)}{n} \tag{2.3}$$

Where μ_r is the mean return. However, research has shown that by assuming zero mean log-returns in daily data the estimations of volatility are more precise. The variance measure is then simplified to just the average of the squared returns. Volatility is often varying over time and exhibit clear signs of clustering, meaning that periods of large price changes tend to happen during the same small period of time. Such persistence of the volatility merits the use of volatility models that takes into account the volatility of past periods. Such models should quickly respond to changing market conditions, by having more weight in recent observations (Penza & Bansal, 2001).

2.4.1 Implied Volatility

The implied volatility (IV) is a forward-looking measure. It is derived from the prices of the at-the-money options on the underlying asset and displays the market's expectation of the asset's volatility in the future. Though the implied volatility shows the market's opinion of the asset's potential movements in the future, it does not forecast in which direction. A high implied volatility implies that the market thinks the asset has a high probability of going either up or down, whereas a low implied volatility indicates the asset will have a very small movement. While the historical volatility only refers to the past events, the implied volatility takes into account current and future events that might have an impact on the underlying asset price. From a mathematical standpoint the implied volatility can be calculated by iterating the Black-Scholes formula with different values of volatility until a solution for the current price is found. A common observation for option prices are that the implied volatility is greater for out-the-money and in-the-money options compared to at-the-money options. This characteristic is referred to as the volatility smile. One of the reasons for the volatility smile is the market's adjustment of the assumptions in the Black-Scholes formula. It assumes normally distributed returns, but since most financial time series are leptokurtic such an assumption will consistently underestimate the probability of experiencing large price swings in the underlying assets. The prices would thus be lower than what can be observed for out-the-money and in-the-money options (Penza & Bansal, 2001).

2.4.2 Simple Moving Average

The simple moving average (SMA) volatility estimation, also called the equally weighted moving average, is the simplest method to estimate short-term volatility. This is achieved by using an equally weighted moving average of the squared returns and is described by Equation 2.4 below:

$$\mathrm{E}(\sigma_t^2) = \frac{\sum_{i=t-n}^{t} r_i^2}{n} \tag{2.4}$$

Where r_i represents the return at time i and n represents the time period over which the moving average is calculated. A standard time window used by both academics and practitioners when dealing with daily data is 30-60 days (Penza & Bansal, 2001). This is a very small period considering many assets have price data spanning several decades. By adopting a 60 day time-window the volatility measure only takes into account the price changes during the past 60 days, while completely disregarding the data prior to the past 60 days, hence why it is a measurement of the *short-term* volatility. Due to this fact, only the price data from the last 60 days needs to be stored in order to estimate the volatility. Applying the same logic to monthly data where the volatility can change considerably from month to month, a short time-period is desirable. However, a time period of 60 days, or only 2 months, would be too few observations for any meaningful moving average estimation. The volatility estimated by a model with only two time-periods would be very spiky and not be able to tell the overall trend of the volatility. A more reasonable period of time would be 12 months. It would still put significant weight into the most recent observation while containing enough observations for a somewhat accurate estimate of the volatility.

This method does however suffer from a major drawback called *ghost features*. The *ghost features* refer to the effect on the estimated volatility following an unusually large positive or large negative return. After a day of unusual returns, the simple moving average volatility will increase sharply. That could be a desirable feature considering that volatility exhibit clustering behavior. The drawback however is the extreme persistence of volatility shocks. An abnormal return will have a lasting effect on the volatility measure over the next n periods and then again sharply decline. In reality the volatility will have declined to normal levels in a much shorter time than n periods. All the returns over the entire period are given the same weight regardless of the fact that some returns - abnormal returns or returns occurring far back in time - might in reality have minor to no impact on the forecast of the volatility of future returns. Also, with the use of a large sample, the calculations of the volatility will be diluted, hence probably less accurate.

2.4.3 Exponentially Weighted Moving Average

Within the RiskMetrics approach, the variance is computed using an Exponentially Weighted Moving Average (EWMA). In contrast to the equally weighted moving average, which puts an equally large weight on every single day within the specified time-period, the exponential weighted moving average puts significantly more weight into observations that happened more recently. The EWMA method can be described by the formula Equation 2.5 below:

$$\mathrm{E}(\sigma_t^2) = \lambda \mathrm{E}(\sigma_{t-1}^2) + (1-\lambda) r_t^2 \tag{2.5}$$

Where λ is the smoothing parameter that must take on a value less than one. The recommendation from RiskMetrics is to set λ equal to a value of 0.94 when using daily data and 0.97 when using monthly data (RiskMetrics Group, 1996). These numbers have been chosen such that the mean square error is as small as possible. In the EWMA approach, each of the squared returns is weighted through a recursive multiplier, rather than having equal weights like in the simple moving average model. The volatility will quickly decline after an abnormal day of return, hence decreases the problem with *ghost features*. However, using the exponentially weighted moving average method with monthly averages will still incorporate *ghost features* due to its inherent nature in the calculation of such monthly averages. The exponential weighted moving average model is the most widely used method for estimating volatility due to its simplicity and low data usage. The only data needed is the most recent observations and the volatility estimations calculated prior to that date. However, it is known to be a rather poor candidate when forecasting the correct quantiles in the presence of fat tail distributions of return (Dowd, 2007; Penza & Bansal, 2001). The RiskMetrics approach to VaR have been shown to only be globally acceptable at the 5% level (Giot & Laurent, 2003; Pafka & Kondor, 2001).

2.4.4 Generalized Autoregressive Conditional Heteroskedasticity

Usually when using economic models the residual, or error term, is assumed to have a constant variance σ_ε^2. The assumption of constant variance of the residual is called *homoskedasticity*. However, in some circumstances when a series exhibit time-variation, the homoskedasticity assumption is incorrect. When a series exhibit time-variation it is often referred to as *heteroskedasticity*. In order to tackle the problem of heteroskedasticity, Engle (1982) developed the autoregressive conditional heteroskedasticity (ARCH) process, which is a method for modelling the conditional variance of a series. The method have since been extended to the generalized autoregressive conditional heteroskedasticity (GARCH) process by Bollerslev (1986) to allow for both autoregressive as well as moving-average parameters when modelling the conditional variance. In a GARCH model, the conditional volatility is modelled by an (autoregressive moving-average) ARMA(P,Q) process where P and Q is the number of lags used for the autoregressive and moving-average parameters. Hence, it is often referred to as a GARCH(P,Q) process in the same manner as the ARMA(P,Q) process.

The most widely used GARCH specification when dealing with financial data is the GARCH(1,1). It is specifically useful when the shocks in the volatility exhibit a high degree of persistence. For the sake of practicality and comparability, the scope of this paper is limited to the GARCH(1,1) specification based on the past 5 years of monthly data. Most of the following derivation is based on Enders (2010). The first step of defining a GARCH(1,1) process is by considering how to model the variable of interest, in this case the logarithmic return variable. By assuming a zero mean return, the model of the return series $\{r_t\}$ is simply equal to the residual:

$$r_t = \varepsilon_t \tag{2.6}$$

From this equation the variance of the residual series $\{\varepsilon_t\}$ can be modelled by the error process:

$$\varepsilon_t = z_t\sqrt{\sigma_t^2} = z_t\sigma_t \tag{2.7}$$

Where σ_t^2 is the conditional variance of the return series and $\{z_t\}$ is a white-noise process with $\mathrm{E}(z_t) = 0$ and $\mathrm{E}(z_t^2) = 1$. Now assume that the conditional variance, σ_t^2, follows an ARMA(1,1) process:

$$\sigma_t^2 = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \beta_1\sigma_{t-1}^2 \tag{2.8}$$

Since $\{z_t\}$ is a white-noise process, both the unconditional and conditional mean of ε_t is:

$$\mathrm{E}(\varepsilon_t) = \mathrm{E}(z_t\sigma_t) = \mathrm{E}(z_t)\mathrm{E}(\sigma_t) = 0 \tag{2.9}$$

This result is thus consistent with the assumption of a zero mean daily return. Furthermore it is possible to estimate the conditional variance of ε_t:

$$\mathrm{E}_{t-1}(\varepsilon_t^2) = \mathrm{E}_{t-1}(z_t^2\sigma_t^2) = \sigma_t^2 \tag{2.10}$$

The conditional variance can be viewed as the short-term fluctuations in the volatility level. This measure is much better for forecasting the volatility as it takes into account the past and current observations. Worth noting is that ε_t^2 is in itself just the conditional variance of the residual and not the conditional variance of the return. However, as seen above, with this model specification - the conditional variance of the residual and the returns will be equal to one another.

Since the variance of the residual series $\{\varepsilon_t\}$, and thus the variance of the return series $\{r_t\}$, can be modelled by an ARMA process it should have the same autocorrelation pattern required by such a model. With that being said, if a GARCH model is adequate - the squared residuals have to exhibit some level of autocorrelation different from zero. If the autocorrelations of the squared residual series are zero, then it would be impossible to model the conditional variance due to the fact that it implies that the variance at any given day is completely random. By looking at the autocorrelation function (ACF) and the partial autocorrelation function (PCF) it is possible to check if the autocorrelations are significantly different from zero. A more formal test can be made such as the Ljung-Box test (Ljung & Box, 1978). It is a portmanteau test which tests the null hypothesis that the time series exhibits no autocorrelation against the hypothesis that is *do* exhibit autocorrelation at the chosen lag. The test statistic is:

$$Q = T(T+2)\sum_{k=1}^{L}\left(\frac{\rho(k)^2}{T-k}\right) \tag{2.11}$$

Where T is the number of observations, L is the number of autocorrelation lags and $\rho(k)$ is the autocorrelation at lag k. The distribution of Q under the null hypothesis is a chi-square distribution with L degrees of freedom.

If the model is shown to exhibit significant autocorrelation it is said to have GARCH errors, which enables modelling and forecasts of the conditional variance such that the variance in one period depend on the variance in prior periods. That is one of the major reasons as to why GARCH modelling of the variance is a strong tool when used on time-varying volatility.

2.4.5 Alternative GARCH Specification

The GARCH framework, while incredibly useful in itself, have since its introduction been modified in various ways in order to satisfy other desirable properties. One such modified GARCH model were proposed by (Glosten et al., 1993). In their paper they found a negative relationship between the conditional expected monthly returns and the conditional variance of the monthly returns. In other words, they found that negative returns gave rise to higher conditional variance than that of positive returns of the same magnitude. This property can not be modelled by the ordinary GARCH model which is symmetric in the sense that positive and negative shocks result in an equal change in conditional variance. Glosten, Jagannathan and Runkle (GJR) instead proposed a slight change to the original GARCH model in order to incorporate leverage effects in the conditional variance:

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \gamma \varepsilon_{t-1}^2 I_{t-1} \tag{2.12}$$

Where I_{t-1} is an indicator function such that:

$$I_{t-1} = \begin{cases} 1, & \text{if } r_{t-1} < 0 \\ 0, & \text{if } r_{t-1} \geq 0 \end{cases} \tag{2.13}$$

Compared to Equation 2.8, which is how the GARCH(1,1) process models the conditional variance, the GJR-GARCH adds an additional variable that introduce asymmetry into the GARCH process.

2.5 Backtesting

Assessing how well a method for calculating VaR performs is done by considering the number of observations that exceeds the VaR level. The number of failures in the VaR model should correspond to the confidence level used in that specific case. A VaR model is considered unsatisfactory if either the number of failures are too low or too high. When the number of failures are too low, the model overestimates the VaR which could be costly if it is used as a basis for hedging. The amount that is hedged would thus be unnecessarily high given the true level of risk. On the other hand if the number of failures are too high, the model is underestimating the level of risk. It is also important to consider the independence of the failures. A well functioning VaR model should have a constant proportion of failures independent from the outcome in the previous period. The backtesting procedure consists of three tests; proportion of failures, independence and conditional coverage, which are described in the next sections.

2.5.1 Proportion of Failures

A reasonable start for the backtesting procedure is to consider the proportion of failures of the VaR model. The number of failures by a model should correspond to the VaR level of that model. When using a $\alpha\%$-VaR it should be expected to find that approximately $\alpha\%$ of the observed returns at time t are below the calculated VaR at time $t-1$. If however the proportion of failures are either too high or too low, the specified model is inappropriate.

2.5.2 Independence

The proportion of failures test only takes into account the number of failures. While that is useful knowledge for backtesting a VaR model it is by construction not taking into account *when* the failures occur. Imagine a VaR model with an α of 5% where the proportion of failures when backtesting are exactly 5%. The backtesting method specified in subsubsection 2.5.1 would suggest that the used VaR model is perfect. However, if the observed failures are all realized during a narrow window of time, the actual risk is much greater than suggested by the model. Consider what would happen to a company if all left-tail events would take place within just a couple of weeks - it would most likely become bankrupt. Should the model be adequate, one should expect an $\alpha\%$ probability to exceed the VaR level in any following period regardless of the outcomes in previous periods. It is therefore important to also take into account the independence of the failures when backtesting a VaR model.

2.5.3 Conditional Coverage

The final backtesting is a procedure that combines the aforementioned two test. The idea is to jointly test if the VaR model both yields a correct number of failures and if the failures are independent.

3 Methodology

3.1 Data

Table 1: Descriptive statistics of monthly log-returns

	EUR/SEK	GBP/USD	CO1	LMCADS03
T	212	244	339	365
Maximum	6.64%	9.04%	37.96%	30.36%
Minimum	-5.38%	-10.22%	-40.74%	-43.93%
Median	-0.02%	-0.08%	0.63%	0.49%
Mean	0.04%	-0.07%	0.35%	0.33%
Standard Deviation	1.59%	2.47%	9.34%	7.29%
Skewness	0.223	-0.434	-0.177	-0.547
Excess Kurtosis	2.707	1.847	2.326	4.990
Degrees of Freedom	6.217	7.248	6.580	5.202

Table 2: Descriptive statistics of 3 months log-returns

	EUR/SEK	GBP/USD	CO1	LMCADS03
T	209	241	336	362
Maximum	11.49%	13.89%	88.48%	49.87%
Minimum	-7.53%	-21.04%	-76.70%	-72.98%
Median	-0.26%	-0.08%	1.85%	1.15%
Mean	0.10%	-0.21%	1.06%	1.04%
Standard Deviation	2.74%	4.46%	18.15%	13.94%
Skewness	0.302	-0.773	-0.406	-0.738
Excess Kurtosis	2.336	3.771	3.525	5.492
Degrees of Freedom	6.569	5.591	5.702	5.093

Table 3: Descriptive statistics of 6 months log-returns

	EUR/SEK	**GBP/USD**	**CO1**	**LMCADS03**
T	206	238	333	359
Maximum	18.83%	16.92%	75.49%	62.99%
Minimum	-11.29%	-31.13%	-112.07%	-101.96%
Median	-0.24%	-0.25%	3.23%	0.33%
Mean	0.18%	-0.39%	2.19%	2.13%
Standard Deviation	3.93%	6.65%	25.29%	21.01%
Skewness	0.702	-1.277	-0.858	-0.458
Excess Kurtosis	3.010	4.405	2.745	3.762
Degrees of Freedom	5.993	5.362	6.186	5.595

Table 4: Descriptive statistics of 12 months log-returns

	EUR/SEK	**GBP/USD**	**CO1**	**LMCADS03**
T	200	232	327	353
Maximum	19.90%	17.25%	94.85%	95.30%
Minimum	-16.27%	-32.87%	-78.49%	-89.37%
Median	0.34%	-0.07%	3.57%	0.67%
Mean	0.39%	-0.70%	3.74%	4.27%
Standard Deviation	5.83%	8.95%	32.31%	28.82%
Skewness	0.072	-1.008	-0.208	0.246
Excess Kurtosis	0.876	1.915	0.027	0.698
Degrees of Freedom	10.847	7.134	226.222	12.597

The data seen in the above tables (Table 1 - Table 4) have been calculated using MATLAB and the sample data outlined in the following subsections. As seen in the tables above, there are considerable differences between the assets. The difference is especially large between the currencies and the commodities. While all four assets have means and medians close to zero, the commodities have much larger standard deviations. All assets except EUR/SEK have a negative skewness, thus having wider and/or fatter left tails. The negative skewness indicates that there is a higher probability of left-tail events. The excess kurtosis is non-zero and positive, meaning that the return distributions have heavier tails than a normal distribution. The degrees of freedom are calculated from the fourth standardized moment of student's t-distribution (Lin & Shen, 2006). It can easily be calculated by solving for ν in the equation:

$$\text{Excess Kurtosis} = \frac{6}{\nu - 4} \qquad \forall\, \nu > 4 \tag{3.1}$$

The high excess kurtosis is also present when looking at the returns over longer periods of time. The degrees of freedom associated with the excess kurtosis seem to hover around 6 for all return horizons from 1 month to 6 months as seen in Table 1, Table 2 and Table 3. For EUR/SEK, GBP/USD and LMCADS03 the degrees of freedom at 12-months increases slightly. However, for CO1, the degrees of freedom increase drastically. With an estimated degrees of freedom at 226 in Table 4, this distribution will be very close to being normally distributed.

3.1.1 EUR/SEK

Figure 3: Graph over the end-of-month prices and log-returns of EUR/SEK.

