Stockholm School of Economics
Master's Thesis in Finance

Portfolio optimization

A trading strategy based on time varying volatility

Henrik Gripenvik
20382@student.hhs.se

Abstract

This thesis evaluates the possibility of making relatively reliable forecasts of future variances, covariances and returns of a portfolio of risky assets. The forecasts are used as input-variables in a trading strategy based on the Markowitz portfolio optimization algorithm forming an ex ante-optimal portfolio of stocks. The results show that the trading strategy, although on average achieving a 14 percent higher monthly Sharpe ratio than the benchmark index, cannot with a satisfactory statistical significance outperform the FTSE100. However, the results should be considered to be of economic significance to an investor since, on average, the strategy renders a three percent higher annual return while maintaining the same level of risk. Furthermore, even greater differences in risk-adjusted returns should be possible to achieve by relaxing the restrictions associated with regular mutual funds and thereby restructure the portfolio mandate like a hedge fund.

Tutor: Stefan Engström

Presentation date: 2006-06-12

Venue: Room 343

Discussants: Daniel Johansson and Andreas Ohlsson

Acknowledgements: I would like to thank my tutor Stefan Engström for his valuable comments and input during the course of this thesis. In addition, I am very grateful for the proofreading by Birgit and Christian Gripenvik.

Contents

List of figures

List of tables

1. Introduction

1.1. Background

Modern portfolio theory has its roots in Markowitz' creative ideas regarding an investor's choice of investment portfolio. He argued that rational investors should hold a portfolio that is efficient in the sense that it offers the highest expected return for a given level of risk. Historically however, mean-variance portfolio optimization was not used to the extent that could have been expected. Rather, professional investors have traditionally tried to identify the securities with the highest expected returns. However, a vast amount of research has painted a rather gloomy picture of the possibilities of consistently achieving a higher risk-adjusted return than that of a passive index portfolio strategy. This in turn has resulted in rather strong opinions regarding the true abilities of fund managers as well as an increasing interest for index funds.

In recent years there has been an upsurge in the interest of portfolio optimization, which can to a large extent be explained by the more or less explosive development in computer power, improved statistical models as well as the accessibility of reliable data. As a result, there are several highly successful hedge funds, which to a large extent rely on quantitative analysis and portfolio optimization in their efforts to generate substantial economic value for their investors. Furthermore, empirical evidence has also shown that there are great opportunities of risk reduction in portfolio optimization. The extent, to which these opportunities can be exploited, depends on the quality of the forecasts of the first and second moments of a return series. Traditionally, because more reliable models were not available at the time, these variables were forecasted by using rather crude methods such as historical averages (Bansal et al., 2004).

Another type of strategy, which focuses on the dynamics of volatility by using advanced statistical models, is now possible due to the increasing computer power. A relatively reliable forecast of the volatility and the covariance between stocks could be exploited for the purpose of forming an ex ante-optimal portfolio of risky assets. Considering the enormous upswing in the public and professional interest of investing in mutual funds this opportunity should be of great interest for professional investors in constant search of that elusive gold mine at the end of the rainbow.

1.2. Purpose

The purpose of the master thesis is to evaluate the skills of a trading strategy, based on forecasted volatilities, covariances and returns, in relation to a relevant benchmark index such as the FTSE100. The evaluation will be based on both the statistical and the economic significance of the results where economic significance is defined as a situation with statistically insignificant results where the results are still of interest for an investor. Furthermore, the thesis will also investigate how the expected minimum variance portfolio, which according to economic theory should perform worse than other portfolios on the efficient frontier, actually performs in relation to the portfolio with the highest expected Sharpe ratio.

1.3. Delimitation

The study is based entirely on U.K. data and forms portfolios of the assets, which make up the index FTSE100[1]. A direct consequence is that no conclusions regarding the ability of the trading strategy beyond the U.K. market can be made. The portfolios will constitute of the 30 largest stocks of the index, which should still make the comparison relevant since these stocks make up approximately 70 percent of the index. The fact that 30 percent are omitted should not invalidate the interpretations of the results since even a very large change in any of these stocks only have a small impact on the FTSE100. The reason for this is that the weights of these stocks, on average, are less than half a percent of the market value.

The constituents of the index changes rather frequently. However, due to an ambition of keeping the workload at a reasonable level, the stocks will be considered to be fixed for up to one year.

The study is based on daily data between 1995-01-01 and 2005-11-25. Furthermore, no consideration will be taken to transaction costs or bid-ask spreads since these can often be regarded as negligible for larger financial institutions.

1.4. The strategy

The trading strategy is essentially a buy and hold strategy[2] where two different holding periods will be considered, one month and a quarter of a year. The reason for using different holding periods is to evaluate which holding period offers the best risk-return relationship measured as the highest Sharpe ratio. The weights of each asset is decided through the use of

[1] FTSE100 is a broad free float stock market index acting as a benchmark index for several equity funds. The constituencies are chosen by market capitalisation and are reviewed quarterly. www.ftse.com.

[2] Buy and hold means that the investor buys the stocks and holds them for a longer time period disregarding short-term fluctuations in the market. (Bansan et al, 2004).

Markowitz' portfolio optimization algorithm by using out-of-sample forecasts[3] of the volatility, covariance and the expected return for the given time period. The procedure is then repeated according to the rebalancing pattern and thereafter evaluated on a monthly basis in relation to FTSE100.

The strategy evaluates two different kinds of optimization procedures, maximizing the expected Sharpe ratio and minimizing the expected volatilitity. The reason for this is the fact that the expected Sharpe ratio, to a larger extent, is affected by forecasts of the expected return, which are notoriously hard to do with any significant results. If the slope of the true efficient frontier is relatively flat, then the loss in return of minimizing the expected volatility could be outweighed by the gain in reduced volatility – thereby yielding a higher risk-adjusted return.

The different portfolios are also affected by the restrictions, which normally apply for fund management. For instance, this means that an individual asset is not allowed to constitute more than 10 percent of the total market value of the portfolio[4].

The data used in this thesis has been collected from the databases SIX Trust[5] and EcoWin and by contact with the FTSE organisation.

1.5. Previous research

Bansal et al (2004) use a dynamic trading strategy in which they develop a method for constructing optimally managed portfolios, which exploits the possibility that asset returns are predictable. The authors characterize the degree of predictability by comparing the performance of portfolios that include conditioning information to those who do not. They find that investing in actively managed funds produces considerable economic gains in relation to fixed weight strategies and thereby conclude that active funds should be of great benefit to investors. Furthermore, the authors find evidence of buy and hold portfolios performing even worse than fixed weight strategies. The intuition is that the buy and hold strategy leads to a less than optimally diversified portfolio in the long term.

Conrad et al (2003) assess the profitability of momentum strategies using a stochastic discount factor approach. They find that a stochastic discount factor can be constructed from a set of industry-sorted portfolios, which explains about half of the profitability of the trading strategy when agents cannot use conditioning information. When allowing for conditioning

[3] Out of sample forecasts are based on historical data until the day/time to be forecasted.

[4] The UCITS-directives regulate that a mutual equity fund is allowed to invest up to 5 percent of its total wealth in any given security with the exception that up to 10 percent is allowed in a security if these securities together do not make up more than 40% of the total portfolio wealth.

[5] SIX Trust is foremost an information- and analysis tool for professional investors working on the Nordic capital markets. (www.six.se).

information they find that the abnormal profits decline and they interpret these results as a sign that at least a part of the returns of momentum portfolios stem from the risk of the strategy rather than just irrational pricing behavior. However, since a significant portion of the momentum profits cannot be explained, they are not able to rule out the existence of residual mispricing.

Guo (1999) compares the performance of a GARCH-based dynamic volatility trading strategy for currency options with an implied stochastic volatility regression (ISVR) based strategy. He uses data from 1983 to 1993 obtained from the Philadelphia stock exchange and finds evidence favouring the ISVR-model. However, both models are able to provide an investor with abnormal returns as well as lower correlation with several asset classes, indicating possible improvements through diversification.

There are a wide variety of volatility models and the research on the performance of different forecasting techniques is rather extensive. Lopez and Walter (2001) evaluate the performance of several covariance matrix forecasts using standard statistical loss functions and a value at risk framework (VaR). Their findings show that, within a VaR framework, simple specifications such as the Exponentially Weighted Moving Average (EWMA) perform better than forecasts from more sophisticated models such as implied volatility models.

Day and Lewis (1992) combines the current market expectations, as reflected in option prices, with past returns information captured by a GARCH model. They do so by embedding an implied volatility variable as an exogenous variable in the GARCH equation. They use data on the S&P100 index together with the corresponding index options and find results indicating that their model is able to generate more reliable volatility forecasts than by just using an implied volatility model.