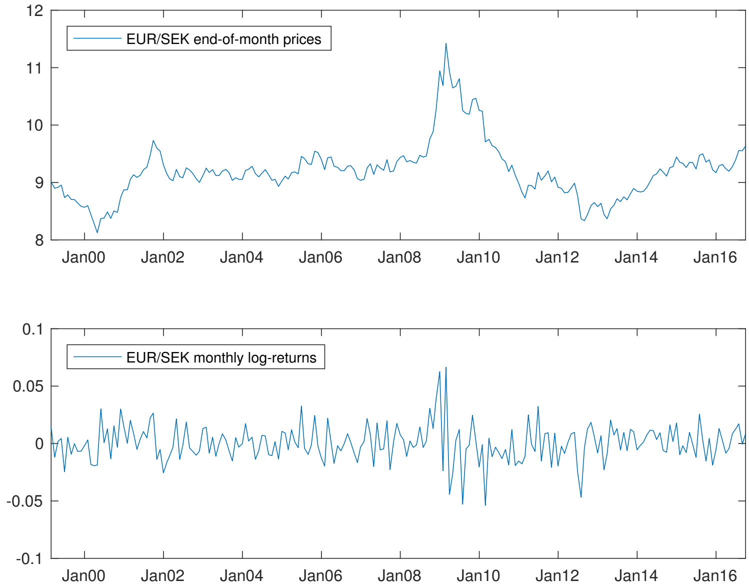

The end-of-month data for the currency pair EUR/SEK is downloaded from Bloomberg with the ticker EURSEK Crncy. The data is ranging from the end of January 1999 to the end of September 2016. The data consists of the price of the underlying currency pair, the implied volatility and the forward rates. The price of the underlying currency pair is the spot price quoted at the last trading day of any given month. It is quoted as the spot price in SEK for buying one unit of EUR. The measure for implied volatility is quoted for one month, three months, six months and one year into the future. Implied volatility is quoted in percentages. The forward rates are also quoted for one month, three months, six months and one year into the future. It is the difference between the spot price and the forward price for each individual period of time. The quotation form is in pips with four decimal places.

3.1.2 GBP/USD

Figure 4: Graph over the end-of-month prices and log-returns of GBP/USD.

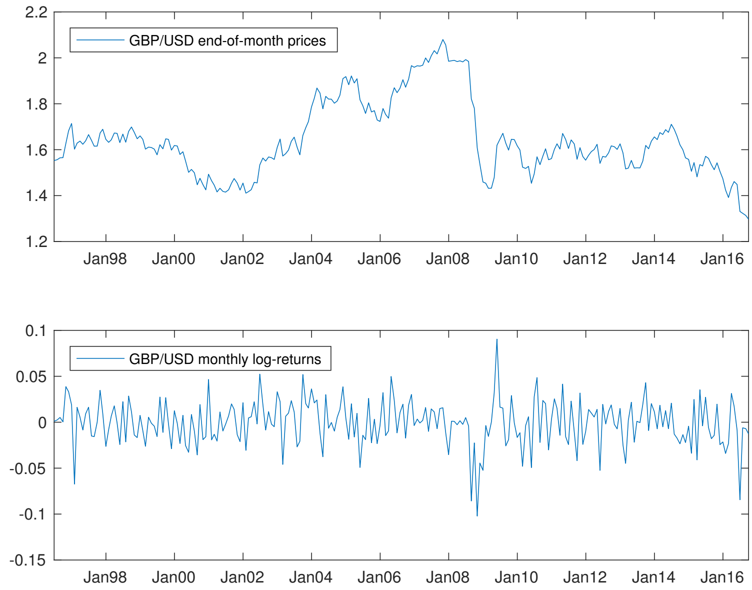

The end-of-month data for the currency pair GBP/USD is downloaded from Bloomberg with the ticker GBPUSD Crncy. The data is ranging from end of the May 1996 to the end of September 2016. The data consists of the price of the underlying currency pair, the implied volatility and the forward rates. The price of the underlying currency pair is the spot price quoted at the last trading day of any given month. It is quoted as the spot price in USD for buying one unit of GBP. The measure for implied volatility is quoted for one month, three months, six months and one year into the future. Implied volatility is quoted in percentages. The forward rates are also quoted for one month, three months, six months and one year into the future. It is the difference between the spot price and the forward price for each individual period of time. The quotation form is in pips with four decimal places.

3.1.3 Brent Crude Oil

Figure 5: Graph over the end-of-month prices and log-returns of CO1.

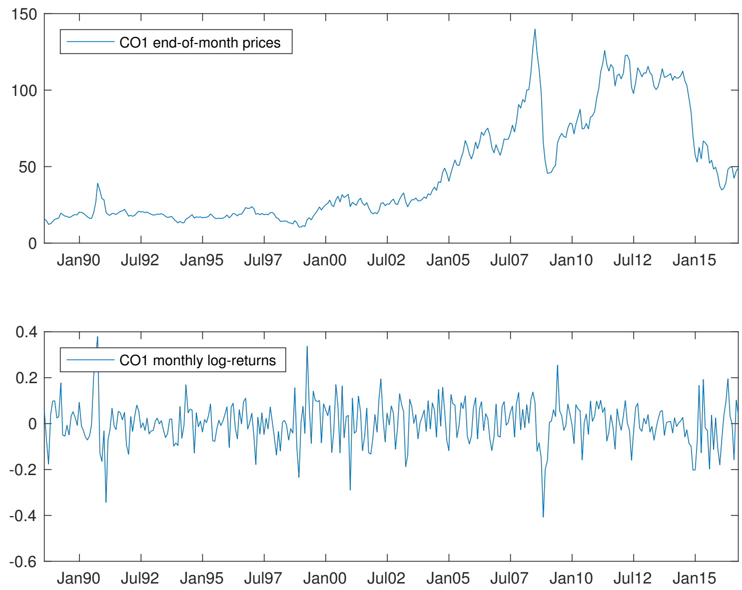

The end-of-month data for Brent crude oil is downloaded from Bloomberg with the ticker CO1 Cmdty, representing the front month contract traded at the intercontinental exchange (ICE). It is the Brent crude oil future contract which is closest to its expiration date, always within one month. The data is ranging from the end of June 1988 to the end of September 2016. The data consists of the price of the underlying futures contract quoted in USD for the contract size of 1,000 barrels. The price of the futures contract is the one quoted at the last trading day of any given month. Due to an insufficient amount of the implied volatility observations in the data, the implied volatility will not be included for Brent crude oil.

3.1.4 Copper

Figure 6: Graph over the end-of-month prices and log-returns of LMCADS03.

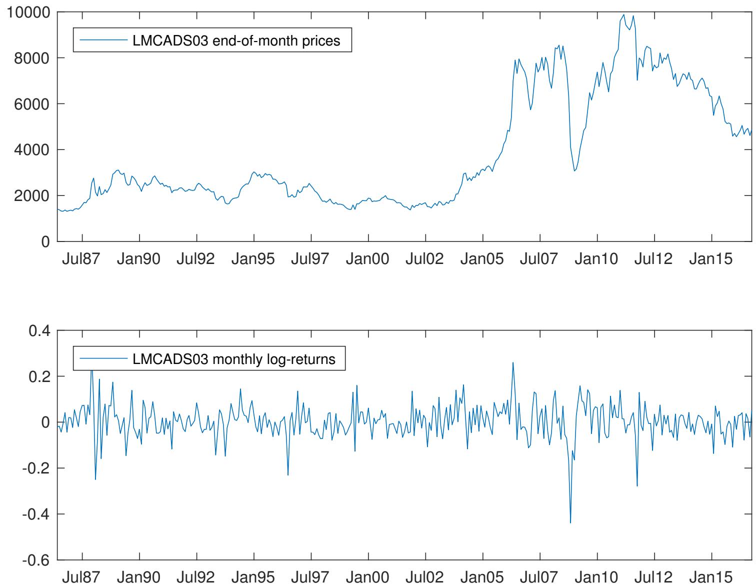

The end-of-month data for copper is downloaded from Bloomberg with the ticker LMCADS03 Cmdty. It represents the three month futures contract traded at the London Metal Exchange (LME). The data is ranging from the end of April 1986 to the end of September 2016. The data consists of the price of the underlying futures contract quoted in USD for the contract size of one metric tonne. The price of the futures contract is the one quoted at the last trading day of any given month. Due to an insufficient amount of the implied volatility observations in the data, the implied volatility will not be included for copper.

3.2 Value-at-Risk Methodologies

The focus of this thesis will be on the parametric VaR approach termed the variance-covariance method. It is a method that assigns a probability distribution to the log-returns of an asset. Most often a normal distribution is used and the method is then referred to as the delta-normal method. However, as seen in the subsection 3.1 section, the asset returns show leptokurtic characteristics which would justify using an alternative distribution such as the t-distribution. A comparison between the normal and Student's t probability distribution function can be seen in Figure 7.

Figure 7: Probability density function of a normal distribution and two t-distributions with different degrees of freedom.

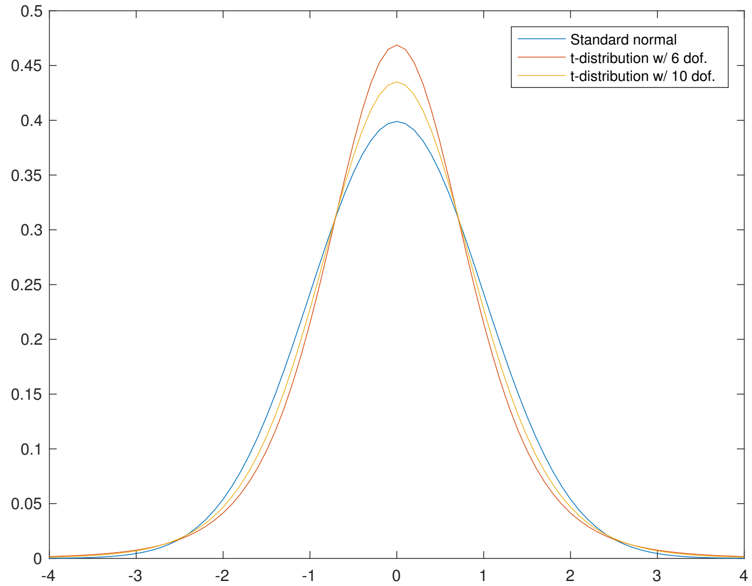

In order to compute the VaR with the delta-normal method in relative terms, the horizon (Δt) and confidence (α) of the measure should be selected. This paper looks at the horizons of 1-month, 3-months, 6-months and 12-months. The chosen confidence levels are the 99%, 95% and 68%. The final parameter, and also the only non-arbitrary variable, is the conditional variance (σ_t^2) of the return series. The VaR at time t can then be computed as:

$$\mathrm{VaR}^{\alpha}_{\Delta t,t|\Phi} = \Phi^{-1}(\alpha)\ \sigma_t\ \sqrt{\Delta t} \tag{3.2}$$

Where Φ^{-1} represents the inverse of the standard normal distribution.

If instead the t-distribution is used in the variance-covariance method an additional parameter needs to be calculated. That parameter is the degree of freedom which determines, among other characteristics, the level of kurtosis in the distribution. In subsection 3.1 the degree of freedom corresponding to each horizon and return series is shown. By adopting the t-distribution it is possible to better capture tail events in the leptokurtic data. However, when using a t-distribution, the standard deviation needs to be adjusted depending on the degrees of freedom used. The reason being that the t-distribution do not have a variance parameter in the same sense as the normal distribution. Instead the variance is determined by the degrees of freedom such that:

$$V(T) = \frac{\nu}{\nu - 2} \qquad \forall\, \nu > 2 \tag{3.3}$$

Where T is a t-distributed variable and ν the degrees of freedom. In order to maintain the same variance as in a standard normal distribution we have to define a new variable that depends on T:

$$Y = aT + x \tag{3.4}$$

The mean is then:

$$E(Y) = b \tag{3.5}$$

And the variance is:

$$V(Y) = a^2 V(T) = a^2 \frac{\nu}{\nu - 2} \tag{3.6}$$

It can now be seen that by setting $b = 0$ and $a = \sqrt{(\nu - 2)/\nu}$ a zero mean and a variance of one is obtained. The VaR at time t using a t-distribution can thus be computed as:

$$\text{VaR}^{\alpha}_{\Delta t, t|t} = t_{\nu}^{-1}(\alpha)\ \sqrt{(\nu - 2)/\nu}\ \sigma_t\ \sqrt{\Delta t} \tag{3.7}$$

Where ν is the estimated degrees of freedom. As ν approaches infinity then $t_\nu(\alpha)$ approaches a normal distribution and $\sqrt{(\nu - 2)/\nu}$ approaches one. The convergence towards a normal distribution appears very fast. With degrees of freedom above 20 there is barely any difference between the t-distribution and the normal distribution. The Jarque-Bera normality test is used to determine whether the data is normally distributed (Jarque & Bera, 1987). It defines the test variable, JB, as:

$$JB = \frac{n}{6}\left(s^2 + \frac{(k - 3)^2}{4}\right) \tag{3.8}$$

Where n is the sample size, s is the skewness and k is the kurtosis. Under the null hypothesis where the data is normally distributed, the test statistic follows a chi-squared distribution with two degrees of freedom for large sample sizes.

3.3 Backtesting

3.3.1 Proportion of Failures

In order to determine whether the proportion of failures are at an adequate level it is possible to test the null hypothesis that $q = \pi$, where q is the true proportion of failures specified by the model and π is the observed proportion of failures. The technique were first developed by Kupiec (1995). Begin by defining the failure series $\{I_t\}$ as:

$$I_{t+1} = \begin{cases} 1, & \text{if } r_{t+1} < VaR_t^\alpha \\ 0, & \text{if } r_{t+1} \geq VaR_t^\alpha \end{cases} \qquad t = 1, ..., T \tag{3.9}$$

Where r_{t+1} is the logarithmic return at time $t+1$ and VaR_t^α is the calculated VaR at time t given the threshold of α. Using this sequence it is possible to construct the likelihood function:

$$L(\pi) = (1-\pi)^{T_0}\pi^{T_1} \tag{3.10}$$

Where T_1 is the number of failures, ie. the sum of I_{t+1} and T_0 is the number of observations without a failure. The π parameter can be estimated by simply dividing the number of failures by the total number of observations in the $\{I_{t+1}\}$ ($\hat{\pi} = T_1/(T_1 + T_0)$). Substituting in the estimate for π into Equation 3.10 yields the likelihood function of the alternative hypothesis:

$$L(\hat{\pi}) = (1 - (T_1/(T_1+T_0))^{T_0}(T_1/(T_1+T_0))^{T_1} \tag{3.11}$$

Under the null hypothesis the likelihood function becomes:

$$L(q) = (1-q)^{T_0}q^{T_1} \tag{3.12}$$

The standard technique for comparing two models are the likelihood-ratio test:

$$\begin{aligned} LR_{pof} &= -2\ln\left(\frac{L(q)}{L(\hat{\pi})}\right) \\ &= -2\ln\left(\frac{(1-q)^{T_0}q^{T_1}}{(1-(T_1/(T_1+T_0))^{T_0}(T_1/(T_1+T_0))^{T_1}}\right) \end{aligned} \tag{3.13}$$

When the number of observations becomes large, it will approximately follow a chi-square distribution with one degree of freedom. As the observed proportion of failures $\hat{\pi}$ approaches

q, the likelihood ratio $(L(q)/L(\hat{\pi}))$ will approach one. The test statistic LR_{pof} will thus approach zero. When $\hat{\pi}$ is different from q the likelihood ratio $(L(q)/L(\hat{\pi}))$ will be between one and zero. Consequently the test statistic (LR_{pof}) will increase in value. Due to these characteristics, a right-tailed chi-squared test can be applied using one degree of freedom as mentioned above. The null hypothesis can then be rejected at a 5% level if the test statistic is above the critical value of 3.8415, which corresponds to a right-tail probability of 5% when using one degree of freedom. Alternatively, the significance of a VaR method can also be directly assessed by calculating the P-value:

$$\text{P-value} = 1 - F_{\chi_1^2}(LR_{pof}) \tag{3.14}$$

Where $F_{\chi_1^2}$ represents the cumulative probability distribution function for chi-squared with one degree of freedom. The null hypothesis is then rejected if P-value $< 5\%$.

3.3.2 Independence

A test for determining the independence of a VaR model were proposed by Christoffersen (1998). His proposed procedure begins by dividing the possible outcomes at any given date into four parts:

$$\Pi_1 = \begin{bmatrix} \pi_{00} & \pi_{01} \\ \pi_{10} & \pi_{11} \end{bmatrix} = \begin{bmatrix} 1-\pi_{01} & \pi_{01} \\ 1-\pi_{11} & \pi_{11} \end{bmatrix} \tag{3.15}$$

Where π_{01} is the probability of observing a failure given that there was no failure in the previous period ($\pi_{01} = P(I_{t+1} = 1|I_t = 0)$), π_{11} is the probability of observing a failure given that there was a failure in the previous period ($\pi_{11} = P(I_{t+1} = 1|I_t = 1)$) and so forth. An assumption with this method is that only the outcome in last period matters for the outcome in the following period. If the observed failures are independent, one would expect that it does not matter whether the observed outcome in last period was a failure or not. The likelihood function of the above process can be written as:

$$L(\Pi_1) = (1-\pi_{01})^{T_{00}}\pi_{01}^{T_{01}}(1-\pi_{11})^{T_{10}}\pi_{11}^{T_{11}} \tag{3.16}$$

Where $T_{ij}, i, j = 0, 1$ is the number of observations when j follows an i. The maximum likelihood estimates are calculated by taking the first derivative of the likelihood function with respect to π_{01} and π_{11} and setting it equal to zero and solving for the corresponding variable:

$$\begin{aligned}\hat{\pi}_{01} &= \frac{T_{01}}{T_{00}+T_{01}} \\ \hat{\pi}_{11} &= \frac{T_{11}}{T_{10}+T_{11}} \\ \hat{\pi}_{00} &= 1-\hat{\pi}_{01} = \frac{T_{00}}{T_{00}+T_{01}} \\ \hat{\pi}_{10} &= 1-\hat{\pi}_{11} = \frac{T_{10}}{T_{10}+T_{11}}\end{aligned} \tag{3.17}$$

Which yields the matrix:

$$\Pi_1 = \begin{bmatrix} \hat{\pi}_{00} & \hat{\pi}_{01} \\ \hat{\pi}_{10} & \hat{\pi}_{11} \end{bmatrix} = \begin{bmatrix} 1-\hat{\pi}_{01} & \hat{\pi}_{01} \\ 1-\hat{\pi}_{11} & \hat{\pi}_{11} \end{bmatrix} = \begin{bmatrix} \frac{T_{00}}{T_{00}+T_{01}} & \frac{T_{01}}{T_{00}+T_{01}} \\ \frac{T_{10}}{T_{10}+T_{11}} & \frac{T_{11}}{T_{10}+T_{11}} \end{bmatrix} \tag{3.18}$$

In the above matrix dependence is allowed, ie. π_{01} does not have to equal π_{11}. However, under the null hypothesis of independence where $\pi_{01} = \pi_{11} = \pi$ the above matrix is instead written as:

$$\hat{\Pi} = \begin{bmatrix} 1-\hat{\pi} & \hat{\pi} \\ 1-\hat{\pi} & \hat{\pi} \end{bmatrix} \tag{3.19}$$

It is now possible to test the null hypothesis by creating the likelihood ratio test:

$$LR_{ind} = -2\ln\left(\frac{L(\hat{\pi})}{L(\hat{\Pi}_1)}\right) \tag{3.20}$$

Where $L(\hat{\pi})$ is the likelihood under the alternative hypothesis in the proportion of failure test:

$$L(\hat{\pi}) = (1-(T_1/(T_1+T_0))^{T_0}(T_1/(T_1+T_0))^{T_1} \tag{3.21}$$

The likelihood ratio constructed for the independence test also follows a chi-squared distribution with one degree of freedom as the sample size become larger. Depending on the number of observations it is possible to run into the situation where there are no consecutive failures ($T_{11} = 0$) and in that case the likelihood function used is instead:

$$L(\hat{\Pi}_1) = (1-\hat{\pi}_{01})^{T_{00}}\hat{\pi}_{01}^{T_{01}} \tag{3.22}$$

The null hypothesis is tested in the same way as in the proportion of failures test where the P-value is calculated:

$$\text{P-value} = 1 - F_{\chi_1^2}(LR_{ind}) \tag{3.23}$$

Where $F_{\chi_1^2}$ represents the cumulative probability distribution function for chi-squared with one degree of freedom. The null hypothesis is then rejected if P-value < 5%.

3.3.3 Conditional Coverage

Christoffersen (1998) also proposed a test to jointly test he proportion of failures and independence of a VaR model. In his paper he outlines a test that considers the likelihood ratio:

$$LR_{cc} = -2\ln\left(\frac{L(q)}{L(\hat{\Pi}_1)}\right) \tag{3.24}$$

Which is tested under the null hypothesis that $\pi_{01} = \pi_{11} = q$. In this case there are two hypothesis that are tested jointly. The correct distribution would therefore be a chi-squared distribution with *two* degrees of freedom. The above likelihood ratio uses both the likelihood function of the null hypothesis in the LR_{pof} test and the likelihood function of the alternative hypothesis in the LR_{ind} test. Thus it is possible to rewrite Equation 3.24 as:

$$\begin{aligned} LR_{cc} &= -2\ln\left(\frac{L(q)}{L(\hat{\Pi}_1)}\right) \\ &= -2\ln\left(\frac{L(q)}{L(\hat{\pi})}\frac{L(\hat{\pi})}{L(\hat{\Pi}_1)}\right) \\ &= -2\ln\left(\frac{L(q)}{L(\hat{\pi})}\right) - 2\ln\left(\frac{L(\hat{\pi})}{L(\hat{\Pi}_1)}\right) \\ &= LR_{pof} + LR_{ind} \end{aligned} \tag{3.25}$$

The test statistic for the joint hypothesis test can be calculated by just adding the test statistics of the proportion of failures test and the independence test. However, the critical value for rejecting the null hypothesis is in this case 5.9915 which corresponds to a right-tail probability of 5% when using two degrees of freedom. As in the previous backtesting sections, it is also possible to calculate the P-value of the test statistic:

$$\text{P-value} = 1 - F_{\chi_2^2}(LR_{cc}) \tag{3.26}$$

Where $F_{\chi_2^2}$ represents the cumulative probability distribution function for chi-squared with two degrees of freedom. The null hypothesis is then rejected if P-value < 5%.

4 Findings

4.1 GARCH Errors

By inspecting the autocorrelation and partial autocorrelation graphs seen in the appendix (Figure 16 to Figure 19), it is possible to see the autocorrelations in the squared residuals. The lag of importance is the first lag due to the GARCH(1,1) specification. If there are no autocorrelation at the first lag, the specified GARCH(1,1) process is unable to model the conditional variance since the series lack heteroskedasticity. By looking at the graphs it can be concluded that all series except GBP/USD seem to exhibit autocorrelation in the first lag. Furthermore, a Ljung-Box test is done in order to ultimately infer the significance of the autocorrelations.

Table 5: Ljung-Box autocorrelation test for the squared residuals.

	Currencies		**Commodities**	
	EUR/SEK	GBP/USD	CO1	LMCADS03
$L = 1$				
Q	10.0999	0.0244	14.6787	6.8061
Critical value	3.8415	3.8415	3.8415	3.8415
P-Value	0.0015***	0.8758	0.0001***	0.0091***
$L = 2$				
Q	27.6808	11.5667	19.3966	16.0029
Critical value	5.9915	5.9915	5.9915	5.9915
P-Value	0.0000***	0.0031***	0.0001***	0.0003***
$L = 3$				
Q	41.0521	11.5851	19.5664	16.0812
Critical value	7.8147	7.8147	7.8147	7.8147
P-Value	0.0000***	0.0089***	0.0002***	0.0011***
$L = 4$				
Q	44.4349	11.8894	32.4384	19.7541
Critical value	9.4877	9.4877	9.4877	9.4877
P-Value	0.0000***	0.0182**	0.0002***	0.0014***

L is the number of lags that are tested. Q is the test statistic, which is calculated according to Equation 2.11. The critical value is based on a 5% significance level and corresponds to the value with which it is a 5% probability of being in the right-tail of a chi-squared distribution with L degrees of freedom. The P-value is the probability of the null hypothesis being correct. The asterisks indicate the level of significance ranging from 10% at one asterisk (*), 5% at two asterisks (**) and 1% at three asterisks (***).

Seen from the Table 5 above, the Ljung-Box test yields the same results as the visual inspection. All lags up to four are significantly autocorrelated with the exception of the first lag in GBP/USD. As such it can be concluded that EUR/SEK, CO1 and LMCADS03 do exhibit GARCH errors. The time series mentioned are heteroskedastic, meaning that the conditional variance is time-variant and have periods with clustered volatility.

4.2 Return Distribution

The distribution of returns are tested for normality using the Jarque-Bera test. The P-values under the null hypothesis that the data is normally distributed can be found in the table below.

Table 6: Jarque-Bera normality test

	EUR/SEK	**GBP/USD**	**CO1**	**LMCADS03**
1 month	0.000	0.000	0.000	0.000
3 months	0.000	0.000	0.000	0.000
6 months	0.000	0.000	0.000	0.000
12 months	0.037	0.000	0.270	0.011

The numbers shown are the P-values of the corresponding test statistic calculated according to Equation 3.8 which is assumed to follow a chi-squared distribution with two degrees of freedom.

None of the four different return series can statistically be determined to follow a normal distribution. Another more fitting distribution would be a t-distribution. It can account for the excess kurtosis by choosing an adequate amount of degrees of freedom. However, the t-distribution is still symmetrical like the normal distribution and can as such not account for any skewness which could be present. The degrees of freedom further on is based on the calculations found in Table 1 to Table 4.

4.3 Backtesting - Normal Distribution

4.3.1 Proportion of Failures

Below in Table 7 the P-values for all the different methods are shown for the two currency pairs and the two commodities.

Table 7: Kupiec proportion of failures test P-values when using a normal distribution

	1 Month			**3 Months**			**6 Months**			**12 Months**		
	1%	5%	32%	1%	5%	32%	1%	5%	32%	1%	5%	32%
EUR/SEK												
EWMA	0.508	0.162	0.445	0.989	0.522	0.055	0.971	0.309	0.394	0.000	0.000	0.817
SMA	0.432	0.749	0.879	0.498	0.319	0.148	NaN	0.002	0.830	NaN	0.003	0.097
GARCH	0.287	0.821	0.416	0.696	0.561	0.155	NaN	0.063	0.368	0.714	0.201	0.192
GJR	0.287	0.883	0.416	0.279	0.561	0.288	0.679	0.894	0.589	0.714	0.678	0.453
IV	0.432	0.002	0.030	NaN	0.009	0.007	NaN	NaN	0.002	NaN	NaN	0.000
GBP/USD												
EWMA	0.668	0.679	0.699	0.012	0.464	0.821	0.000	0.000	0.188	0.000	0.000	0.001
SMA	0.043	0.326	0.805	0.003	0.010	0.609	0.001	0.181	0.927	0.000	0.001	0.494
GARCH	0.431	0.365	0.537	0.003	0.211	0.098	0.013	0.057	0.597	0.000	0.044	0.224
GJR	0.431	0.365	0.537	0.003	0.211	0.069	0.003	0.057	0.714	0.000	0.044	0.295
IV	0.668	0.416	0.647	0.668	0.464	0.070	0.011	0.845	0.507	0.010	0.246	0.156
CO1												
EWMA	0.072	0.729	0.102	0.000	0.164	0.030	0.000	0.000	0.545	0.000	0.000	0.706
SMA	0.001	0.172	0.581	0.000	0.022	0.187	0.000	0.000	0.146	0.000	0.000	0.036
GARCH	0.094	0.286	0.417	0.000	0.066	0.029	0.000	0.000	0.003	0.000	0.049	0.001
GJR	0.003	0.189	0.581	0.000	0.038	0.029	0.000	0.000	0.004	0.000	0.049	0.002
LMCADS03												
EWMA	0.230	0.507	0.252	0.225	0.522	0.338	0.000	0.000	0.096	0.000	0.000	0.001
SMA	0.014	0.574	0.823	0.001	0.031	0.848	0.000	0.005	0.519	0.000	0.000	0.384
GARCH	0.052	0.545	0.490	0.005	0.222	0.809	0.005	0.207	0.805	0.014	0.002	0.625
GJR	0.134	0.545	0.347	0.017	0.144	0.625	0.001	0.049	0.625	0.004	0.000	0.909

A P-value larger than 5% indicates that the volatility method passes the test with the right amount of failures.

Overall, most of the volatility methods manage to pass the proportion of failures test at a 60% success rate. However, the SMA only passes around 50% out of the 48 tests in total. While the implied volatility shows significantly better results within the GBP/USD data compared to the other methods, it performs worse than the rest within the EUR/SEK data. When it comes to the Brent crude oil, all methods perform rather poor with a success rate of 50% or less; with the SMA's performance at the bottom. For the copper data, GARCH performs the best while the SMA again performs the worst.In Figure 8 to Figure 11 the comparisons of observed failure rates to assumed failure rates for the four assets can be seen. Most of the volatility methods tend to overestimate the risk, with the largest exception being the EWMA greatly underestimating the risk at all significance levels in the longer horizons.

Figure 8: Comparison of the observed failure rate to the assumed failure level (α) for EUR/SEK using a normal distribution.

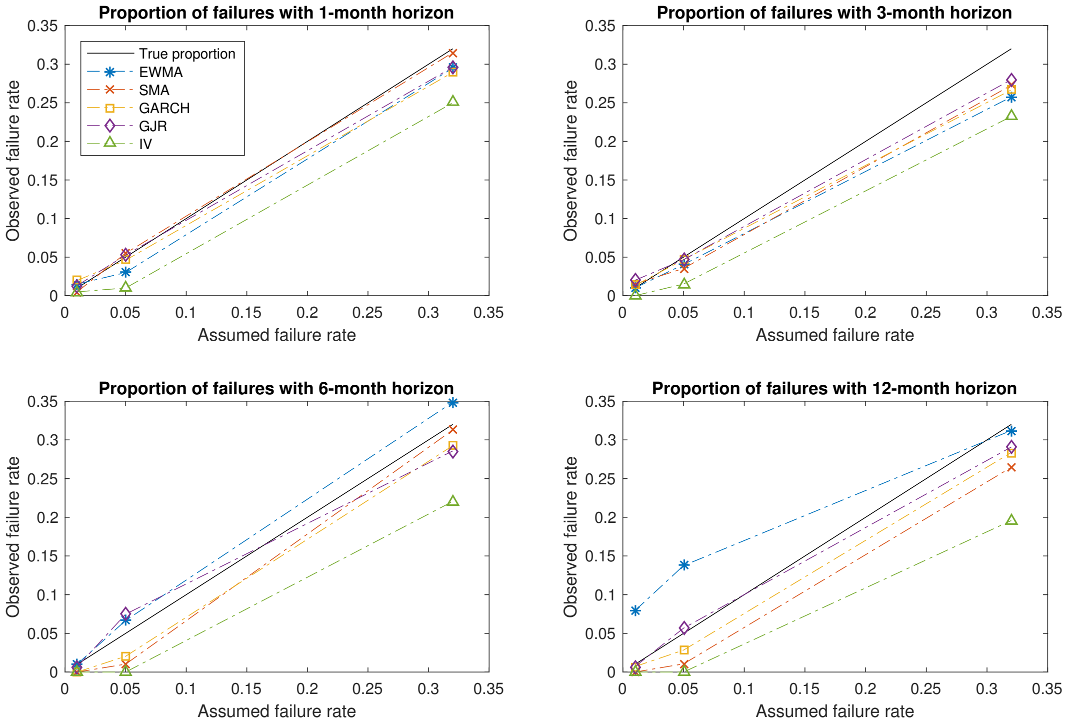

Figure 9: Comparison of the observed failure rate to the assumed failure level (α) for GBP/USD using a normal distribution.

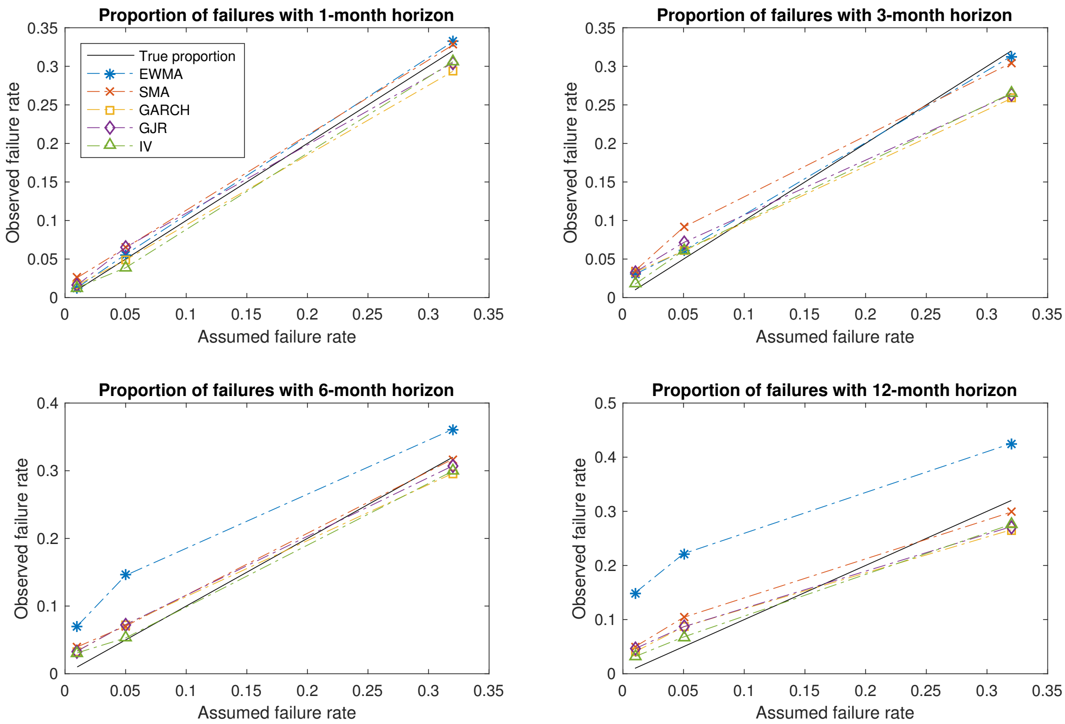

Figure 10: Comparison of the observed failure rate to the assumed failure level (α) for CO1 using a normal distribution.