Lamoureux and Lastrapes (1993) examine the behavior of measured variances from the options market and the underlying stock market. They found evidence that implied volatility performed very well but were unable to reject the hypothesis that predictions from GARCH models contain additional information regarding future volatility.

Ghysels et al (2004) develop a mixed data-sampling model (MIDAS) which is a regression model based on time series data sampled at different time frequencies. They use high frequency U.S. market data of the Dow Jones Composite Portfolio as well as individual stocks during the time period of 1993 to 2003. The authors find that their MIDAS estimator is a better forecaster of the stock market variance than rolling window or GARCH estimators.

Kim et al. (1998) compare several different GARCH models with a stochastic volatility model (SV), where the variance is specified to follow some latent stochastic process. Their findings show that the SV model offers better forecasts of the variance than most GARCH models.

Author	**Method**	**Results**
Bansal et al.	Dynamic trading strategy	Higher risk-adjusted return than a buy and hold strategy.
Conrad et al.	Stochastic discount factor	Momentum effects may still be present.
Guo	Dynamic volatility based strategies for currency options.	Higher risk-adjusted return in relation to a more naive model.
Lopez and Walter	Forecasting covariance-matrices by using EWMA	Better forecasts than implicit volatility.
Day and Lewis	Volatility forecasts by using a GARCH model	Better forecasts than implicit volatility.
Lamoureux and Lastrapes	Forecasting volatility using implied volatility.	Implied volatility gives relatively reliable forecasts.
Ghysels et al.	Volatility forecasts by using MIDAS	Higher predictive power than GARCH.
Kim et al.	Volatility forecasts by using Stochastic volatility	Higher predictive power than GARCH.

Table 1: Summary of previous research

1.6. Outline and reader's guide

The remainder of this thesis is structured as follows: Chapter 2 reviews the concept of portfolio optimization as well as an introduction to common properties of financial time series. Chapter 3 introduces the different models used in the trading strategy. Chapter 4 presents the empirical results together with an analysis of the data. Finally, chapter 5 presents the conclusions of the thesis.

2. Portfolio optimization

The introduction of modern portfolio theory by Harry Markowitz has lead to a mathematical explanation of the expression “don’t put all your eggs in one basket”. One of the most fundamental conclusions in Markowitz’ portfolio choice theory is that rational investors should not choose assets only because of their unique properties such as the expected return and variance, but should also consider the covariation between the different assets. As the number of assets in a portfolio increases, the covariance makes up an increasingly greater part of an individual asset’s contribution to the total risk of a portfolio. This can be seen in the figure below where the variance terms make up the diagonal elements of the variance-covariance matrix. (Markowitz, 1952).

$$\begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{11} & \cdots & a_{nn} \end{pmatrix}$$

Figure1: Variance-covariance matrix

For each incremental asset, *1* variance term and *n-1* covariance terms are added to the matrix. As long as an asset does not correlate perfectly with the other assets in the portfolio, the total variance will be reduced. In an investment perspective, this can be seen in terms of diversification. The idea is that a portfolio should consist of a large amount of assets, which belong to different lines of business with the purpose of spreading the risk exposure and achieving lower correlation. The effect of diversification is common knowledge within the field of financial theory and a great number of researchers have found supporting evidence. For instance, Solnik (1974) shows that the risk of a well-diversified portfolio initially decreases dramatically and then converges to an undiversifiable level of risk.

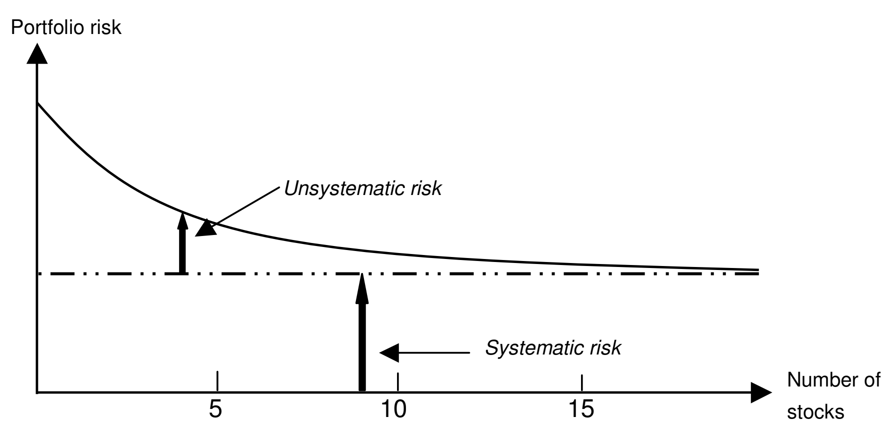

Figure 2: The effect of diversification

By using the optimization procedure for a given universe of securities, an efficient frontier of risky assets may be formed where the portfolios on the frontier are efficient in the sense that they offer the highest return for any given level of risk.

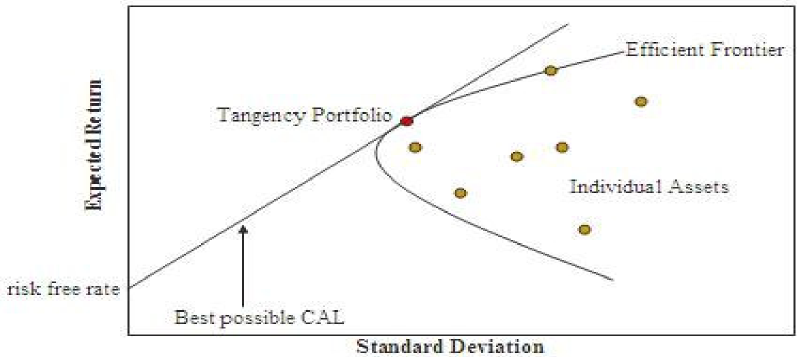

Figure 3: The efficient frontier

This is the foundation of the strategy, which relies on predicted changes of the efficient frontier and forms ex ante-optimal portfolios based on historical time series data of the stocks in the FTSE100 index.

2.1. Financial time series

Time series of financial data such as stock returns are characterised by the fact that the probability distribution have fatter tails in comparison to the normal distribution.

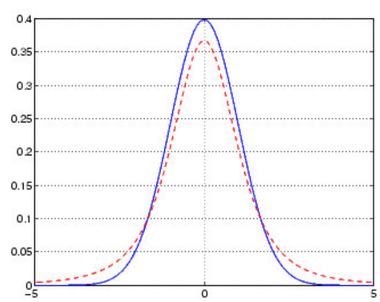

Figure 4: Normal distribution in comparison to a heavier-tailed distribution.

Furthermore, high frequency return series, such as daily or weekly data, are often uncorrelated but not independent whereas the volatility[6] of stock returns, however, are correlated over time.

[6] Measured as the squared daily return.

This phenomenon is called volatility clustering, which means that small changes in the stock price have a tendency to be followed by another small change and vice versa. An example of this can be seen in the figure below. (Aas et al., 2004).

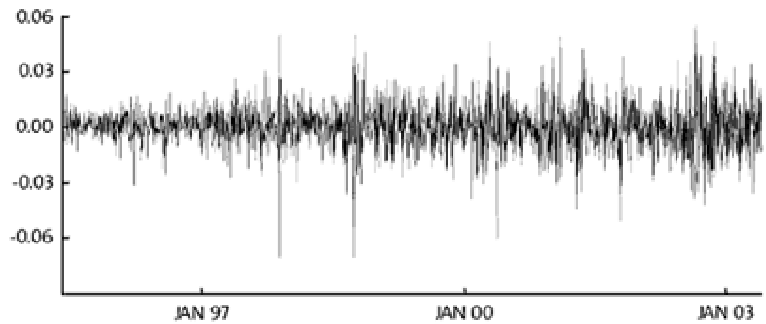

Figure 5: Daily returns of the Standard & Poor's 500 stock index

Volatility clustering explains a great part, but not the entire effect, of heavier tails in the distribution (Knight et al., 2001). This in turn, can be a signal that stock return series can be better explained with a heavier-tailed distribution.

Beside these visual properties of stock return distributions it is common that the volatility of stock returns is asymmetric, which means that negative shocks have a tendency to result in higher increases in volatility than positive ones. This is often referred to as a leverage effect (Ibid.).

2.2. Forecasting

Given that a time series can be seen as a stochastic process, i.e. as a sequence of random observations, it may contain a certain level of correlation from one observation to another. The correlation structure can be used for the purpose of making forecasts of future values based on historical observations. By exploiting the structure a time series can be broken down to a deterministic part and a pure random part. Mathematically this can be expressed as:

$$r_t = f(t-1, X) + \varepsilon_t \tag{1}$$

where $f(t-1, X)$ represents the deterministic part of the present return as a function of all the information available at time *t-1*. This includes the historical random error terms, $(\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots)$, the historical observations, $(r_{t-1}, r_{t-2}, \ldots)$, and all other relevant explanatory data, *X*. The random part, ε_t, represents the error in r_t, but can also be seen as the forecasting error.

A common assumption regarding ε_t when modelling financial time series is that the error term follows an uncorrelated normal distribution with a mean of zero. Despite the fact that the error terms are uncorrelated they are dependent[7], which means that there are possibilities for statistical methods to capture this behaviour. (Aas et al., 2004). One of the more suitable models for this purpose is the Autoregressive Conditional Heteroscedasticity model (ARCH).

[7] Dependence means that the estimate depends on previous values. (Brooks, 2002).

3. Model specification

3.1. Forecasting volatility

During the last two decades a large number of ARCH models have evolved with the purpose of modelling time varying variance. The ARCH model was developed by Engle (1982) and then generalised by two independent researchers, Bollerslev (1986) and Taylor (1986) (Brooks, 2002). The model captures the clustering effect in time series by assigning greater weights to more recently observed data. Furthermore, the model expresses the conditional variance as a linear function of historically squared error terms, which mathematically can be expressed as:

$$\sigma_t = A_0 + A_1\varepsilon_{t-1}^2 + ... + A_q\varepsilon_{t-q}^2 \qquad where\ A_i \geq 0 \tag{2}$$

As seen in the formula, the conditional variance increases with the magnitude of the squared error terms regardless of their signs. The lag length, *q*, decides the length of time that shocks persist in the conditional variance[8].

Bollerslev and Taylor later generalised the ARCH model (GARCH) and included past variance terms in the conditional variance, which equipped the model with a longer memory. The general GARCH(*p,q*) model can mathematically be expressed as:

$$\sigma_t^2 = K_0 + \sum_{i=1}^{q} A_1\varepsilon_{t-i}^2 + \sum_{j=1}^{p} G_j\sigma_{t-j}^2 \tag{3}$$

where the present and conditional variance depend on *q* lags of the squared residuals as well as on *p* lags of the conditional variance. Empirical evidence suggests that a GARCH(1,1) process is often sufficient to capture the volatility process in most financial time series. (Bollerslev et al., 1992).

Since its discovery, the model has later been revised into many different versions with the purpose of providing efficient solutions to problems, which the original model could not handle in a satisfactory way.

[8] Lags are past observed values of a variable, e.g. two lags indicate the values of the variable of the two previous time periods.

The forecasts of the variance of the constituting stocks are made with an EGARCH model, which was in its original form developed by Nelson (1991). The model can be expressed as:

$$\ln\sigma_t^2 = K + \sum_{i=1}^{P} G_i \ln\sigma_{t-i}^2 + \sum_{j=1}^{Q} A_j \left[\frac{|\varepsilon_{t-j}|}{\sigma_{t-j}} - E\left(\frac{|\varepsilon_{t-j}|}{\sigma_{t-j}} \right) \right] + \sum_{j=1}^{Q} L_j \left(\frac{\varepsilon_{t-j}}{\sigma_{t-j}} \right) \tag{4}$$

$$\textit{where} \quad E\left(\frac{|\varepsilon_{t-j}|}{\sigma_{t-j}} \right) = E(|z_{t-j}|) = \sqrt{\frac{v-2}{\pi}} \cdot \frac{\Gamma\left(\frac{v-1}{2}\right)}{\Gamma\frac{v}{2}} \qquad \textit{with } v > 2 \textit{ degrees of freedom}$$

The EGARCH parameters *K, G, A* and *L* represent a constant, previous variance terms (GARCH effect), previous residual terms (ARCH effect) and finally the leverage effect.

The reason for using this model in favour of other similar ones is that it considers the leverage effect, mentioned above, in a satisfactory way. Whenever there is a negative relationship between the volatility and the return the coefficient, L_j, will be negative. Furthermore, because the model uses the logarithm of the conditional variance it is not necessary to inflict any restrictions on the parameters in order to make sure that non-negative results are achieved.

3.2. Modelling covariances

Traditionally, the covariance between two risky assets, *x* and *y,* has been modelled by using a relatively naive model, which generates forecasts by using equally weighted historical data. The model estimates the covariance as:

$$COV_{xy,t} = \frac{1}{n}\sum_{i=1}^{n} r_{x,i} r_{y,i} \tag{5}$$

However, this model is relatively blunt since past observations are not necessarily representative for the given time period. Furthermore, all observations are assigned equal weights regardless of where in time they are situated. Because of these inefficiencies the covariance is forecasted using an EWMA model which expresses the covariance as:

$$\sigma(x,y)_t = (1-\lambda)\sum_{i=0}^{\infty} \lambda^i rx_{t-1} ry_{t-1} \tag{6}$$

In contrast to the naive model, the EWMA model assigns different weights to each observation by introducing an exponential smoother, λ, which dictates the speed with which the relative weight decreases with time. The exponential smoother takes a value between zero and one and is usually set to 0,94, which in studies performed by RiskMetrics[9] has shown to give the best forecast of future realised daily volatilities. (Mina et al., 2001).

3.3. Forecasting expected return

In order to make forecasts of the expected return, the last input variable in the Markowitz algorithm, an ARMA(*p,q*) is used. The model consists of an autoregressive part, ϕ, where the present value of the return depends on previous values as well as on an error term, u_t. Mathematically it can be expressed as:

$$r_t = \mu + \sum_{i=1}^{p} \phi_i r_{t-i} + u_t \qquad (7)$$

Furthermore, the model consists of a MA-component, θ, where the return is expressed as:

$$r_t = \mu + \sum_{i=1}^{q} \theta_i u_{t-i} + u_t \qquad (8)$$

This means that the return can be described as a linear combination of white noise terms[10]. The AR, as well as the MA-component requires that the return series is stationary and invertible which, in a simplified sense, means that the probability of *r* taking a value within a given interval should be the same for the whole time period. Finally, the ARMA(*p,q*) model combines these properties and model the return as:

$$r_t = \mu_t + \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{i=1}^{q} \theta_i u_{t-i} \qquad (9)$$

[9] RiskMetrics is a company, previously belonging to J.P. Morgan, with a focus on risk management.

[10] A white noise process has a mean and variance of $(0, \sigma^2)$ and a lack of autocorrelation.

4. Empirical results

The study is based on four different portfolios, which are optimized with respect to expected Sharpe ratio or expected minimum variance with a forecasting period of one month and a quarter of a year. A more detailed disclosure of the different portfolio data can be found in Appendix A.

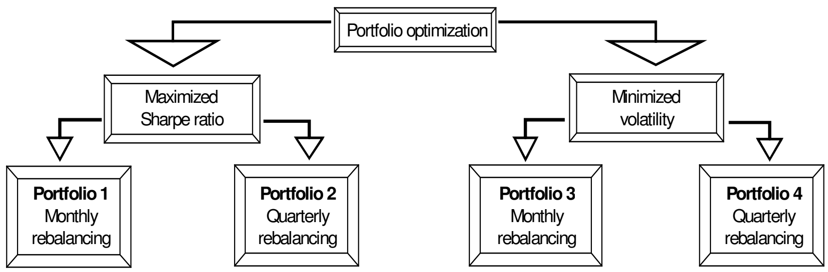

Figure 6: Portfolio optimization strategies

4.1. Strategies with optimized Sharpe ratio

In absolute terms, these strategies have performed very well in relation to FTSE100. Portfolio 1 (monthly rebalancing) grows 187 percent, which can be compared to 103 percent for the index over the total measurement period. On an annual basis portfolio 1 generates an average return of 12 percent while the FTSE100 has offered an average return of 9 percent[11]. At the same time, portfolio 2 (quarterly rebalancing) has grown 162 percent with an average annual return of 11 percent.

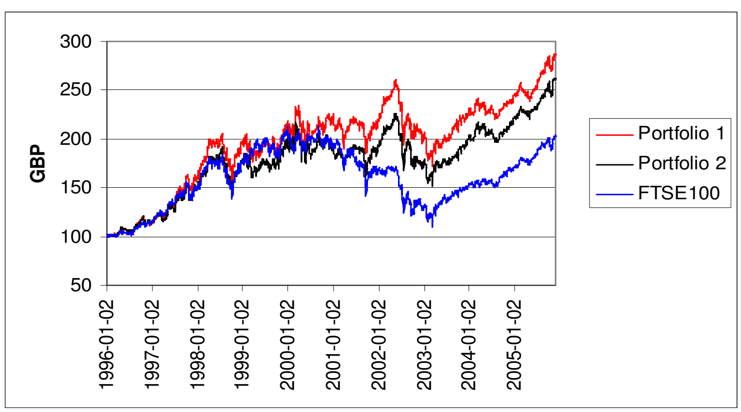

Figure 7: Price development of portfolios 1, 2 and FTSE100

[11] Geometric mean.