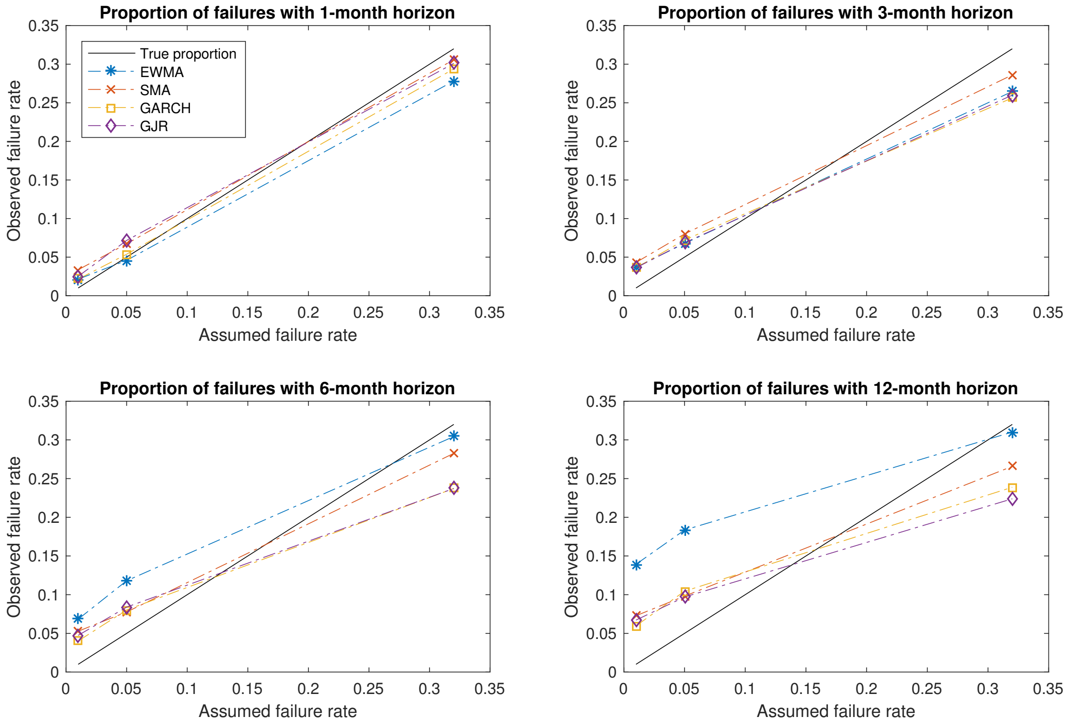

Figure 11: Comparison of the observed failure rate to the assumed failure level (α) for LMCADS03 using a normal distribution.

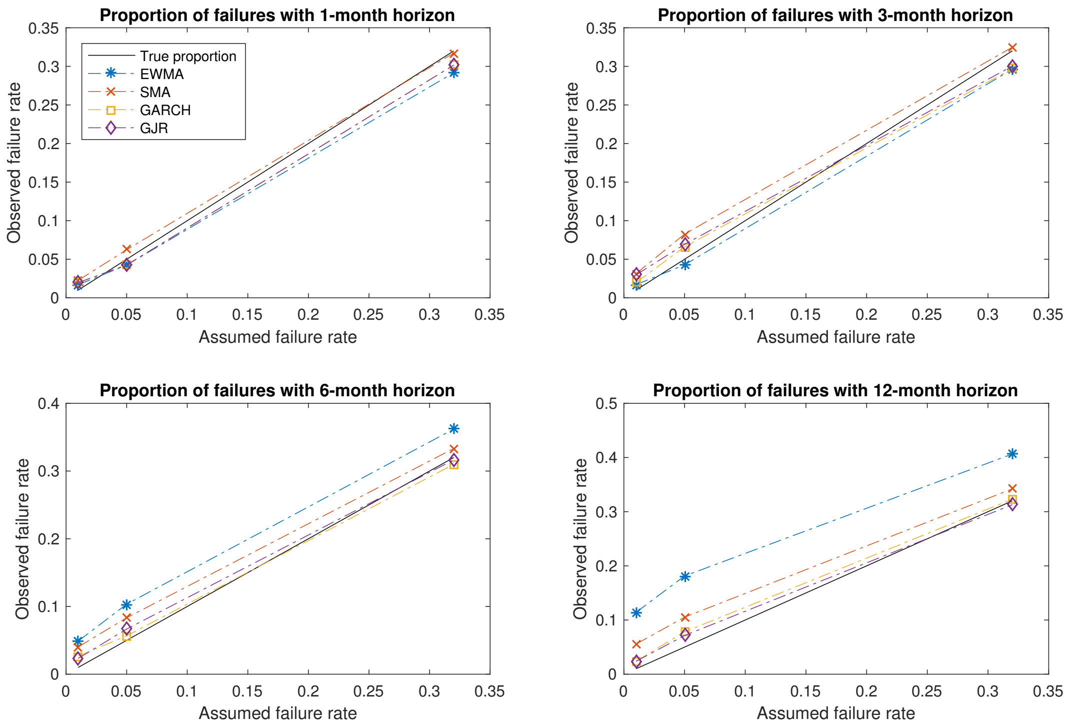

4.3.2 Independence

Table 8: Independence test P-values when using a normal distribution

	1 Month		
	1%	5%	32%
EUR/SEK			
EWMA	0.999	1.000	0.947
SMA	0.966	0.010	0.591
GARCH	0.999	0.307	0.259
GJR	0.999	0.418	0.469
IV	0.966	0.997	0.795
GBP/USD			
EWMA	0.999	1.000	0.430
SMA	1.000	1.000	0.553
GARCH	0.999	1.000	0.187
GJR	0.999	1.000	0.187
IV	0.999	1.000	0.956
CO1			
EWMA	0.128	0.025	0.007
SMA	0.003	0.002	0.008
GARCH	1.000	0.119	0.064
GJR	0.001	0.160	0.076
LMCADS03			
EWMA	1.000	0.661	0.825
SMA	1.000	0.432	0.087
GARCH	1.000	0.526	0.691
GJR	1.000	0.526	0.317

A P-value larger than 5% indicates independence.

Overall results show that it is within the 1-month horizon the independence is displayed in most cases. In Table 8, it can be observed that within the EUR/SEK, GBP/USD and the copper data, all the volatility estimation methods pass the tests at all significance levels at the 1-month horizon. However, within the Brent Crude Oil, not all methods pass the independence test. While the GARCH method passes the test at all significance levels during the 1-month horizon, the SMA fails at all significance levels. The GJR method manages to pass the test at the 5% and 32% level; the EWMA only passes the test at the 1% level during the 1-month horizon.

4.3.3 Conditional coverage

Table 9: Conditional coverage test P-values when using a normal distribution

	1 Month		
	1%	5%	32%
EUR/SEK			
EWMA	0.803	0.376	0.746
SMA	0.733	0.035	0.855
GARCH	0.567	0.578	0.380
GJR	0.567	0.712	0.552
IV	0.733	0.007	0.092
GBP/USD			
EWMA	0.912	0.918	0.680
SMA	0.129	0.618	0.814
GARCH	0.733	0.663	0.347
GJR	0.733	0.663	0.347
IV	0.912	0.718	0.899
CO1			
EWMA	0.062	0.077	0.007
SMA	0.000	0.003	0.024
GARCH	0.246	0.168	0.200
GJR	0.001	0.157	0.179
LMCADS03			
EWMA	0.486	0.729	0.501
SMA	0.050	0.627	0.225
GARCH	0.151	0.681	0.728
GJR	0.325	0.681	0.390

A P-value larger than 5% indicates that the volatility method yields both a correct number of failures and independence.

For the EUR/SEK, most of the volatility estimation methods manage to yield both a correct number of failures and independence. However, EWMA and the implied volatility did not perform satisfactory at the 5% significance level during the 1-month horizon. Within the GBP/USD data, all volatility methods pass all the tests during the 1-month horizon. For the Brent crude oil, the GARCH method passes at all significance levels during the 1-month horizon. The SMA on the other hand did not pass at any significance levels. The EWMA manages to pass at the 1% and the 5% level and the GJR passes at the 5% and the 32% levels during the 1-month horizon. When it comes to copper, all the methods perform satisfactory and passes the conditional coverage test at all significance levels during the 1-month horizon.

4.4 Backtesting - Student's t-distribution

4.4.1 Proportion of Failures

Table 10: Kupiec proportion of failures test P-values when using the best fitting t-distribution

	1 Month			3 Months			6 Months			12 Months		
	1%	5%	32%	1%	5%	32%	1%	5%	32%	1%	5%	32%
EUR/SEK												
EWMA	NaN	0.305	0.648	NaN	0.522	0.148	0.971	0.309	0.394	0.000	0.000	0.817
SMA	0.432	0.749	0651	0.498	0.522	0.411	NaN	0.010	0.927	NaN	0.003	0.133
GARCH	0.287	0.821	0.911	0.662	0.850	0.377	NaN	0.167	0.297	0.714	0.201	0.350
GJR	0.709	0.883	0.911	0.279	0.561	0.725	0.679	0.105	0.224	0.714	0.469	0.453
IV	0.432	0.002	0.220	NaN	0.009	0.026	NaN	NaN	0.015	NaN	0.000	0.000
GBP/USD												
EWMA	0.668	0.326	0.279	0.308	0.311	0.536	0.000	0.000	0.083	0.000	0.000	0.001
SMA	0.125	0.326	0.505	0.041	0.005	0.369	0.011	0.017	0.537	0.000	0.001	0.804
GARCH	0.431	0.790	0.985	0.014	0.531	0.317	0.013	0.192	0.964	0.002	0.044	0.378
GJR	0.431	0.130	0.860	0.014	0.211	0.317	0.013	0.192	0.783	0.002	0.044	0.474
IV	0.829	0.416	0.598	0.658	0.464	0.712	0.039	0.623	0.960	0.010	0.151	0.494
CO1												
EWMA	0.174	0.929	0.429	0.026	0.022	0.071	0.000	0.000	0.815	0.000	0.000	0.989
SMA	0.003	0.111	0.780	0.001	0.006	0.338	0.000	0.035	0.545	0.000	0.001	0.036
GARCH	0.232	0.582	0.581	0.003	0.038	0.055	0.001	0.034	0.009	0.000	0.000	0.003
GJR	0.094	0.118	0.673	0.001	0.038	0.055	0.000	0.018	0.020	0.000	0.000	0.001
LMCADS03												
EWMA	0.806	0.507	0.729	0.452	0.913	0.675	0.000	0.000	0.019	0.000	0.000	0.000
SMA	0.102	0.211	0.493	0.014	0.018	0.324	0.001	0.003	0.059	0.000	0.000	0.120
GARCH	0.304	0.948	0.864	0.298	0.144	0.802	0.290	0.440	0.622	0.272	0.040	0.540
GJR	0.304	0.948	0.591	0.298	0.054	0.537	0.290	0.133	0.460	0.272	0.068	0.716

A P-value larger than 5% indicates that the volatility method yields the right number of observed failures as the assumed ones.

In general, most of the volatility estimation models performed quite well for the EUR/SEK data; GARCH and GJR pass all the tests at all significance levels and horizons. However, the performance of the implied volatility is rather weak within the EUR/SEK while it shows superior performance within the GBP/USD, with a passing level of more than 80%.

For Brent crude oil, the performances of the volatility methods are weaker than for the other assets, with a large amount of failures in the 3-months to 12-months horizon. The SMA is the only method not passing all of the tests in the 1-month horizon.

When it comes to copper, the GJR performs the best and passes all of the tests, followed by the GARCH only failing at the 5% level in the 12-months horizon. The EWMA performs satisfactory within the 1-month and 3-months horizon but starts to fail after that. The SMA performs rather bad and passes only 50% of the tests.

Seen in Figure 12 to Figure 15 are the comparison of the observed failure rates to the assumed failure rates for the four assets using a t-distribution.

Figure 12: Comparison of the observed failure rate to the assumed failure level (α) for EUR/SEK using the best fitting t-distribution.

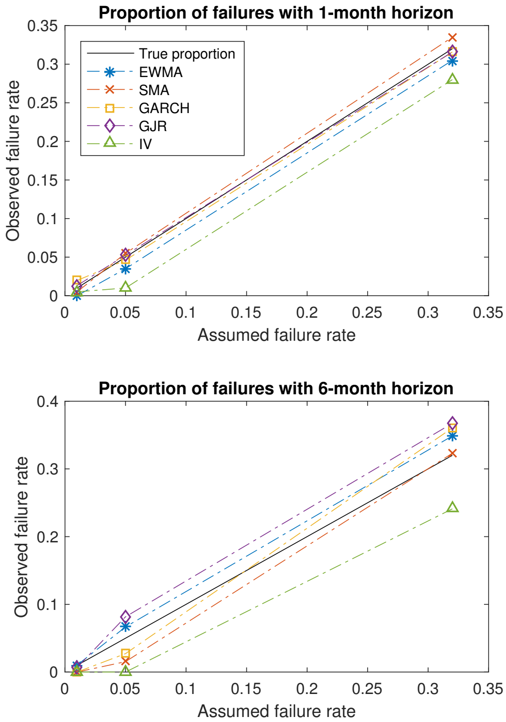

Figure 13: Comparison of the observed failure rate to the assumed failure level (α) for GBP/USD using the best fitting t-distribution.

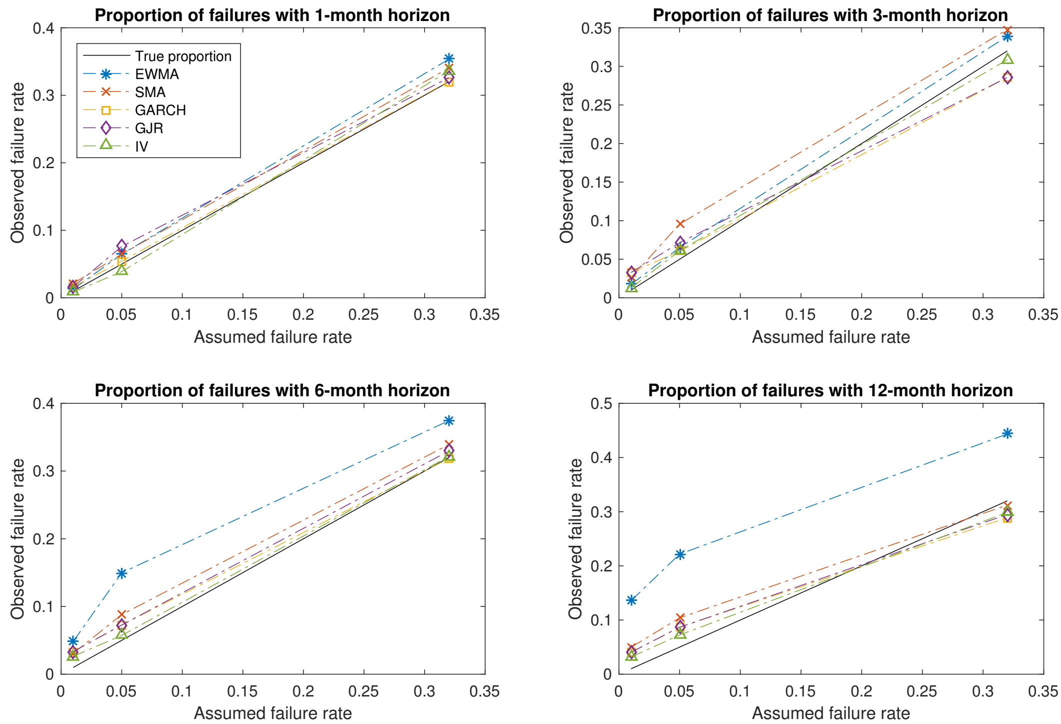

Figure 14: Comparison of the observed failure rate to the assumed failure level (α) for CO1 using the best fitting t-distribution.

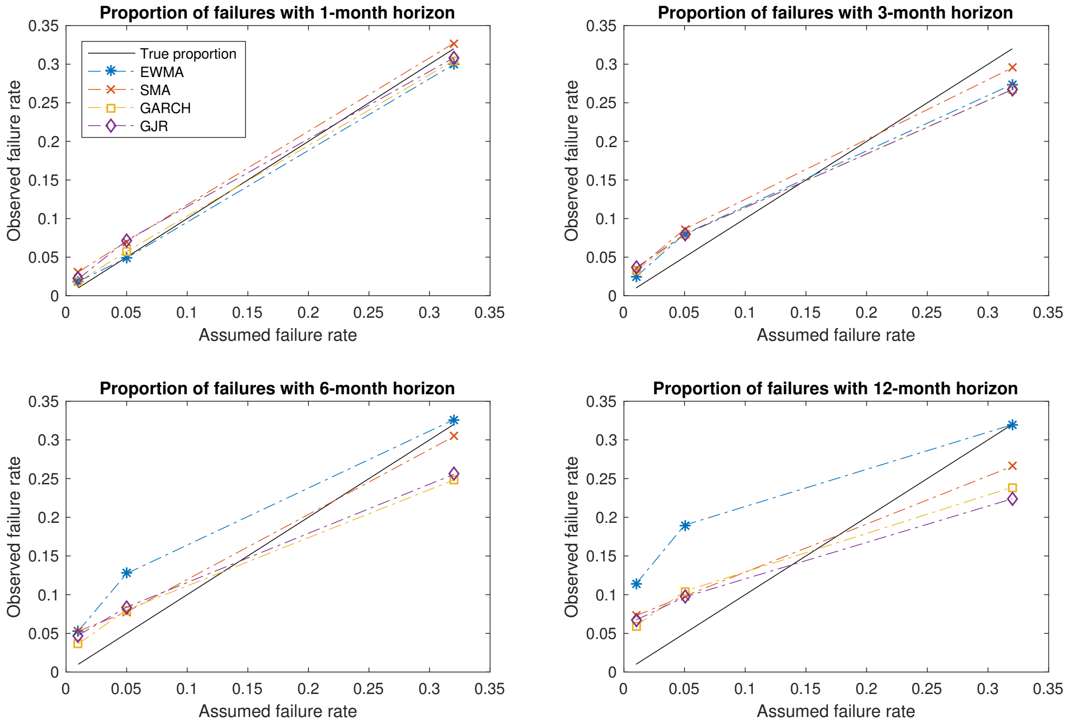

Figure 15: Comparison of the observed failure rate to the assumed failure level (α) for LMCADS03 using the best fitting t-distribution.