Year	Portfolio 1	Portfolio 2	FTSE100	Portfolio 1 - FTSE100	Portfolio 2 - FTSE100
1996	15%	16%	17%	-2%	-1%
1997	41%	30%	29%	12%	2%
1998	21%	17%	18%	4%	0%
1999	12%	14%	21%	-9%	-6%
2000	3%	-4%	-8%	11%	4%
2001	-2%	0%	-14%	12%	14%
2002	-6%	-8%	-21%	15%	13%
2003	12%	14%	17%	-5%	-2%
2004	6%	9%	11%	-5%	-2%
2005	16%	17%	18%	-2%	-1%
Average	12%	11%	9%	3%	2%

Table 2: Yearly returns of portfolios 1, 2 and FTSE100

In an evaluation of the higher returns it is, however, more interesting to examine the quality of the return forecasts. A linear regression has been performed where the expected return acts as the explanatory variable and the realised return is the dependent one. The return forecasts have very low predictive power, only 0,3 percent at the best, and are not statistically significant. In light of this it must be considered remarkable that the portfolios have an almost consistently better price development than the index. One possible explanation to this fact could be that there has systematically existed a "large-cap" effect in the selection of stocks made each year[12]. However, in an examination of the FT30 index, where an average yearly dividend of 3,8 percent has been added back to the index, it is evident that this is not the case[13]. The FT30 performs a lot worse than the FTSE100 with a growth of only 32 percent for the entire time period, which makes it reasonable to conclude that such an effect has not been present during the time period.

A comparison of the absolute return is, however, not very interesting in itself since the return of a portfolio usually is in line with the total risk that an investor is forced to accept. In a comparison of the total standard deviation it turns out that the portfolios on average, have the same risk as the index.

Year	Portfolio 1	Portfolio 2	FTSE100	Portfolio 1 - FTSE100	Portfolio 2 - FTSE100
1996	10%	9%	9%	0%	0%
1997	17%	17%	15%	1%	2%
1998	20%	21%	21%	-1%	0%
1999	19%	19%	18%	1%	1%
2000	23%	24%	19%	4%	4%
2001	17%	17%	22%	-5%	-5%
2002	22%	23%	28%	-5%	-5%
2003	18%	18%	19%	-2%	-1%
2004	10%	11%	10%	0%	0%
2005	11%	10%	9%	2%	1%
Average	17%	17%	17%	0%	0%

Table 3: Yearly standard deviations of portfolios 1, 2 and FTSE100

[12] The fact that only the 30 largest companies are being used could imply that there is a bias if larger companies have systematically performed better during the time period.
[13] No total return-version of the FT30 index exists, which means that no dividends are reinvested. The average cash dividend in relation to the stock price of the FTSE100 index has been 3,8 percent. (www.di.se).

On average, the volatility is 17 percent for both portfolios as well as for the index. The quality of the forecasts of the total standard deviation is relatively good. Portfolios 1 and 2 have R^2-measures of 36,7 and 37,1 percent. These results are highly significant with p-values[14] of $2{,}8 \cdot 10^{-13}$ for portfolio 1 and $3{,}05 \cdot 10^{-5}$ for portfolio 2.

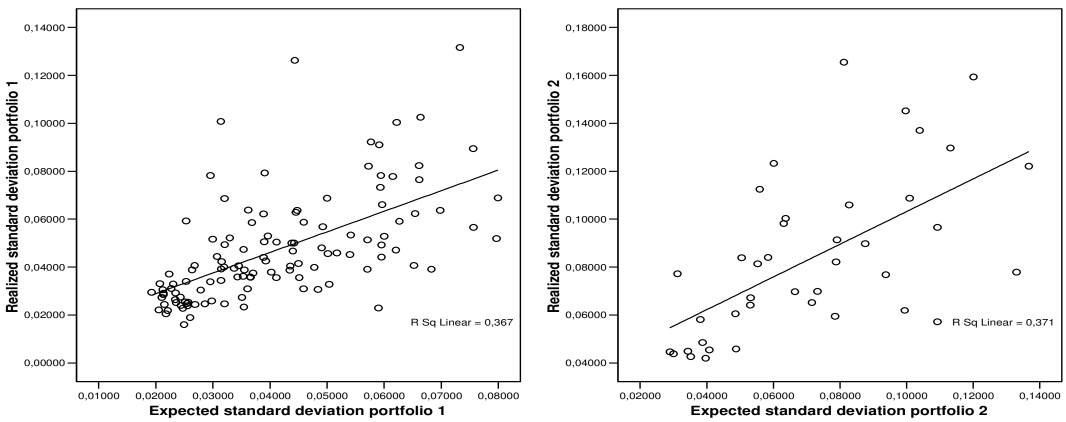

Figure 8: Linear regression of the predictive power of the volatility forecasts of portfolios 1 and 2

The fact that both portfolios have higher average returns and lower standard deviations than the index indicates that the portfolios, on average, offer a higher risk-adjusted return than the FTSE100 during the entire measurement period. In comparison to the FTSE100, this translates into a 14 percent higher monthly Sharpe ratio for portfolio 1 and 4 percent higher for portfolio 2.

Year	Portfolio 1	Portfolio 2	FTSE100	Portfolio 1 - FTSE100	Portfolio 2 - FTSE100
1996	0,20	0,25	0,35	-0,15	-0,10
1997	0,48	0,36	0,44	0,04	-0,09
1998	0,26	0,23	0,22	0,04	0,01
1999	0,09	0,09	0,30	-0,21	-0,21
2000	0,13	-0,03	-0,09	0,21	0,06
2001	-0,04	-0,04	-0,27	0,23	0,23
2002	0,10	0,09	-0,25	0,36	0,34
2003	0,23	0,23	0,31	-0,08	-0,08
2004	0,14	0,21	0,23	-0,08	-0,01
2005	0,42	0,47	0,57	-0,15	-0,10
Average	0,21	0,19	0,18	0,03	0,01

Table 4: Monthly average Sharpe ratios of portfolios 1, 2 and FTSE100

In evaluating the portfolios there are however great difficulties since one can only observe the performance of the portfolios ex post and then hope that random effects are not mistaken for or conceals the strategy's true abilities. A t-test of the differences in risk-adjusted returns is therefore performed, which shows that the results are not statistically significant[15].

[14] P-value is equivalent to the lowest significance level to which a null hypothesis can be rejected. (Gujarati, 2003).
[15] See Appendix A.

These results are also supported by running the regression:

$$R_{it} - R_{FTSE100t} = \alpha_i + \beta_i \left(R_{FTSE100t} - rf_t \right) + \varepsilon_t \qquad (10)$$

following the methodology of Engström (2005) where R_{it} is equal to the return of portfolio *i* and α_i is the Jensen's alpha measure. Portfolios 1 and 2 obtain alpha values of 0,34 and 0,27 percent respectively on a monthly basis with p-values of 0,125 and 0,252, i.e. the results are not statistically significant at the conventional level of 5 percent.

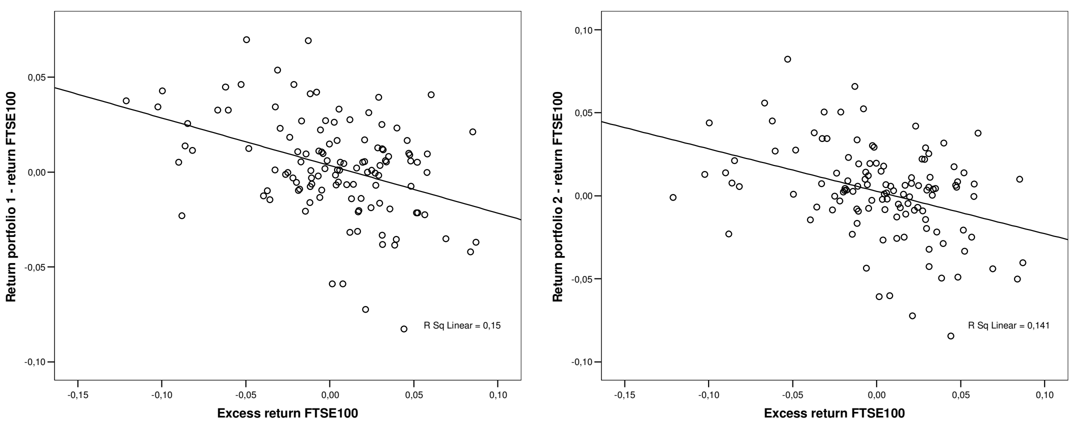

Figure 9: Linear regression of the performance of portfolio 1 and 2 in excess of the FTSE100

The implication of these results is that it is not possible to make any statistical inferences regarding the abilities of the trading strategies. This situation is very common since it is necessary to have very large amounts of data in order to even ensure that large observed differences are significantly different from zero[16]. (Bodie et al., 2005). The statistical insignificance of the results does, however, not necessarily make them economically insignificant. The difference in performance is likely to be considered relevant by an investor who is trying to outperform his/her benchmark index since the monetary value of achieving, on average, 3 percent higher annual returns while maintaining the level of risk should be substantial.

[16] Bodie et al (2005) exemplify with a fund manager who obtains a constant alpha-value of 0,2 percent per month while at the same time the statistical properties are assumed to be constant. In this extremely favourable situation it is still necessary to obtain 32 years of data in order to ensure that there actually exists an alpha different from zero.

4.2. Strategies with minimized volatility

Following the reasoning in section 1.4, The Strategy, two portfolios with minimized expected volatility are formed. Portfolio 3 is rebalanced on a monthly basis and portfolio 4 is rebalanced quarterly.

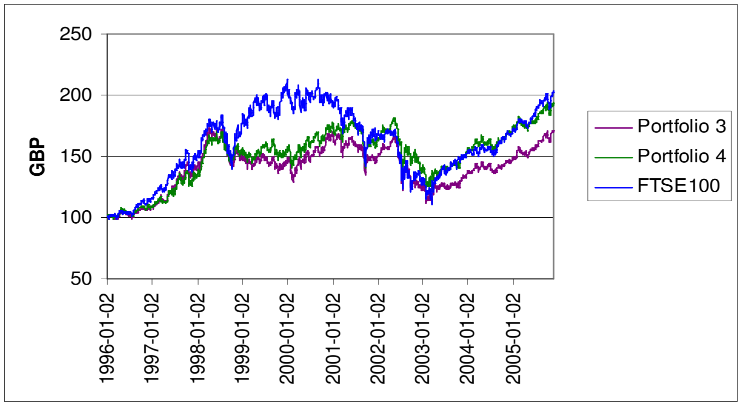

Figure 10:Price development of portfolios 3, 4 and FTSE100

On an absolute level, these portfolios perform worse than, or almost as good as the FTSE100. Portfolio 3 grows 71 percent over the whole measurement period while portfolio 4 and the index grow 93 and 103 percent respectively. This can be translated into an average yearly return of 7 and 8 percent for portfolios 3 and 4, which is to be compared to 9 percent for the index[17].

Year	Portfolio 3	Portfolio 4	FTSE100	Portfolio 3 - FTSE100	Portfolio 4 - FTSE100
1996	9%	10%	17%	-8%	-7%
1997	30%	23%	29%	1%	-5%
1998	6%	15%	18%	-11%	-3%
1999	-1%	2%	21%	-22%	-19%
2000	14%	12%	-8%	22%	20%
2001	-12%	-8%	-14%	2%	6%
2002	-14%	-12%	-21%	6%	9%
2003	8%	10%	17%	-9%	-7%
2004	9%	11%	11%	-3%	0%
2005	14%	10%	18%	-4%	-7%
Average	7%	8%	9%	-2%	-1%

Table 5: Yearly returns of portfolios 3, 4 and FTSE100

The volatilities of the portfolios are lower than their maximized-Sharpe-ratio counterparts, which is in line with what can be expected. Portfolios 3 and 4 both have an average standard deviation of 14 percent, which is to be compared to 17 percent for the index.

[17] Return averages measured as geometric averages.

Year	Portfolio 3	Portfolio 4	FTSE100	Portfolio 3 - FTSE100	Portfolio 4 - FTSE100
1996	8%	8%	9%	-1%	-1%
1997	13%	14%	15%	-2%	-2%
1998	16%	15%	21%	-5%	-6%
1999	14%	14%	18%	-4%	-4%
2000	16%	16%	19%	-3%	-3%
2001	15%	16%	22%	-6%	-6%
2002	22%	22%	28%	-6%	-6%
2003	16%	17%	19%	-3%	-3%
2004	9%	9%	10%	-2%	-1%
2005	8%	8%	9%	0%	-1%
Average	14%	14%	17%	-3%	-3%

Table 6: Yearly standard deviations of portfolios 3, 4 and FTSE100

The forecasts of the total volatilities of the portfolios are again relatively good. Portfolio 3 has a R^2-measure of 38 percent and portfolio 4 achieves 33 percent. These results are statistically significant with p-values of $8{,}84 \cdot 10^{-14}$ and $1{,}06 \cdot 10^{-4}$ for portfolio 3 and 4 respectively.

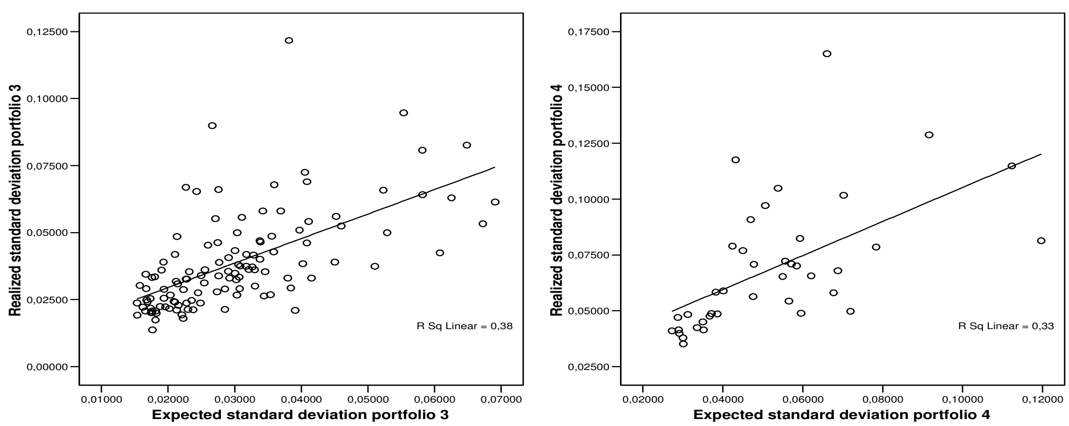

.

Figure 11: Linear regression of the predictive power of the volatility forecasts of portfolios 3 and 4

On a risk-adjusted basis, both portfolios perform worse than their Sharpe-ratio-maximized counterparts. Portfolios 3 and 4 offer a monthly Sharpe ratio of 0,14 and 0,17 respectively, which is to be compared to 0,18 for the index. In relation to FTSE100 this implies an almost 25 percent lower Sharpe ratio for portfolio 3 and barely 8 percent lower for portfolio 4.

Year	Portfolio 3	Portfolio 4	FTSE100	Portfolio 3 - FTSE100	Portfolio 4 - FTSE100
1996	0,11	0,14	0,35	-0,24	-0,21
1997	0,50	0,36	0,44	0,05	-0,08
1998	0,09	0,24	0,22	-0,13	0,02
1999	-0,11	-0,04	0,30	-0,40	-0,34
2000	0,30	0,25	-0,09	0,39	0,34
2001	-0,26	-0,15	-0,27	0,02	0,13
2002	-0,08	-0,08	-0,25	0,17	0,17
2003	0,16	0,24	0,31	-0,15	-0,07
2004	0,20	0,34	0,23	-0,02	0,11
2005	0,44	0,37	0,57	-0,14	-0,21
Average	0,14	0,17	0,18	-0,04	-0,01

Table 7: Monthly average Sharpe ratios of portfolios 3, 4 and FTSE100

Running regression 10 for the two portfolios gives monthly alpha values of –0,09 and 0,03 percent for portfolios 3 and 4 with the corresponding p-values of 0,6 and 0,9, which make them statistically insignificant. The conclusion of these results is that even though the volatility is reduced it is not enough to offset the high price of lower returns, which is consistent with economic theory.