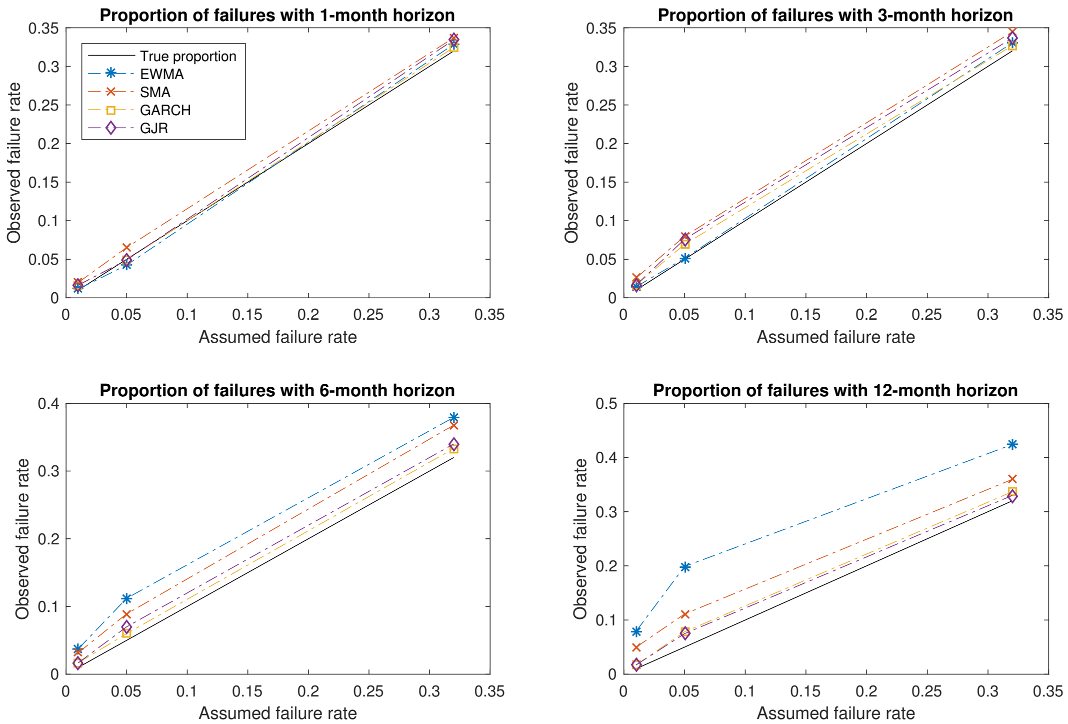

4.4.2 Independence

Table 11: Independence test P-values when using the best fitting t- distribution

	1 Month		
	1%	5%	32%
EUR/SEK			
EWMA	1.000	1.000	0.640
SMA	0.966	0.010	0.175
GARCH	0.999	0.307	0.104
GJR	0.996	0.418	0.635
IV	0.996	0.997	0.866
GBP/USD			
EWMA	0.999	1.000	0.720
SMA	1.000	1.000	0.698
GARCH	0.999	1.000	0.265
GJR	0.999	0.940	0.059
IV	0.997	1.000	0.768
CO1			
EWMA	0.088	0.038	0.010
SMA	0.002	0.003	0.040
GARCH	1.000	0.293	0.135
GJR	0.105	0.211	0.111
LMCADS03			
EWMA	1.000	0.661	0.252
SMA	1.000	0.678	0.171
GARCH	1.000	0.713	0.586
GJR	1.000	0.168	0.291

A P-value larger than 5% indicates that the volatility method passes the independence test.

Overall, most of the volatility methods display independence in the 1-month horizon. However, the SMA and the EWMA did not pass all of the tests. The EWMA fails at the 5% and the 32% level for Brent crude oil. The SMA not only failed at the 5% level for EUR/SEK but also failed at all significance levels for Brent crude oil.

4.4.3 Conditional Coverage

Table 12: Conditional coverage test P-values when using the best fitting t-distribution

	1 Month		
	1%	5%	32%
EUR/SEK			
EWMA	1.000	0.591	0.807
SMA	0.733	0.035	0.359
GARCH	0.567	0.578	0.265
GJR	0.933	0.712	0.881
IV	0.733	0.007	0.464
GBP/USD			
EWMA	0.912	0.618	0.522
SMA	0.309	0.618	0.743
GARCH	0.733	0.965	0.538
GJR	0.733	0.318	0.166
IV	0.977	0.718	0.833
CO1			
EWMA	0.093	0.116	0.027
SMA	0.000	0.003	0.118
GARCH	0.489	0.494	0.282
GJR	0.066	0.134	0.258
LMCADS03			
EWMA	0.970	0.729	0.489
SMA	0.262	0.420	0.310
GARCH	0.590	0.932	0.850
GJR	0.590	0.381	0.496

A P-value larger than 5% indicates that the volatility method passes the conditional coverage test.

Overall, the volatility estimation methods performed quite all right, with most of them passing all of the tests at the 1-month horizon. However, when it comes to the EUR/SEK data the implied volatility method fails to pass the test at the 5% level. For Brent crude oil, the SMA fails the tests at both the 1% and 5% levels while the EWMA fails at the 32% level during the 1-month horizon.

5 Discussion

Overall, the findings from the proportion of failures test were unsatisfactory when using a normal distribution. These findings were consistent for most volatility estimation techniques, although some methods were accurate in very specific cases. Generally no volatility estimation technique can be said to be superior to another. Only approximately 60% of the tests showed significant results in the proportion of failures test. However, the SMA method performed much worse than any other method both in the proportion of failures test as well as in the independence test with a success rate of only around 52%. It was especially weak when applied to the Brent crude oil (CO1) and GBP/USD with a success rate of 33% and 42% respectively. The major drawback with the SMA is its inability to handle volatility shocks. A volatility shock will only be partly incorporated into the SMA estimation in addition to persisting for the length of the estimation, which is one year in this case. These features make the SMA a simple, but inaccurate estimator of volatility, which do not justify any practical application. When using a normal distribution, the GARCH volatility estimation method provided the highest accuracy with a success rate of 67%. While it showed low accuracy for the Brent crude oil, it showed perfect accuracy when applied to EUR/SEK. However, the CO1 return series proved extremely hard to model using our parametric methods. Considering all volatility techniques as well as both the normal distribution and t-distribution, the success rate was a merely 35%. The reason for this is unclear. The CO1 series showed no more skewness or kurtosis than the other series. The fitted t-distribution showed to be a weak fit for the left tail as seen in Figure 22. The reason could be the oligiopoly present in the oil market, where one large participant such as OPEC can quickly influence the prices by changing their oil production.

The implied volatility performed poorly for EUR/SEK where it overestimated the volatility level and consequently the VaR level as well. When applied to GBP/USD it performed better with only a slight overestimation of VaR at the 32% level. The overestimation of the implied volatility has been found in other assets such as the S&P100 index (Fleming, 1998). In theory, the implied volatility should produce superior results as it can both incorporate historical data as well as future events that may affect the volatility level. Research into the predictive nature of the implied volatility has shown that it has a superior predictive power compared to historical volatility measures for the S&P indices (Chiras & Manaster, 1978; Day & Lewis, 1992; Becker et al., 2006; Christensen & Prabhala, 1998). Similar results have been presented for currency options and commodity futures option as well (Xu & Taylor, 1995; Szakmary et al., 2003). The opinions differ to some extent with Canina & Figlewski (1993) finding no evidence for superior predictability of the implied volatility from S&P100 index options. On the other hand, regarding EUR/SEK, perhaps an inefficient option market could explain the lack of accuracy of the implied volatility found in this study. Since an inefficient market do not reflect the market's aggregated expectation of the future volatility it is bound to have less predictive power. Most of the research regarding the efficiency of option markets are carried out on very liquid options such as the S&P indices and USD denominated currencies. In less liquid options such as the EUR/SEK options (Bank for International Settlements, 2016), the market efficiency might be lower which would explain the inaccuracy when applied in the VaR calculations. For GBP/USD options, which is much more liquid, the accuracy is almost perfect.

Another interesting note is the improved fit when combining the implied volatility with a t-distribution. The volatility smile, which is present for most options, suggests that the implied volatility is lower for at-the-money options compared to options with high/lower strike prices. An explanation for this behavior is that the market have incorporated fatter tails for the distribution of the returns on the underlying assets. The market's aggregate assumption is that tail events are more likely to occur than implied by a normal distribution (Heston, 1993; Jackwerth & Rubinstein, 1996). In other words, the leptokurtic and skewed behavior of financial return series would give rise to a volatility smile effect when using the Black-Scholes model, and its assumption of normality, to compute for the implied volatilities. The implied volatility is thus quoted with a non-normal distribution in mind. It would hence be more intuitive to use a non-normal distribution with excess kurtosis, such as the t-distribution, for calculating the return quantiles when the implied volatility is used. An interesting topic of future research would be to explore the link between the liquidity of options and the predictive power of the implied volatility.

The results from the study do in fact point toward an improvement when using a t-distribution, albeit just a minor refinement. An improvement can be seen regardless of which volatility estimation methods used. One of the outcomes from using a t-distribution is a potential change in the observed failures. For probability levels further out in the left tail the observed proportion of failures will decrease to a greater extent than for probability levels closer to the mean. This can be explained by the larger difference between a 1% t-distribution (-3.14 at 6 degrees of freedom) and a 1% normal distribution (-2.33) than between a 32% t-distribution (-0.49 at 6 degrees of freedom) and a 32% normal distribution (-0.47). The result is in theory a steeper slope in graphs similar to Figure 8 - Figure 15. A caveat with using the t-distribution is that while it can easily incorporate excess kurtosis, it is symmetrical and thus cannot account for the skewness seen in many financial return series such as the data used in this thesis. In Table 1 - Table 4 significant skewness are found in all four datasets. This skewness will affect the fitting procedure for determining the optimal t-distribution. Due to having to account for the shape of both tails, the optimization procedure produces biased estimates, especially in the cases where only one tail is of interest. A solution to this problem would be to individually model the left and right tail by "mirroring" the negative and positive returns around the zero (with the assumption that the mean return equals zero). The same tail shape would then be present in both the right and left tails for the two new distributions. Another option is to fit a distribution to the data instead of trying to fit the data to a distribution. This sort of procedure is know as the extreme value theory (EVT). McNeil & Frey (2000) found that the EVT procedure provides higher accuracy when applied to GBP/USD, but also among other assets as well.

6 Conclusion

The leptokurtic nature of most financial returns have been a long known fact. Still, most parametric VaR methods assume that data is normally distributed. This practice will greatly underestimate the probability of tail-events, which should produce inaccurate VaR estimations. In this study, the leptokurtic t-distribution is explored as an alternative to the normal distribution in parametric VaR calculations. A t-distribution is fitted to the data in order to produce a distribution more true to reality.

The results indicate slight improvements when using a t-distribution, although the improvements were less than expected. Since the t-distribution only incorporates the excess kurtosis it is bound to yield biased results due to the skewness that is present in most financial time series. The results displayed no discernible difference between the asset classes. However, some variations were found *within* the commodities. The VaR measure was much more accurate for the LMCADS03 compared to the CO1. A possible explanation could be the event risk present in the oil market. Also, not surprisingly, the VaR models lose accuracy at longer horizons due to the time-variation of the volatility. The resulting VaR measures are practically useless on horizons longer than 3 months.

The area in which the t-distribution prevails is when considering high and low probability levels. At the 5% level, the normal distribution and the t-distribution produces similar critical values. As the probability level increases or decreases, the difference between the distributions become larger. As shown in this study, this behavior provides greater precision of the VaR measure since the normal distribution tend to overestimate the risk at high probability levels and underestimate the risk at lower probability levels. The implementation of a t-distribution is straightforward and can easily be applied in a simple spreadsheet.

7 Appendices

7.1 Autocorrelation Plots

Figure 16: Autocorrelation and partial autocorrelation plots for both the return series and the squared return series of the monthly EUR/SEK sample data.

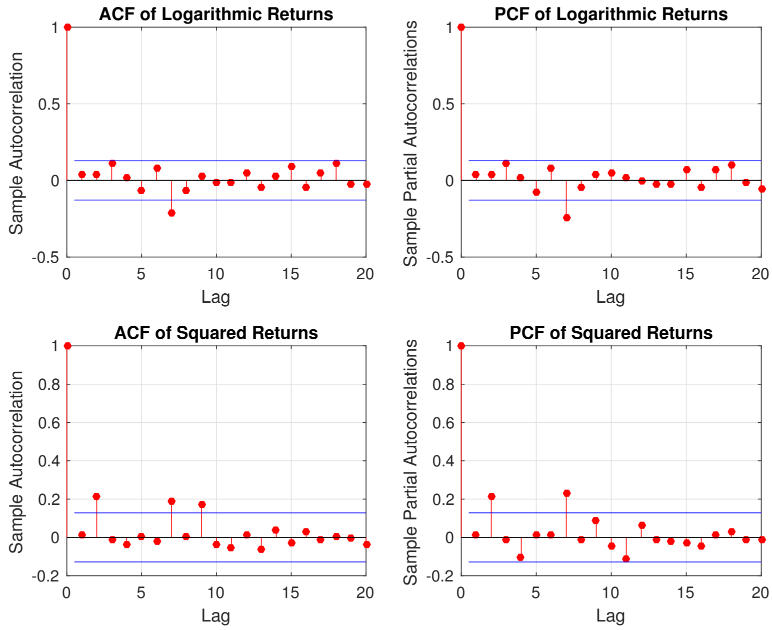

Figure 17: Autocorrelation and partial autocorrelation plots for both the return series and the squared return series of the monthly GBP/USD sample data.

Figure 18: Autocorrelation and partial autocorrelation plots for both the return series and the squared return series of the monthly CO1 sample data.

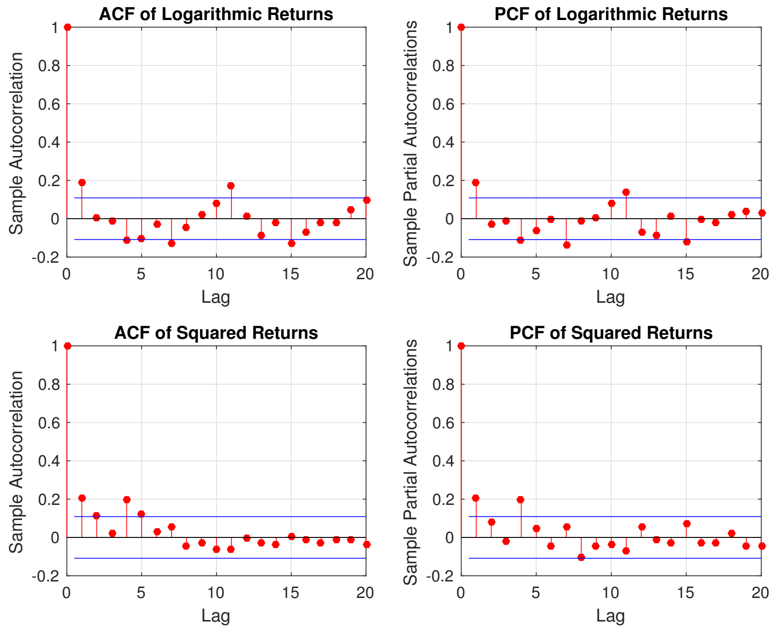

Figure 19: Autocorrelation and partial autocorrelation plots for both the return series and the squared return series of the monthly LMCADS03 sample data.

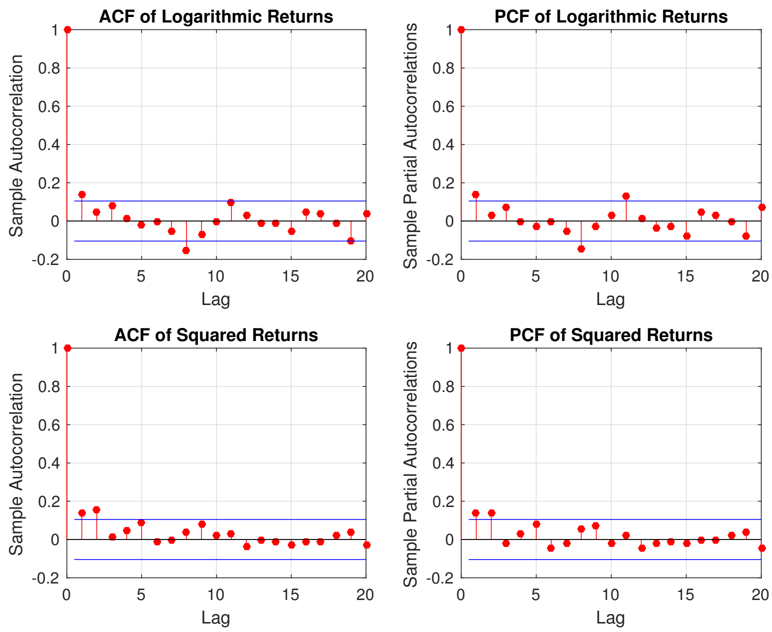

7.2 QQ Plots

Figure 20: QQ plots of EUR/SEK sample log-returns for different periods and distributions.