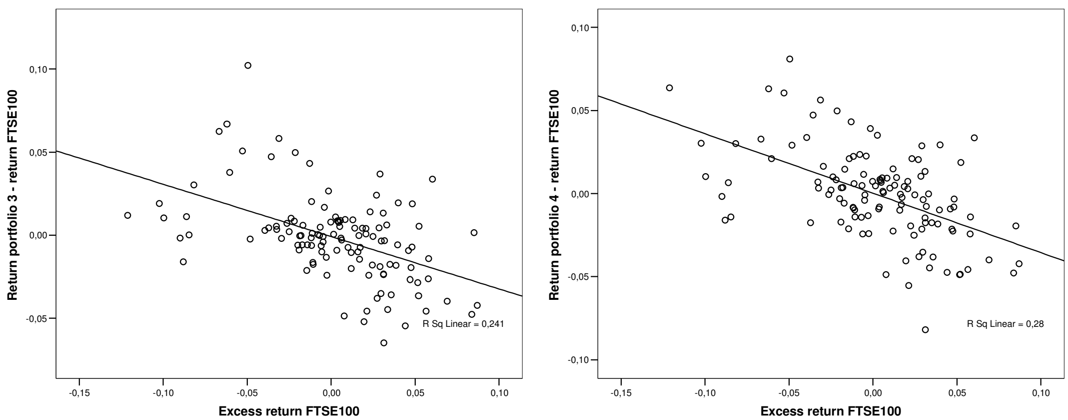

Figure 12: Linear regression of the performance of portfolio 3 and 4 in excess of the FTSE100

4.3. *The usefulness of the strategies*

The reason why the portfolios within each optimization procedure perform quite similar is of course that the forecasts generate very similar weights allocated to each stock. A possible explanation for this could be the restrictions on the amount of capital allocated to each asset. All the portfolios can at any given time take positive positions in all of the 30 different stocks. However, on average the portfolios concentrate their possession to only 19 stocks. Since the concentration to each stock is limited, it is quite possible that the different strategies would generate less similar portfolio weights if these limits were not present.

It is also interesting to note that the results seem to indicate that, in relation to FTSE100, all portfolios perform their best during the years when the index is experiencing negative returns.

The fact that both Sharpe-ratio-optimized portfolios, on average, were able to perform a higher risk-adjusted return than their benchmark index could be due to purely random effects. However, it could also be explained in terms of systematic and unsystematic risk. In this case the higher risk-adjusted return could be explained as a consequence of a successful portfolio optimization. This would imply that the FTSE100 is assumed not to be optimally diversified, thereby containing a degree of unsystematic risk. The assumption may be considered bold but not unreasonable. The reason for this is that the 10 largest stocks in the FTSE100 consistently make up more than 50 percent of the total index-value. Furthermore, the industry representation is relatively skewed in the sense that a few industries[18], which are strongly correlated with each other, are heavily represented among the companies with the greatest weights.

[18] The three largest industries constitute about 50 percent of the total weight during the time period.

5. Conclusions

In contrast to several other academic reports regarding trading strategies, which show that there are possibilities in improving the risk- and return relationship, the results of this study does not find such an improvement to be statistically significant. This implies that no reliable conclusions regarding the potential abilities of the trading strategies can be made. However, this does not mean that the study is merely a fruitless attempt to achieve higher risk-adjusted returns. An investor with the ambition of, on average, performing well in relation to the FTSE100 index would have been able to do so, during this time-period of ten years, by doing monthly rebalances and maximizing the expected Sharpe ratio. This means that the results should be considered to be of economic significance. Furthermore, even greater differences in performance should be attainable by relaxing the portfolio restrictions commonly associated with mutual funds and thereby structure the portfolio mandate like a hedge fund.

The results also indicate that the cost, in terms of lower realized returns, of minimizing volatility is relatively high – rendering a lower risk-adjusted return than maximizing the expected Sharpe ratio, which is in line with the reasoning of modern portfolio theory.

Finally, some concerns about the results are also justified. In the process of excluding trading costs, such as bid-ask spreads and commissions, the results will of course be biased upwards.

References

Books

Barnett, W., Kirman, A., Salmon, M., *Nonlinear dynamics and economics: Proceedings of the Tenth International Symposium in Economic Theory and Econometrics* Cambridge university press, tenth edition, United Kingdom, 1996.

Bodie Z., Kane A., Marcus A., *Investments*, McGraw-Hill, sixth edition, 2005.

Box, G.E.P. , G.M. Jenkins, och G.C. Reinsel, *Time series analysis: Forecasting and control*, Third edition, Prentice Hall, 1994.

Brockwell P. och Davis R., *Introduction to time series and forecasting*, Springer-Verlag, New York, 1996.

Brooks Chris, *Introductory econometrics for finance*, first edition, Cambridge university press, Cambridge, 2002.

Gourieroux, C., *ARCH models and financial applications*, Springer-Verlag, 1997.

Gujarati Damodar, *Basic econometric,* fourth edition, McGraw-Hill, 2003.

Knight J., Satchell S., *Return distribution in finance*, Reed educational and professional publishing Ltd, 2001.

Straumann, D., *Estimation in conditionally heteroscedastic time series models*, Springer, Germany, 2005.

Articles

Aas K., Dimakos X., *Statistical modelling of financial timer series: an introduction*, Norwegian computer center, 2004.

Bansal R., Dahlquist M., Harvey R., *Dynamic trading strategies and portfolio choice*, Working paper 10820, NBER working paper series, 2004.

Bollerslev, T. Chou, R.Y. Kroner, K.F. ”*ARCH modelling in finance”*, Journal of econometrics, 52, 5-59. ,1986.

Bollerslev, T., Chou R., Kroner K., *ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence*, Journal of Econometrics, Vol.52, No.1, 1992.

Christie, A.A., *The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects,* Journal of Financial Economics, 10, 1982.

Conrad, J. Dittmar, R. Dong-Hyun, A., *Risk Adjustment and Trading Strategies,* Review of financial studies, 16, 2003.

Guo D., *Dynamic Volatility Trading Strategies in the Currency Option Market Using Stochastic Volatility Forecasts*, Centre Solutions, Zurich Financial Service Group, 1999.

Day T.E., Lewis C.M., *Stock Market Volatility and the Information Content of Stock Index Options*, Journal of Econometrics, 52, 267–287. 1992.

Engström S., *Does active portfolio management create value? – An evaluation of fund managers' decisions*, SSE/EFI working paper series in Economics and Finance no 553.

Ghysels E., Santa-Clara P., Valkanov R., *Predicting volatility: Getting the most out of return data sampled at different frequencies*, Working paper 10914, NBER working paper series, 2004.

Jones, R.H. *Fitting autoregressions*, J.Amer. Stat. Assoc., 70, 1975.

Kim S., Shepard N., Chib S., *Stochastic volatility: likelihood inference and comparison with ARCH models*, Review of economic studies, 65, 1998.

Lamoureux C, Lastrapes W, *Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities*, Review of Financial Studies 6, 293-326. 1993.

Lopez J., Walter C., *Evaluating Covariance Matrix Forecasts in a Value-at-Risk Framework*, Working paper 2000-21, Federal Reserve Bank of San Francisco, 2000.

Markowitz H. *Portfolio selection*, Journal of finance, 1952.

Mina J., Xiao J., *Return to RiskMetrics: The evolution of a standard*, RiskMetrics, 2001.

Nelson, D., *Conditional heteroscedasticity in asset returns: A new approach*, Econometrica, 59, 347-370, 1991.

Shibata, R. *Selection of the order of an autoregressive model by Akaike's information criterion*, Biometrika, 1976.

Solnik B., *Why not diversify internationally rather than domestically?*, Financial analyst journal 30, 1974.

Internet

Dagens Industri, *http://di.se/Nyheter/?page=%2fAvdelningar%2fArtikel.aspx%3fO%3dIndex %26ArticleId%3d2003%5c08%5c12%5c84022%26src%3ddi, 2005-11-02.*

FTSE, *http://www.ftse.com*, 2005-10-05.

Scandinavian information exchange, *http://www.six.se*, 2005-10-03.