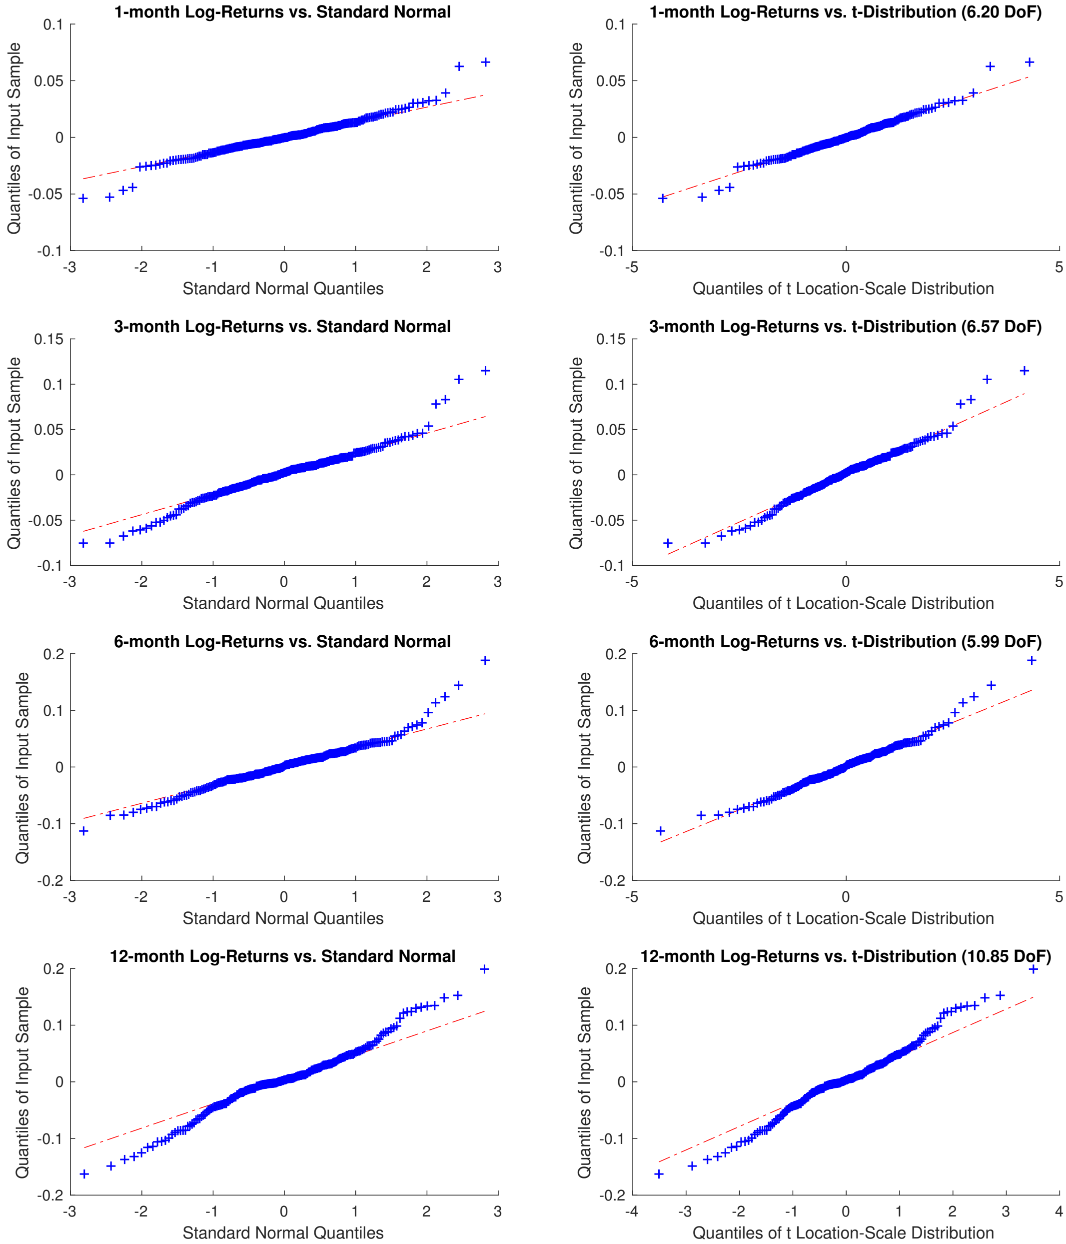

Figure 21: QQ plots of GBP/USD sample log-returns for different periods and distributions.

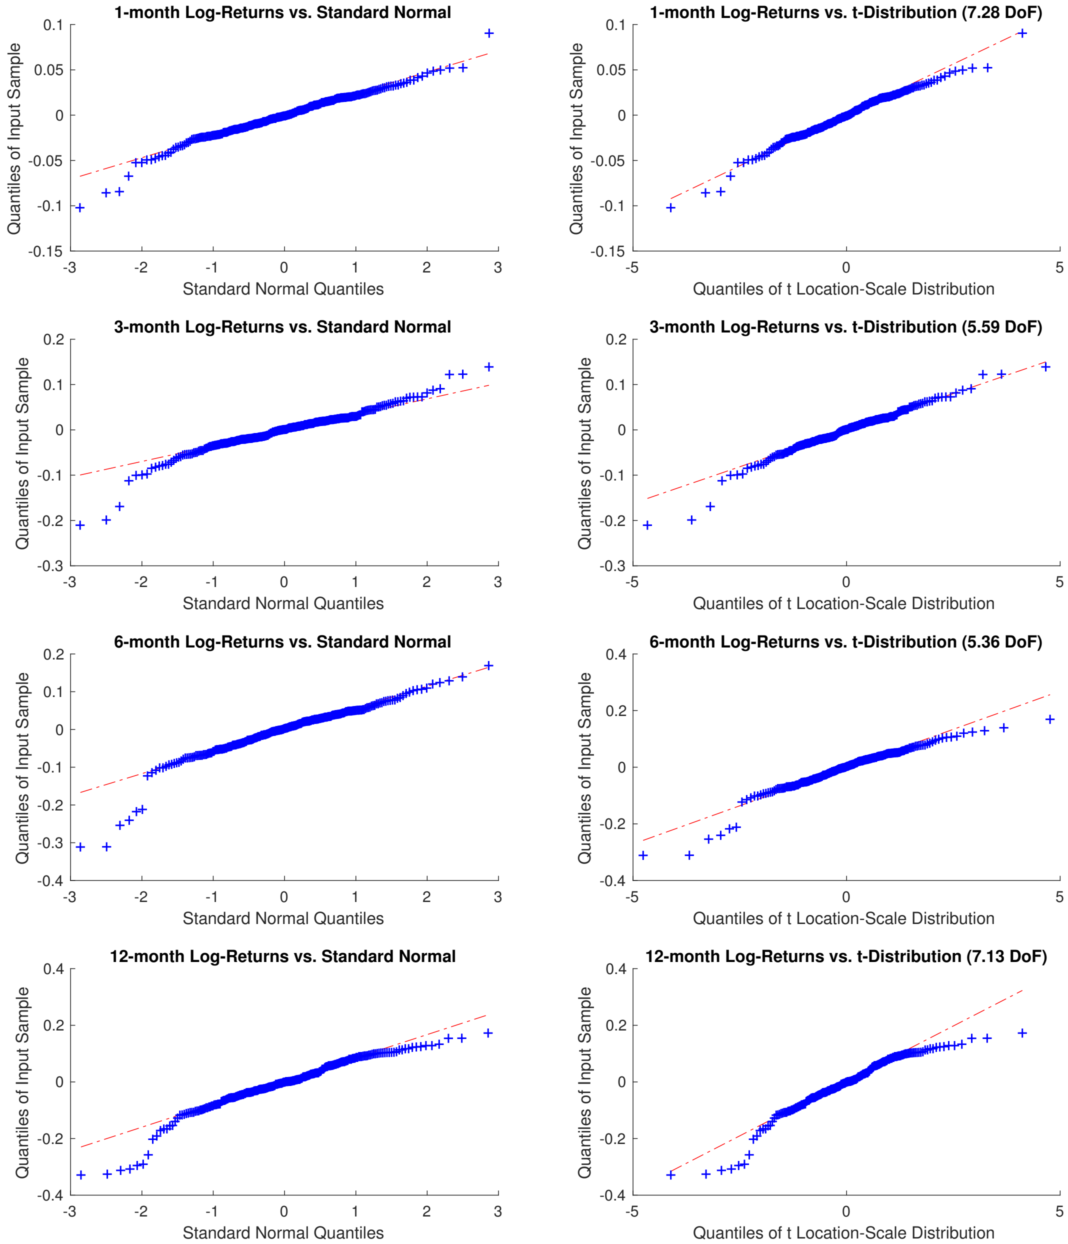

Figure 22: QQ plots of CO1 sample log-returns for different periods and distributions.

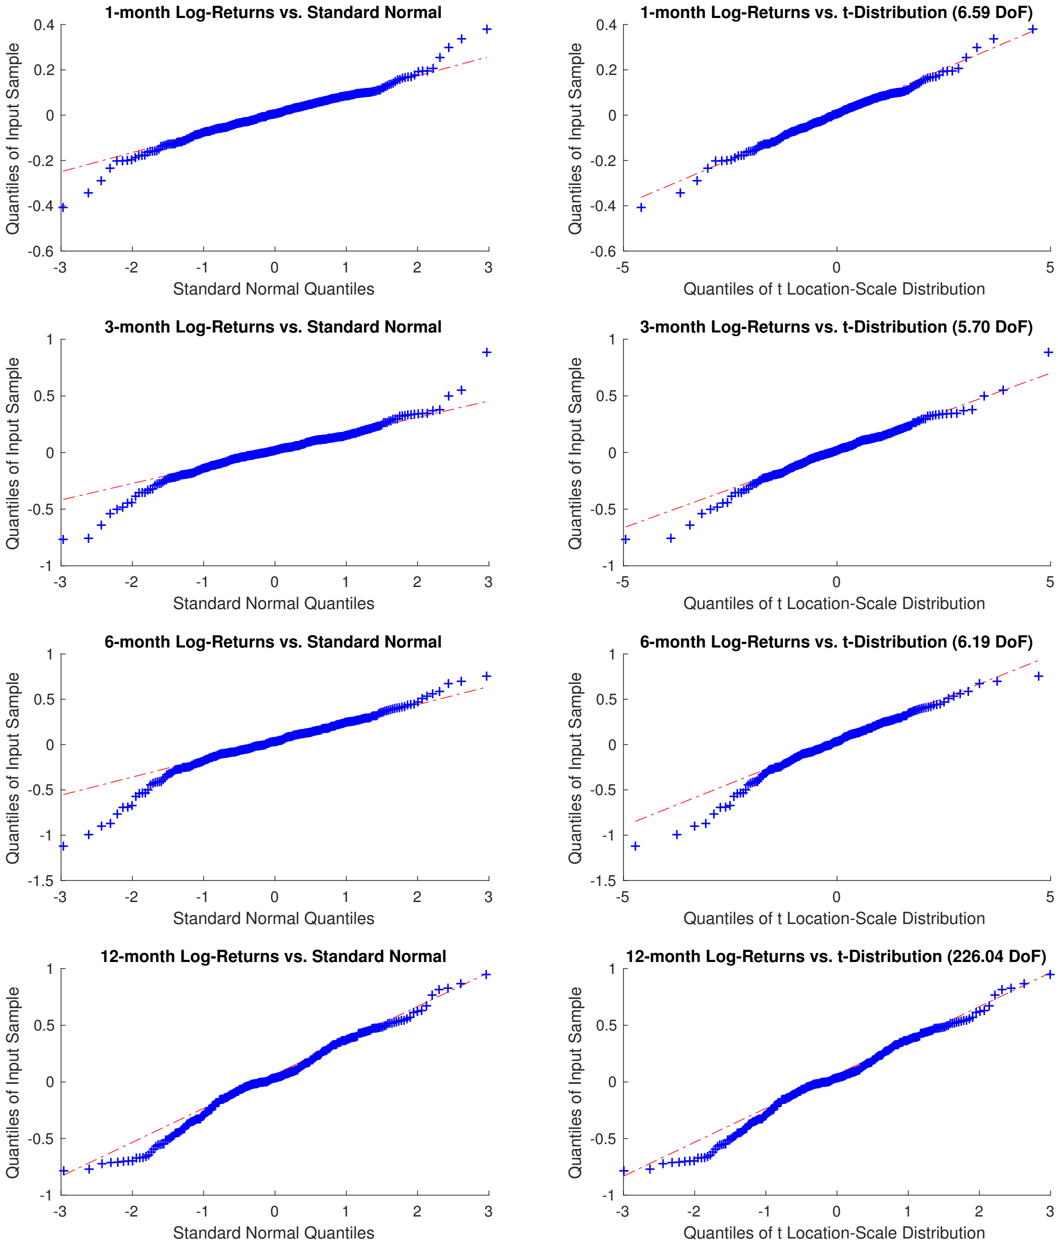

Figure 23: QQ plots of LMCADS03 sample log-returns for different periods and distributions.

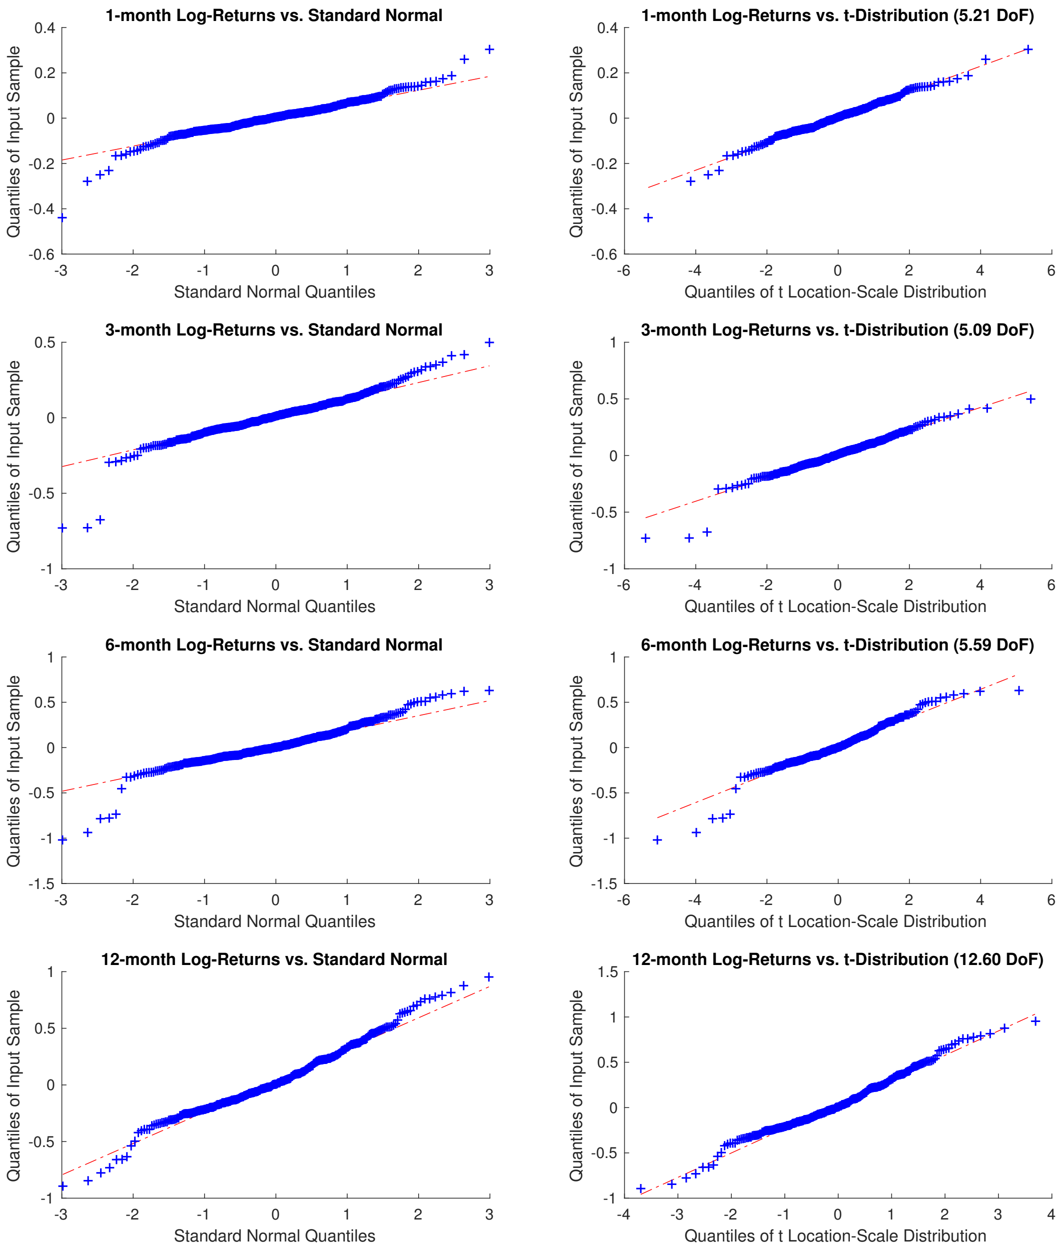

7.3 Backtesting

7.3.1 Normal Distribution

Table 13: Backtesting EUR/SEK for 1% VaR

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	200	200	152	152	200	200	200
π	1.50%	0.50%	1.97%	1.97%	0.50%	1.00%	1.00%
LR_{pof}	0.5082	0.4315	0.2869	0.2869	0.4315	1.0000	1.0000
LR_{ind}	0.9997	0.9656	0.9995	0.9995	0.9656	0.9970	0.9970
LR_{cc}	0.8034	0.7332	0.5672	0.5672	0.7332	1.0000	1.0000
3 Months							
n	198	198	150	150	198	198	198
π	1.01%	1.52%	1.33%	2.00%	0.00%	1.52%	1.52%
LR_{pof}	0.9886	0.4984	0.6962	0.2786	1.0000	0.4984	0.4984
LR_{ind}	0.9653	0.0144	0.0110	0.0356	1.0000	0.0144	0.0144
LR_{cc}	0.9990	0.0398	0.0365	0.0611	1.0000	0.0398	0.0398
6 Months							
n	195	195	147	147	195	195	195
π	1.03%	0.00%	0.00%	0.68%	0.00%	0.00%	0.00%
LR_{pof}	0.9714	1.0000	1.0000	0.6792	1.0000	1.0000	1.0000
LR_{ind}	0.9969	1.0000	1.0000	0.9598	1.0000	1.0000	1.0000
LR_{cc}	0.9994	1.0000	1.0000	0.9169	1.0000	1.0000	1.0000
12 Months							
n	189	189	141	141	189	189	189
π	7.94%	0.00%	0.71%	0.71%	0.00%	0.00%	0.00%
LR_{pof}	0.0000	1.0000	0.7143	0.7143	1.0000	1.0000	1.0000
LR_{ind}	0.0000	1.0000	0.9589	0.9589	1.0000	1.0000	1.0000
LR_{cc}	0.0000	1.0000	0.9339	0.9339	1.0000	1.0000	1.0000