Appendix A1 – Results of portfolio 1

Return

Year	Portfolio 1	FTSE100	Portfolio 1 - FTSE100
1996	15%	17%	-2%
1996	41%	29%	12%
1996	21%	18%	4%
1996	12%	21%	-9%
1996	3%	-8%	11%
1996	-2%	-14%	12%
1996	-6%	-21%	15%
1996	12%	17%	-5%
1996	6%	11%	-5%
1996	16%	18%	-2%
Average	12%	9%	3%

Standard deviation

Year	Portfolio 1	FTSE100	Portfolio 1 - FTSE100
1996	10%	9%	0%
1997	17%	15%	1%
1998	20%	21%	-1%
1999	19%	18%	1%
2000	23%	19%	4%
2001	17%	22%	-5%
2002	22%	28%	-5%
2003	18%	19%	-2%
2004	10%	10%	0%
2005	11%	9%	2%
Average	17%	17%	0%

Sharpe ratio

Year	Portfolio 1	FTSE100	Portfolio 1 - FTSE100
1996	0,20	0,35	-0,15
1997	0,48	0,44	0,04
1998	0,26	0,22	0,04
1999	0,09	0,30	-0,21
2000	0,13	-0,09	0,21
2001	-0,04	-0,27	0,23
2002	0,10	-0,25	0,36
2003	0,23	0,31	-0,08
2004	0,14	0,23	-0,08
2005	0,42	0,57	-0,15
Average	0,21	0,18	0,03

T-test and correlation with FTSE100

Year	T-value	Critical value
1996	-0,1	1,8
1997	0,0	1,8
1998	0,1	1,8
1999	-0,3	1,8
2000	0,3	1,8
2001	0,3	1,8
2002	0,4	1,8
2003	-0,1	1,8
2004	-0,1	1,8
2005	-0,4	1,8
Average	0,0	1,8

Year	Correlation
1996	0,96
1997	0,98
1998	0,94
1999	0,50
2000	0,45
2001	0,59
2002	0,90
2003	0,96
2004	0,72
2005	0,96
Average	0,80

Appendix A2 – Results of portfolio 2

Return

Year	Portfolio 2	FTSE100	Portfolio 2 - FTSE100
1996	16%	17%	-1%
1996	30%	29%	2%
1996	17%	18%	0%
1996	14%	21%	-6%
1996	-4%	-8%	4%
1996	0%	-14%	14%
1996	-8%	-21%	13%
1996	14%	17%	-2%
1996	9%	11%	-2%
1996	17%	18%	-1%
Average	11%	9%	2%

Standard deviation

Year	Portfolio 2	FTSE100	Portfolio 2 - FTSE100
1996	9%	9%	0%
1997	17%	15%	2%
1998	21%	21%	0%
1999	19%	18%	1%
2000	24%	19%	4%
2001	17%	22%	-5%
2002	23%	28%	-5%
2003	18%	19%	-1%
2004	11%	10%	0%
2005	10%	9%	1%
Average	17%	17%	0%

Sharpe ratio

Year	Portfolio 2	FTSE100	Portfolio 2 - FTSE100
1996	0,25	0,35	-0,10
1997	0,36	0,44	-0,09
1998	0,23	0,22	0,01
1999	0,09	0,30	-0,21
2000	-0,03	-0,09	0,06
2001	-0,04	-0,27	0,23
2002	0,09	-0,25	0,34
2003	0,23	0,31	-0,08
2004	0,21	0,23	-0,01
2005	0,47	0,57	-0,10
Average	0,19	0,18	0,01

T-test and correlation with FTSE100

Year	T-value	Critical value
1996	-0,1	1,8
1997	-0,1	1,8
1998	0,0	1,8
1999	-0,4	1,8
2000	0,1	1,8
2001	0,3	1,8
2002	0,4	1,8
2003	-0,1	1,8
2004	0,0	1,8
2005	-0,1	1,8
Average	0,0	1,8

Year	Correlation
1996	0,96
1997	0,98
1998	0,94
1999	0,50
2000	0,48
2001	0,41
2002	0,91
2003	0,96
2004	0,68
2005	0,57
Average	0,74

Appendix A3 – Results of portfolio 3

Return

Year	Portfolio 3	FTSE100	Portfolio 3 - FTSE100
1996	9%	17%	-8%
1996	30%	29%	1%
1996	6%	18%	-11%
1996	-1%	21%	-22%
1996	14%	-8%	22%
1996	-12%	-14%	2%
1996	-14%	-21%	6%
1996	8%	17%	-9%
1996	9%	11%	-3%
1996	14%	18%	-4%
Average	7%	9%	-2%

Standard deviation

Year	Portfolio 3	FTSE100	Portfolio 3 - FTSE100
1996	8%	9%	-1%
1997	13%	15%	-2%
1998	16%	21%	-5%
1999	14%	18%	-4%
2000	16%	19%	-3%
2001	15%	22%	-6%
2002	22%	28%	-6%
2003	16%	19%	-3%
2004	9%	10%	-2%
2005	8%	9%	0%
Average	14%	17%	-3%

Sharpe ratio

Year	Portfolio 3	FTSE100	Portfolio 3 - FTSE100
1996	0,11	0,35	-0,24
1997	0,50	0,44	0,05
1998	0,09	0,22	-0,13
1999	-0,11	0,30	-0,40
2000	0,30	-0,09	0,39
2001	-0,26	-0,27	0,02
2002	-0,08	-0,25	0,17
2003	0,16	0,31	-0,15
2004	0,20	0,23	-0,02
2005	0,44	0,57	-0,14
Average	0,14	0,18	-0,04

T-test and correlation with FTSE100

Year	T-value	Critical value
1996	-0,23	1,8
1997	0,05	1,8
1998	-0,13	1,8
1999	-0,76	1,8
2000	0,40	1,8
2001	0,02	1,8
2002	0,21	1,8
2003	-0,16	1,8
2004	-0,03	1,8
2005	-0,17	1,8
Average	-0,08	1,8

Year	Correlation
1996	0,94
1997	0,98
1998	0,87
1999	0,20
2000	0,32
2001	0,93
2002	0,96
2003	0,91
2004	0,85
2005	0,82
Average	0,78

Appendix A4 – Results of portfolio 4

Return

Year	Portfolio 4	FTSE100	Portfolio 4 - FTSE100
1996	10%	17%	-7%
1996	23%	29%	-5%
1996	15%	18%	-3%
1996	2%	21%	-19%
1996	12%	-8%	20%
1996	-8%	-14%	6%
1996	-12%	-21%	9%
1996	10%	17%	-7%
1996	11%	11%	0%
1996	10%	18%	-7%
Average	7%	9%	-2%

Standard deviation

Year	Portfolio 4	FTSE100	Portfolio 4 - FTSE100
1996	8%	9%	-1%
1997	14%	15%	-2%
1998	15%	21%	-6%
1999	14%	18%	-4%
2000	16%	19%	-3%
2001	16%	22%	-6%
2002	22%	28%	-6%
2003	17%	19%	-3%
2004	9%	10%	-1%
2005	8%	9%	-1%
Average	14%	17%	-3%

Sharpe ratio

Year	Portfolio 4	FTSE100	Portfolio 4 - FTSE100
1996	14%	35%	-21%
1997	36%	44%	-8%
1998	24%	22%	2%
1999	-4%	30%	-34%
2000	25%	-9%	34%
2001	-15%	-27%	13%
2002	-8%	-25%	17%
2003	24%	31%	-7%
2004	34%	23%	11%
2005	37%	57%	-21%
Average	17%	18%	-1%

T-test and correlation with FTSE100

Year	T-value	Critical value
1996	-0,20	1,8
1997	-0,08	1,8
1998	0,02	1,8
1999	-0,49	1,8
2000	0,40	1,8
2001	0,20	1,8
2002	0,23	1,8
2003	-0,08	1,8
2004	0,10	1,8
2005	-0,37	1,8
Average	-0,03	1,8

Year	Correlation
1996	0,92
1997	0,98
1998	0,84
1999	0,31
2000	0,36
2001	0,81
2002	0,95
2003	0,95
2004	0,92
2005	0,59
Average	0,76

Appendix B – Generation of GARCH-forecasts

1. Model selection and structure

Under the assumption that a return series is heteroscedastic, i.e. that the volatility varies with time, it is appropriate to use a GARCH-model for forecasting the volatility. Whether a time series is to be considered heteroscedastic or not can be determined by using the Engle's ARCH-test. More information regarding this test can be found in Appendix C together with a selection of the results of the test.

In order to increase the reliability of the forecasts the model needs to be adjusted to the data at hand. As a result of the discussion in paragraph 2.1 - Financial time series, an EGARCH-model is used for modelling the variance of the return series. Along this reasoning the series is therefore assumed to follow a t-distribution.

1.1. Initial parameter estimates

In order to optimize the parameters which make out the ARMA- and EGARCH processes a vector of initial parameter estimates is needed. (Brockwell et al., 1996). The parameters which are needed [C, ϕ_j, θ_i, K, G_i, A_j, L_j], stem from the equations (4) and (9). The first three parameters are ARMA-parameters and the remaining four are EGARCH.