Table 14: Backtesting EUR/SEK for 5% VaR

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	200	200	152	152	200	200	200
π	3.00%	5.50%	4.61%	5.26%	1.00%	3.00%	3.00%
LR_{pof}	0.1622	0.7493	0.8211	0.8826	0.0017	0.1622	0.1622
LR_{ind}	1.0000	0.0101	0.3067	0.4177	0.9970	0.1545	0.1545
LR_{cc}	0.3765	0.0349	0.5781	0.7122	0.0071	0.1366	0.1366
3 Months							
n	198	198	150	150	198	198	198
π	4.04%	3.54%	4.00%	4.00%	1.52%	5.05%	4.55%
LR_{pof}	0.5222	0.3193	0.5610	0.5610	0.0087	0.9740	0.7658
LR_{ind}	0.0009	0.0003	0.2150	0.0130	0.0144	0.0051	0.0022
LR_{cc}	0.0031	0.0008	0.3915	0.0387	0.0160	0.0199	0.0090
6 Months							
n	195	195	147	147	195	195	195
π	6.67%	1.03%	2.04%	4.76%	0.00%	2.05%	2.05%
LR_{pof}	0.3085	0.0021	0.0630	0.8938	1.0000	0.0329	0.0329
LR_{ind}	0.0000	0.9969	0.9995	0.0013	1.0000	0.0568	0.0568
LR_{cc}	0.0000	0.0087	0.1775	0.0056	1.0000	0.0168	0.0168
12 Months							
n	189	189	141	141	189	189	189
π	13.76%	1.06%	2.84%	4.26%	0.00%	2.12%	2.12%
LR_{pof}	0.0000	0.0027	0.2010	0.6775	1.0000	0.0408	0.0408
LR_{ind}	0.0000	0.0085	0.0016	0.0000	1.0000	0.0009	0.0009
LR_{cc}	0.0000	0.0003	0.0031	0.0000	1.0000	0.0005	0.0005

Table 15: Backtesting EUR/SEK for 32% VaR

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	200	200	152	152	200	200	200
π	29.50%	31.50%	28.95%	28.95%	25.00%	30.00%	31.00%
LR_{pof}	0.4452	0.8793	0.4156	0.4156	0.0300	0.5420	0.7611
LR_{ind}	0.9466	0.5905	0.2590	0.4688	0.7948	0.9431	0.7379
LR_{cc}	0.7455	0.8553	0.3797	0.5523	0.0918	0.8282	0.9029
3 Months							
n	198	198	150	150	198	198	198
π	25.76%	27.27%	26.67%	28.00%	23.23%	26.26%	26.26%
LR_{pof}	0.0548	0.1480	0.1546	0.2879	0.0065	0.0781	0.0781
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	195	195	147	147	195	195	195
π	34.87%	31.28%	28.57%	29.93%	22.05%	26.67%	26.67%
LR_{pof}	0.3935	0.8295	0.3680	0.5887	0.0021	0.1046	0.1046
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	31.22%	26.46%	26.95%	29.08%	24.87%	26.46%	25.40%
LR_{pof}	0.8171	0.0965	0.1918	0.4531	0.0001	0.0469	0.0315
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 16: Backtesting GBP/USD for 1% VaR

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	232	232	184	184	232	232	232
π	1.29%	2.59%	1.63%	1.63%	1.29%	1.29%	1.29%
LR_{pof}	0.6677	0.0428	0.4309	0.4309	0.6677	0.6677	0.6677
LR_{ind}	0.9997	1.0000	0.9996	0.9996	0.9997	0.9997	0.9997
LR_{cc}	0.9120	0.1286	0.7333	0.7333	0.9120	0.9120	0.9120
3 Months							
n	230	230	182	182	230	230	230
π	3.04%	3.48%	3.85%	3.85%	1.74%	1.74%	1.74%
LR_{pof}	0.0122	0.0032	0.0033	0.0033	0.3079	0.3079	0.3079
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	227	227	179	179	227	227	227
π	7.05%	3.96%	3.35%	3.91%	3.08%	2.64%	2.64%
LR_{pof}	0.0000	0.0007	0.0128	0.0030	0.0114	0.0389	0.0389
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	14.93%	4.98%	5.78%	5.78%	3.17%	3.17%	3.17%
LR_{pof}	0.0000	0.0000	0.0000	0.0000	0.0098	0.0098	0.0098
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 17: Backtesting GBP/USD for 5% VaR

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	232	232	184	184	232	232	232
π	5.60%	6.47%	6.52%	6.52%	3.88%	4.74%	4.74%
LR_{pof}	0.6788	0.3262	0.3646	0.3646	0.4157	0.8544	0.8544
LR_{ind}	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
LR_{cc}	0.9178	0.6175	0.6630	0.6630	0.7180	0.9835	0.9835
3 Months							
n	230	230	182	182	230	230	230
π	6.09%	9.13%	7.14%	7.14%	6.09%	5.65%	5.22%
LR_{pof}	0.4638	0.0096	0.2113	0.2113	0.4638	0.6564	0.8806
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	227	227	179	179	227	227	227
π	14.54%	7.05%	8.38%	8.38%	5.29%	5.29%	5.29%
LR_{pof}	0.0000	0.1811	0.0574	0.0574	0.8445	0.8445	0.8445
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	22.17%	10.41%	8.67%	8.67%	6.79%	6.79%	6.79%
LR_{pof}	0.0000	0.0012	0.0438	0.0438	0.2464	0.2464	0.2464
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 18: Backtesting GBP/USD for 32% VaR

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	232	232	184	184	232	232	232
π	33.19%	32.76%	29.89%	29.89%	30.60%	31.90%	31.90%
LR_{pof}	0.6986	0.8048	0.5373	0.5373	0.6471	0.9730	0.9730
LR_{ind}	0.4301	0.5532	0.1872	0.1872	0.9563	0.6254	0.4899
LR_{cc}	0.6796	0.8136	0.3465	0.3465	0.8992	0.8871	0.7874
3 Months							
n	230	230	182	182	230	230	230
π	31.30%	30.43%	26.37%	25.82%	26.52%	27.83%	28.70%
LR_{pof}	0.8207	0.6093	0.0978	0.0686	0.0701	0.1694	0.2780
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	227	227	179	179	227	227	227
π	36.12%	31.72%	30.17%	30.73%	29.96%	30.40%	30.40%
LR_{pof}	0.1875	0.9274	0.5973	0.7139	0.5066	0.6029	0.6029
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	42.53%	29.86%	27.75%	28.32%	27.60%	28.96%	28.96%
LR_{pof}	0.0010	0.4935	0.2244	0.2947	0.1555	0.3282	0.3282
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 19: Backtesting CO1 for 1% VaR

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	328	328	280	280	328	328
π	2.14%	3.36%	2.15%	3.23%	1.83%	1.83%
LR_{pof}	0.0719	0.0007	0.0939	0.0030	0.1742	0.1742
LR_{ind}	0.1283	0.0031	1.0000	0.0012	0.0878	0.0878
LR_{cc}	0.0623	0.0000	0.2459	0.0010	0.0925	0.0925
3 Months						
n	326	326	278	278	326	326
π	3.69%	4.31%	4.33%	5.05%	2.77%	2.77%
LR_{pof}	0.0002	0.0000	0.0000	0.0000	0.0084	0.0084
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	323	323	275	275	323	323
π	6.83%	5.28%	4.01%	4.74%	3.42%	3.42%
LR_{pof}	0.0000	0.0000	0.0002	0.0000	0.0006	0.0006
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	317	317	269	269	317	317
π	13.92%	7.28%	7.46%	7.46%	5.06%	5.38%
LR_{pof}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 20: Backtesting CO1 for 5% VaR

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	328	328	280	280	328	328
π	4.59%	6.73%	6.45%	6.81%	5.50%	5.50%
LR_{pof}	0.7285	0.1723	0.2860	0.1875	0.6801	0.6801
LR_{ind}	0.0253	0.0016	0.1188	0.1603	0.0772	0.0772
LR_{cc}	0.0771	0.0027	0.1677	0.1566	0.1927	0.1927
3 Months						
n	326	326	278	278	326	326
π	6.77%	8.00%	7.58%	7.94%	6.15%	6.77%
LR_{pof}	0.1639	0.0219	0.0661	0.0377	0.3562	0.1639
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	323	323	275	275	323	323
π	11.80%	7.76%	8.39%	8.76%	6.83%	6.52%
LR_{pof}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	317	317	269	269	317	317
π	18.35%	9.81%	10.82%	10.45%	7.59%	7.59%
LR_{pof}	0.0000	0.0005	0.0485	0.0485	0.0001	0.0003
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 21: Backtesting CO1 for 32% VaR

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	328	328	280	280	328	328
π	27.83%	30.58%	29.75%	30.47%	29.36%	29.97%
LR_{pof}	0.1016	0.5808	0.4172	0.5812	0.3020	0.4285
LR_{ind}	0.0074	0.0080	0.0642	0.0764	0.0016	0.0021
LR_{cc}	0.0073	0.0236	0.1298	0.1787	0.0040	0.0064
3 Months						
n	326	326	278	278	326	326
π	26.46%	28.62%	25.99%	25.99%	27.38%	27.38%
LR_{pof}	0.0295	0.1865	0.0290	0.0290	0.0705	0.0705
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	323	323	275	275	323	323
π	30.43%	28.26%	23.72%	24.09%	26.40%	26.71%
LR_{pof}	0.5453	0.1458	0.0025	0.0040	0.0283	0.0386
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	317	317	269	269	317	317
π	31.01%	26.58%	22.76%	23.51%	25.63%	25.63%
LR_{pof}	0.7060	0.0358	0.0008	0.0022	0.0133	0.0133
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 22: Backtesting LMCADS03 for 1% VaR

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	354	354	306	306	354	354
π	1.70%	2.55%	2.30%	1.97%	1.70%	1.70%
LR_{pof}	0.2296	0.0144	0.0518	0.1338	0.2296	0.2296
LR_{ind}	1.0000	1.0000	1.0000	1.0000	0.0805	0.0805
LR_{cc}	0.4860	0.0500	0.1509	0.3249	0.1055	0.1055
3 Months						
n	352	352	304	304	352	352
π	1.71%	3.13%	2.97%	2.64%	2.28%	2.28%
LR_{pof}	0.2251	0.0013	0.0053	0.0172	0.0390	0.0390
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	349	349	301	301	349	349
π	4.89%	4.31%	3.00%	3.33%	1.72%	1.72%
LR_{pof}	0.0000	0.0000	0.0050	0.0014	0.2183	0.2183
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	343	343	295	295	343	343
π	11.40%	5.85%	2.72%	3.06%	1.75%	1.75%
LR_{pof}	0.0000	0.0000	0.0144	0.0043	0.2052	0.2052
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 23: Backtesting LMCADS03 for 5% VaR

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	354	354	306	306	354	354
π	4.25%	5.67%	4.26%	4.26%	3.97%	3.97%
LR_{pof}	0.5070	0.5738	0.5446	0.5446	0.3558	0.3558
LR_{ind}	0.6611	0.4323	0.5262	0.5262	0.5747	0.5747
LR_{cc}	0.7289	0.6272	0.6809	0.6809	0.5578	0.5578
3 Months						
n	352	352	304	304	352	352
π	4.27%	7.69%	6.60%	6.93%	4.56%	4.56%
LR_{pof}	0.5223	0.0314	0.2220	0.1441	0.7002	0.7002
LR_{ind}	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0005	0.0000	0.0000	0.0000	0.0001	0.0001
6 Months						
n	349	349	301	301	349	349
π	10.34%	8.62%	6.67%	7.67%	4.89%	4.89%
LR_{pof}	0.0001	0.0048	0.2065	0.0486	0.9213	0.9213
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	343	343	295	295	343	343
π	18.13%	10.82%	9.25%	10.54%	4.39%	4.39%
LR_{pof}	0.0000	0.0000	0.0015	0.0001	0.5949	0.5949
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 24: Backtesting LMCADS03 for 32% VaR

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	354	354	306	306	354	354
π	29.18%	31.44%	30.16%	29.51%	29.18%	29.18%
LR_{pof}	0.2519	0.8228	0.4896	0.3474	0.2519	0.2519
LR_{ind}	0.8249	0.0866	0.6913	0.3168	0.1350	0.1350
LR_{cc}	0.5061	0.2246	0.7279	0.3896	0.1697	0.1697
3 Months						
n	352	352	304	304	352	352
π	29.63%	32.48%	31.35%	30.69%	30.20%	30.20%
LR_{pof}	0.3378	0.8478	0.8089	0.6245	0.4673	0.4673
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	349	349	301	301	349	349
π	36.21%	33.62%	32.67%	32.00%	31.90%	31.90%
LR_{pof}	0.0960	0.5187	0.8048	1.0000	0.9670	0.9670
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	343	343	295	295	343	343
π	40.64%	34.21%	33.33%	32.33%	32.75%	32.75%
LR_{pof}	0.0008	0.3836	0.6253	0.9085	0.7671	0.7671
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

7.3.2 Student's t-distribution

Table 25: Backtesting EUR/SEK for 1% VaR using the best fitting t-distribution

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	200	200	152	152	200	200	200
π	0.00%	0.50%	1.97%	1.32%	0.50%	1.00%	0.50%
LR_{pof}	1.0000	0.4315	0.2869	0.7090	0.4315	1.0000	0.4315
LR_{ind}	1.0000	0.9656	0.9995	0.9960	0.9956	0.9970	0.9656
LR_{cc}	1.0000	0.7332	0.5672	0.9327	0.7332	1.0000	0.7332
3 Months							
n	198	198	150	150	198	198	198
π	0.00%	1.52%	0.67%	2.00%	0.00%	1.52%	1.52%
LR_{pof}	1.0000	0.4984	0.6623	0.2786	1.0000	0.4984	0.4984
LR_{ind}	1.0000	0.0144	0.9602	0.0356	1.0000	0.0144	0.0144
LR_{cc}	1.0000	0.0398	0.9079	0.0611	1.0000	0.0398	0.0398
6 Months							
n	195	195	147	147	195	195	195
π	1.03%	0.00%	0.00%	0.68%	0.00%	0.00%	0.00%
LR_{pof}	0.9714	1.0000	1.0000	0.6792	1.0000	1.0000	1.0000
LR_{ind}	0.9969	1.0000	1.0000	0.9598	1.0000	1.0000	1.0000
LR_{cc}	0.9994	1.0000	1.0000	0.9169	1.0000	1.0000	1.0000
12 Months							
n	189	189	141	141	189	189	189
π	6.88%	0.00%	0.71%	0.71%	0.00%	0.00%	0.00%
LR_{pof}	0.0000	1.0000	0.7143	0.7143	1.0000	1.0000	1.0000
LR_{ind}	0.0000	1.0000	0.9589	0.9589	1.0000	1.0000	1.0000
LR_{cc}	0.0000	1.0000	0.9339	0.9339	1.0000	1.0000	1.0000

Table 26: Backtesting EUR/SEK for 5% VaR using the best fitting t-distribution

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	200	200	152	152	200	200	200
π	3.50%	5.50%	4.61%	5.26%	1.00%	3.00%	3.00%
LR_{pof}	0.3047	0.7493	0.8211	0.8826	0.0017	0.1622	0.1622
LR_{ind}	1.0000	0.0101	0.3067	0.4177	0.9970	0.1545	0.1545
LR_{cc}	0.5905	0.0349	0.5781	0.7122	0.0071	0.1366	0.1366
3 Months							
n	198	198	150	150	198	198	198
π	4.04%	4.04%	4.67%	6.00%	1.52%	5.56%	5.05%
LR_{pof}	0.5222	0.5222	0.8498	0.5854	0.0087	0.7243	0.9740
LR_{ind}	0.0009	0.0000	0.3113	0.0004	0.0144	0.0000	0.0051
LR_{cc}	0.0031	0.0001	0.5883	0.0017	0.0016	0.0001	0.0199
6 Months							
n	195	195	147	147	195	195	195
π	6.67%	1.54%	2.72%	8.16%	0.00%	3.08%	3.08%
LR_{pof}	0.3085	0.0098	0.1667	0.1054	1.0000	0.1860	0.1860
LR_{ind}	0.0000	0.0263	0.9999	0.0000	1.0000	0.1591	0.0074
LR_{cc}	0.0000	0.0030	0.3843	0.0000	1.0000	0.1547	0.0116
12 Months							
n	189	189	141	141	189	189	189
π	13.76%	1.06%	2.84%	6.38%	0.53%	2.65%	2.12%
LR_{pof}	0.0000	0.0027	0.2010	0.4691	0.0003	0.1040	0.0408
LR_{ind}	0.0000	0.0085	0.0016	0.0000	0.9646	0.0000	0.0009
LR_{cc}	0.0000	0.0003	0.0034	0.0000	0.0017	0.0000	0.0005

Table 27: Backtesting EUR/SEK for 32% VaR using the best fitting t-distribution

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	200	200	152	152	200	200	200
π	30.50%	33.50%	31.58%	31.58%	28.00%	32.50%	33.00%
LR_{pof}	0.6479	0.6506	0.9113	0.9113	0.2197	0.8797	0.7624
LR_{ind}	0.6395	0.1745	0.1040	0.6353	0.8661	0.3438	0.4088
LR_{cc}	0.8074	0.3590	0.2650	0.8881	0.4643	0.6316	0.6792
3 Months							
n	198	198	150	150	198	198	198
π	27.27%	29.29%	28.67%	30.67%	24.75%	28.28%	29.29%
LR_{pof}	0.1480	0.4105	0.3768	0.7253	0.0252	0.2569	0.4105
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	195	195	147	147	195	195	195
π	34.87%	32.31%	36.05%	36.73%	24.10%	29.74%	29.74%
LR_{pof}	0.3935	0.9267	0.2970	0.2240	0.0153	0.4966	0.4966
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	31.22%	26.98%	28.37%	29.08%	20.63%	25.40%	26.46%
LR_{pof}	0.8171	0.1334	0.3501	0.4531	0.0005	0.0469	0.0965
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 28: Backtesting GBP/USD for 1% VaR using the best fitting t-distribution

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	232	232	184	184	232	232	232
π	1.29%	2.16%	1.63%	1.63%	0.86%	1.29%	1.29%
LR_{pof}	0.6677	0.1253	0.4309	0.4309	0.8288	0.6677	0.6677
LR_{ind}	0.9997	1.0000	0.9996	0.9996	0.9974	0.9997	0.9997
LR_{cc}	0.9120	0.3088	0.7333	0.7333	0.9769	0.9120	0.9120
3 Months							
n	230	230	182	182	230	230	230
π	1.74%	2.61%	3.30%	3.30%	1.30%	1.74%	1.74%
LR_{pof}	0.3079	0.0412	0.0139	0.0139	0.6577	0.3079	0.3079
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	227	227	179	179	227	227	227
π	4.85%	3.08%	3.35%	3.35%	2.64%	2.64%	2.64%
LR_{pof}	0.0000	0.0114	0.0128	0.0128	0.0389	0.0389	0.0389
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	13.57%	4.98%	4.05%	4.05%	3.17%	3.17%	3.17%
LR_{pof}	0.0000	0.0000	0.0024	0.0024	0.0098	0.0098	0.0098
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 29: Backtesting GBP/USD for 5% VaR using the best fitting t-distribution

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	232	232	184	184	232	232	232
π	6.47%	6.47%	5.43%	7.61%	3.88%	4.74%	5.17%
LR_{pof}	0.3262	0.3262	0.7895	0.1303	0.4157	0.8554	0.9046
LR_{ind}	1.0000	1.0000	1.0000	0.9402	1.0000	1.0000	1.0000
LR_{cc}	0.6175	0.6175	0.9650	0.3175	0.7180	0.9835	0.9928
3 Months							
n	230	230	182	182	230	230	230
π	6.52%	9.57%	6.04%	7.14%	6.09%	5.65%	5.65%
LR_{pof}	0.3108	0.0045	0.5309	0.2113	0.4638	0.6564	0.6564
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	227	227	179	179	227	227	227
π	14.98%	8.81%	7.26%	7.26%	5.73%	6.17%	6.61%
LR_{pof}	0.0000	0.0168	0.1919	0.1919	0.6230	0.4355	0.2884
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	22.17%	10.41%	8.67%	8.67%	7.24%	6.79%	7.24%
LR_{pof}	0.0000	0.0012	0.0438	0.0438	0.1509	0.2464	0.1509
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 30: Backtesting GBP/USD for 32% VaR using the best fitting t-distribution

	Out-of-Sample					In-Sample	
	EWMA	SMA	GARCH	GJR	IV	GARCH	GJR
1 Month							
n	232	232	184	184	232	232	232
π	35.34%	34.05%	32.07%	32.61%	33.62%	34.05%	33.62%
LR_{pof}	0.2789	0.5052	0.9849	0.8597	0.5983	0.5052	0.5983
LR_{ind}	0.7198	0.6982	0.2652	0.0591	0.7682	0.6982	0.3779
LR_{cc}	0.5218	0.7429	0.5375	0.1659	0.8334	0.7429	0.5900
3 Months							
n	230	230	182	182	230	230	230
π	33.91%	34.78%	28.57%	28.57%	30.87%	33.48%	33.48%
LR_{pof}	0.5360	0.3691	0.3165	0.3165	0.7124	0.6322	0.6322
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months							
n	227	227	179	179	227	227	227
π	37.44%	33.92%	31.84%	32.96%	32.16%	33.04%	33.04%
LR_{pof}	0.0827	0.5371	0.9642	0.7834	0.9592	0.7377	0.7377
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months							
n	189	189	141	141	189	189	189
π	44.34%	31.22%	28.90%	29.48%	29.86%	30.32%	30.77%
LR_{pof}	0.0001	0.8037	0.3780	0.4741	0.4935	0.5899	0.6939
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 31: Backtesting CO1 for 1% VaR using the best fitting t-distribution

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	328	328	280	280	328	328
π	1.83%	3.06%	1.79%	2.15%	0.92%	1.22%
LR_{pof}	0.1742	0.0026	0.2315	0.0939	0.8790	0.6950
LR_{ind}	0.0878	0.0016	1.0000	0.1054	0.9998	1.0000
LR_{cc}	0.0925	0.0001	0.4888	0.0663	0.9885	0.9260
3 Months						
n	326	326	278	278	326	326
π	2.46%	3.38%	3.25%	3.61%	2.46%	2.77%
LR_{pof}	0.0256	0.0007	0.0029	0.0007	0.0256	0.0084
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	323	323	275	275	323	323
π	5.28%	5.28%	3.65%	4.74%	3.42%	3.42%
LR_{pof}	0.0000	0.0000	0.0007	0.0000	0.0006	0.0006
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	317	317	269	269	317	317
π	11.39%	7.28%	5.97%	6.72%	5.06%	5.38%
LR_{pof}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 32: Backtesting CO1 for 5% VaR using the best fitting t-distribution

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	328	328	280	280	328	328
π	4.89%	7.03%	5.73%	7.17%	5.81%	5.81%
LR_{pof}	0.9290	0.1109	0.5818	0.1176	0.5116	0.5116
LR_{ind}	0.0379	0.0028	0.2927	0.2107	0.0189	0.0189
LR_{cc}	0.1156	0.0032	0.4940	0.1343	0.0513	0.0513
3 Months						
n	326	326	278	278	326	326
π	8.00%	8.62%	7.94%	7.94%	7.38%	7.38%
LR_{pof}	0.0219	0.0064	0.0377	0.0377	0.0646	0.0646
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	323	323	275	275	323	323
π	12.73%	7.76%	8.03%	8.39%	6.83%	6.83%
LR_{pof}	0.0000	0.0346	0.0337	0.0183	0.1520	0.1520
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	317	317	269	269	317	317
π	18.99%	9.81%	10.45%	9.70%	7.59%	7.59%
LR_{pof}	0.0000	0.0005	0.0000	0.0000	0.0485	0.0485
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 33: Backtesting CO1 for 32% VaR using the best fitting t-distribution

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	328	328	280	280	328	328
π	29.97%	32.72%	30.47%	30.82%	31.50%	31.50%
LR_{pof}	0.4285	0.7801	0.5812	0.6728	0.8456	0.8456
LR_{ind}	0.0104	0.0403	0.1353	0.1114	0.0062	0.0062
LR_{cc}	0.0273	0.1176	0.2816	0.2576	0.0233	0.0233
3 Months						
n	326	326	278	278	326	326
π	27.38%	29.54%	26.71%	26.71%	28.00%	28.00%
LR_{pof}	0.0705	0.3380	0.0553	0.0553	0.1177	0.1177
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 Months						
n	323	323	275	275	323	323
π	32.61%	30.43%	24.82%	25.55%	29.19%	29.50%
LR_{pof}	0.8152	0.5453	0.0091	0.0195	0.2762	0.3333
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	317	317	269	269	317	317
π	31.96%	26.58%	23.88%	22.39%	25.63%	25.95%
LR_{pof}	0.9885	0.0358	0.0034	0.0005	0.0133	0.0188
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 34: Backtesting LMCADS03 for 1% VaR using the best fitting t-distribution

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	354	354	306	306	354	354
π	1.13%	1.98%	1.64%	1.64%	0.85%	0.85%
LR_{pof}	0.8055	0.1017	0.3042	0.3042	0.7710	0.7710
LR_{ind}	1.0000	1.0000	1.0000	1.0000	0.9999	0.9999
LR_{cc}	0.9702	0.2620	0.5899	0.5899	0.9585	0.9585
3 Months						
n	352	352	304	304	352	352
π	1.42%	2.56%	1.65%	1.65%	1.42%	1.42%
LR_{pof}	0.4524	0.0139	0.2983	0.2983	0.4524	0.4524
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0008	0.0008
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0027	0.0027
6 Months						
n	349	349	301	301	349	349
π	3.74%	3.16%	1.67%	1.67%	1.72%	1.44%
LR_{pof}	0.0001	0.0012	0.2895	0.2895	0.2183	0.4421
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	343	343	295	295	343	343
π	7.89%	4.97%	1.70%	1.70%	1.75%	1.75%
LR_{pof}	0.0000	0.0000	0.2723	0.2723	0.2052	0.2052
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 35: Backtesting LMCADS03 for 5% VaR using the best fitting t-distribution

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	354	354	306	306	354	354
π	4.25%	6.52%	4.92%	4.92%	4.25%	4.25%
LR_{pof}	0.5070	0.2110	0.9475	0.9475	0.5070	0.5070
LR_{ind}	0.6611	0.6776	0.7126	0.1684	0.1460	0.6611
LR_{cc}	0.7289	0.4195	0.9324	0.3865	0.2789	0.7289
3 Months						
n	352	352	304	304	352	352
π	5.13%	7.98%	6.93%	7.59%	5.41%	4.84%
LR_{pof}	0.9126	0.0180	0.1441	0.0537	0.7259	0.8923
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0008	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0027	0.0002
6 Months						
n	349	349	301	301	349	349
π	11.21%	8.91%	6.00%	7.00%	5.17%	5.17%
LR_{pof}	0.0000	0.0025	0.4404	0.1328	0.8833	0.8833
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	343	343	295	295	343	343
π	19.88%	11.11%	7.82%	7.48%	4.68%	4.68%
LR_{pof}	0.0000	0.0000	0.0395	0.0679	0.7827	0.7827
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 36: Backtesting LMCADS03 for 32% VaR using the best fitting t-distribution

	Out-of-Sample				In-Sample	
	EWMA	SMA	GARCH	GJR	GARCH	GJR
1 Month						
n	354	354	306	306	354	354
π	32.86%	33.71%	32.46%	33.44%	31.73%	31.73%
LR_{pof}	0.7293	0.4927	0.8637	0.5906	0.9127	0.9127
LR_{ind}	0.2518	0.1714	0.5860	0.2911	0.1903	0.1903
LR_{cc}	0.4885	0.3102	0.8496	0.4956	0.4216	0.4216
3 Months						
n	352	352	304	304	352	352
π	33.05%	34.47%	32.67%	33.66%	32.48%	32.48%
LR_{pof}	0.6746	0.3238	0.8020	0.5366	0.8478	0.8478
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0008	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0027	0.0002
6 Months						
n	349	349	301	301	349	349
π	37.93%	36.78%	33.33%	34.00%	32.76%	32.76%
LR_{pof}	0.0194	0.0588	0.6218	0.4601	0.7621	0.7621
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 Months						
n	343	343	295	295	343	343
π	42.40%	35.96%	33.67%	32.99%	33.33%	33.33%
LR_{pof}	0.0001	0.1197	0.5403	0.7158	0.5984	0.5984
LR_{ind}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LR_{cc}	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

8 References

Aggarwal, R., Edward, J., & Mellen, L. E. (1990). The Nature of Currency Black Markets: Empirical Test of Weak and Semistrong Form Efficiency. *The International Trade Journal*, 5(1), 1–24.

Bank for International Settlements (2016). Triennial central bank survey: global foreign exchange market turnover in 2016. `http://www.bis.org/publ/rpfxf16fxt.pdf`. Online; accessed on 2016-12-05.

Becker, R., Clements, A. E., & White, S. I. (2006). On the informational efficiency of S&P500 implied volatility. *The North American Journal of Economics and Finance*, 17(2), 139–153.

Bekaert, G., Erb, C. B., Harvey, C. R., & Viskanta, T. E. (1998). Distributional Characteristics of Emerging Market Returns and Asset Allocation. *The Journal of Portfolio Management*, 24(2), 102–116.

Blattberg, R. C. & Gonedes, N. J. (1974). A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices. *The Journal of Business*, 47(2), 244–280.

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.

Brooks, C. & Persand, G. (2002). Model choice and value-at-risk performance. *Financial Analysts Journal*, 58(5), 87–97.

Canina, L. & Figlewski, S. (1993). The informational content of implied volatility. *Review of Financial studies*, 6(3), 659–681.

Chiras, D. P. & Manaster, S. (1978). The information content of option prices and a test of market efficiency. *Journal of Financial Economics*, 6(2), 213–234.

Christensen, B. J. & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50(2), 125–150.

Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, (pp. 841–862).

Cornew, R. W., Town, D. E., & Crowson, L. D. (1984). Stable Distributions, Futures Prices, and the Measurement of Trading Performance. *Journal of Futures Markets*, 4(4), 531–557.

Danielsson, J. & De Vries, C. G. (2000). Value-at-risk and extreme returns. *Annales d'Economie et de Statistique*, (pp. 239–270).

Day, T. E. & Lewis, C. M. (1992). Stock market volatility and the information content of stock index options. *Journal of Econometrics*, 52(1-2), 267–287.

Dowd, K. (2007). *Measuring market risk*. John Wiley & Sons.

Enders, W. (2010). *Applied Econometric Time Series*, chapter 3. John Wiley & Sons, Inc., 3 edition.

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica: Journal of the Econometric Society*, 50(4), 987–1007.

Fama, E. F. (1965). The Behavior of Stock-Market Prices. *The Journal of Business*, 38(1), 34–105.

Fleming, J. (1998). The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of empirical finance*, 5(4), 317–345.

Giot, P. & Laurent, S. (2003). Market risk in commodity markets: a var approach. *Energy Economics*, 25(5), 435–457.

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779–1801.

Gray, J. B. & French, D. W. (1990). Empirical Comparisons of Distributional Models for Stock Index Returns. *Journal of Business Finance & Accounting*, 17(3), 451–459.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies*, 6(2), 327–343.

Hoppe, R. (1998). VaR and the Unreal World. *Risk*, 11, 45–50.

Huang, Y. C. & Lin, B.-J. (2004). Value-at-Risk Analysis for Taiwan Stock Index Futures: Fat Tails and Conditional Asymmetries in Return Innovations. *Review of Quantitative Finance and Accounting*, 22(2), 79–95.

Hull, J. & White, A. (1998). Incorporating volatility updating into the historical simulation method for value-at-risk. *Journal of Risk*, 1(1), 5–19.

Jackwerth, J. C. & Rubinstein, M. (1996). Recovering probability distributions from option prices. *The Journal of Finance*, 51(5), 1611–1631.

Jarque, C. M. & Bera, A. K. (1987). A test for normality of observations and regression residuals. *International Statistical Review/Revue Internationale de Statistique*, (pp. 163–172).

Jorion, P. (1997). *Value at risk*. McGraw-Hill, New York.

Kim, D. & Kon, S. J. (1994). Alternative Models for the Conditional Heteroscedasticity of Stock Returns. *Journal of Business*, 67(4), 563–598.

Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2).

Lin, C.-H. & Shen, S.-S. (2006). Can the student-t distribution provide accurate value at risk? *The Journal of Risk Finance*, 7(3), 292–300.

Ljung, G. M. & Box, G. E. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2), 297–303.

Lopez, J. A. (1998). Methods for evaluating value-at-risk estimates. *Economic Policy Review*, 4(3).

Mandelbrot, B. B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, 36(4), 394–419.

Manganelli, S. & Engle, R. F. (2001). Value at Risk Models in Finance. ECB Working Paper No. 75.

McNeil, A. J. & Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of empirical finance*, 7(3), 271–300.

Pafka, S. & Kondor, I. (2001). Evaluating the riskmetrics methodology in measuring volatility and value-at-risk in financial markets. *Physica A: Statistical Mechanics and its Applications*, 299(1), 305–310.

Papaioannou, M. & Gatzonas, E. (2002). Assessing market and credit risk of country funds: A value-at-risk analysis. *Global Risk Management: Financial, Operational and Insurance Strategies, International Finance Review*, 3, 61–79.

Penza, P. & Bansal, V. K. (2001). *Measuring market risk with value at risk*, volume 17. John Wiley & Sons.

Praetz, P. D. (1972). The Distribution of Share Price Changes. *The Journal of Business*, 45(1), 49–55.

RiskMetrics Group (1996). *Riskmetrics - Technical Document.* J.P.Morgan and Reuters, New York, 4th edition.

Rockafellar, R. T. & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7), 1443–1471.

Saunders, A. & Allen, L. (1999). Credit risk measurement.

Szakmary, A., Ors, E., Kim, J. K., & Davidson, W. N. (2003). The predictive power of implied volatility: Evidence from 35 futures markets. *Journal of Banking & Finance*, 27(11), 2151–2175.

Taleb, N. (1997). Against value-at-risk: Nassim taleb replies to philippe jorion. *Derivatives Strategy*, 2(4).

Tucker, A. L. & Pond, L. (1988). The Probability Distribution of Foreign Exchange Price Changes: Tests of Candidate Processes. *The Review of Economics and Statistics*, 70(4), 638–647.

Xu, X. & Taylor, S. J. (1995). Conditional volatility and the informational efficiency of the PHLX currency options market. *Journal of Banking & Finance*, 19(5), 803–821.