1.1.1. ARMA *(p,q)*

The autoregressive component, ϕ_j, is estimated by creating an auto-covariance matrix and solving the Yule-Walker equations. By using the estimated coefficients the series is filtered in order to obtain a pure moving average process. The auto-covariance sequence of the moving average is estimated and used in order to iterate the moving average component, θ_i. Finally, the variance of the residual is estimated. (Box et al., 1994)

1.1.2. EGARCH *(p,q)*

In contrast to the statistical analysis made in order to estimate the ARMA-parameters, the EGARCH-parameters are solely based on empirical analysis of financial time series and are thereby Ad Hoc in their nature.

The following assumptions have been made

$$G_1 + G_2 + \ldots + G_P = 0{,}9$$
$$A_1 + A_2 + \ldots A_Q = 0{,}2$$
$$L_i = 0 \qquad 1 \le i \le Q$$

The initial estimates for the EGARCH(1,1) are thereby made according to:

$$\ln \sigma_t^2 = K + 0.9 \ln \sigma_{t-1}^2 + 0.2[|z_{t-1}| - E(|z_{t-1}|)] \tag{11}$$

The estimate of the constant *K* is made by using the relationship between the independent variance of the residual process, σ^2, and the G_i-parameter in the EGARCH(1,1)-model according to:

$$K = (1 - G_1)\ln \sigma^2$$
$$K = (1 - 0{,}9)\ln \sigma^2$$
$$K = 0{,}1 \ln \sigma^2$$

1.2. Maximum likelihood estimation

When the starting parameters are estimated a conditional[19] log-likelihood estimation[20] (MLE) can be made. The estimation optimizes the output by iteratively adjusting the data being used as input in the log-likelihood function until a satisfactory degree of exactitude has been reached or, alternatively, the model does not manage to converge a solution. The probability of convergence is limited due to four factors, which are presented in the table below together with their respective limitations.

Modelled restrictions		**Limitation**
Maximum number of iterations (MaxIT)	=	400
Maximum number of MLE-evaluations (MaxMLE)	=	800
Tolerance level of restriction breach (MaxREST)	=	1×10^{-6}
Tolerance level of the function value (MaxFUNK)	=	1×10^{-6}

All factors terminate the convergence optimization by finishing the process prematurely given that a convergence has not already occurred when the limitation is breached.

[19] Conditioned, in this case, means that the model demands a selection of historical data in order to start.
[20] This type of optimization is often referred to as non-linear programming.

MaxIT, as well as MaxMLE, are purely mechanical in their nature. The value of their respective restrictions is chosen in light of the fact that, due to their size, they rarely prevent the process to converge prematurely. A discontinued conversion, as a result of the restrictions being breached, is therefore a signal that the chosen model does not describe the data in a satisfactory way.

The tolerance parameters (MaxREST and MaxFUNK) affect how and when conversion is reached and are therefore capable to affect the solution. MaxREST describes the model tolerance against a violation of the given restrictions[21]. Concretely, MaxREST represents the maximum value with which a parameter estimation can breach a restriction and still enable a successful conversion. MaxFUNK describes the model tolerance placed upon the log-likelihood function. A successful conversion is achieved when the change in the value of the log-likelihood function is less than the MaxFUNK.

2. Forecasting

The generated forecasts are estimated by using the minimum mean-square error of the conditional return series (r_t) and the standard deviation of the residuals (ε_t). Practically this is done by examining the ARMA- and EGARCH-models through a linear filter from which the model, by iterations, generate conditioned expectations of the equations one forecast period at a time.

As the models are retrospective in their nature they need presampled data in order to start the iteration process. The presampled data plays the same input roll as in the estimation of the ARMA- and EGARCH-parameters.

[21] The restrictions control, among other things, that the return series is stationary and that the variance is always positive by verifying that the eigen-values of the AR- and MA-polynomial is within the unit-circle. Furthermore, they also make sure that the eigen-values of the EGARCH polynomial are within the unit-circle.

Appendix C – Analysis of data

1. Extreme values

All GARCH-models perform their best when the market is relatively stable. (Gourieroux, 1997). This means that the models are not capable to model unexpected irregular phenomenons such as stock market crashes. Since the forecasts are based on historical data, this type of shocks can have a serious impact for a relatively long period of time, which of course can generate distorted results. (Aas et al., 2004). As a result, all return series have been removed of extreme values on the 5 percent level using Grubbs test[22].

The removal of extreme values is, however, not entirely unproblematic since it has a negative influence on the volatility and affects the mean value. Furthermore, there is always a risk that new extreme values will replace the removed ones.

2. Autocorrelation Function

In order to estimate the degree of dependence in the return series, and thereby choose the most appropriate model, an autocorrelation function (ACF) has been used. An ACF plot around zero for all lags except lag zero indicates that the time series can be properly modelled as an IID noise (Brockwell, 1996). If the ACF plot exhibit decaying amplitude as the size of the lag increases, this signals that the time series contains a trend component. (Straumann, 2005). By a similar procedure, one can also examine whether the series has a reoccurring trend component (Ibid.). This means that the ACF can be used as an indicator for non-stationarity.

Concretely this means that if a few ACF markers are outside the 95 percent interval during the first 40 lags or, alternatively, one of them exists far beyond the interval, then the assumption of IID-noise can be rejected, which means that a MA(q)-model can be used.

The appearance of the outliers and their location should be decisive for which type of MA-model is to be used. If, for instance, the ACF-plot looks like the figure below, then a MA(1) is appropriate since no outliers are located after the first lag[23].

---

[22] For more information of how the test is formally executed, see for instance: www.itl.nist.gov/div898/handbook/eda/section3/eda35h.htm

[23] Formally, 95 percent of the autocorrelated data after lag 1 should be within the confidence interval when using a long time series.

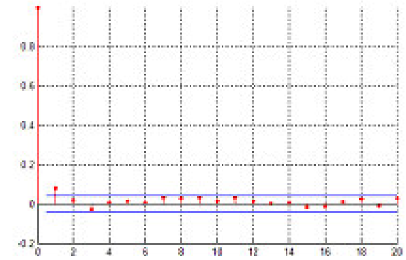

Figure 13: ACF of the return

3. Partial Autocorrelation function (PACF)

Just as the ACF can be used to determine the order of the MA-model, the PACF can be used to determine the order of AR-model to be used. The evaluation is done in a similar way as with the MA-model.

ACF, as well as PACF, should however be used with care since individual ACF values can exhibit large variances and also be autocorrelated. (Box et al., 1994).

4. ACF of the quadratic return

In order to examine the correlation structure of the second moment of a return series, i.e. the variance, an ACF of the quadratic return can be used. As previously mentioned, the volatility is assumed to be autocorrelated. This can easily be seen in the figure below. in paragraph

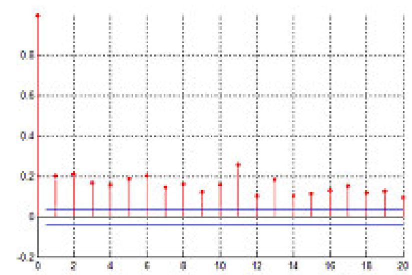

Figure 14: ACF of the variance

If the ACF fades away as the lag-length increases, this is a sign that the variance process is close to be non-stationary. (Brockwell, 1996)

5. Engle's ARCH-test

An Engle's ARCH-test has been done in order to quantify the degree of heteroscedasticity. (Barnett, 1996). If the null hypothesis of no ARCH-effect is accepted then the time series of sample residuals is assumed to be following a normal distribution.

The test is carried out at the 5 percent level and examines 10, 15 and 20 lags. The results, which show clear signs of heteroscedasticity in all time series, can be seen after the figures of the return, ACF and PACF.

6. AIC

As with the ARCH-test for heterscedasticity, an AIC-test can be performed in order to determine the order of the lags *p* and *q*. The test tries different orders of *p* and *q* by simultaneously taking the number of estimated parameters into account (Brockwell, 1996). By doing so, the test can to some extent be compared to the "R^2-Adjusted" obtained from a regular multiple regression because it takes the number of estimated parameters into account.

The test does, however, have a built in bias. This is because it tends to estimate too high a value of the parameter *p*. (Jones, 1975, Shibata, 1976). This overfitting of the data has lead to common use of complementary models, such as the AICC and the BIC. (Brockwell, 1996). The implication of this is that the results of the AIC-test are to be considered more of a guideline rather than an absolute answer to the appropriate values of *p* and *q*. As a result, a general rule of thumb has been created, which says that the simplest model which in a satisfactory way can explain the data should be used.

The results of the tests, where a lower AIC-value is better than a higher one, are presented at the end of this appendix. They show that the relevant time series behave quite similar, which has lead to *p* and *q* being chosen to be equal to one.

7. Returns, ACF and PACF

BAA

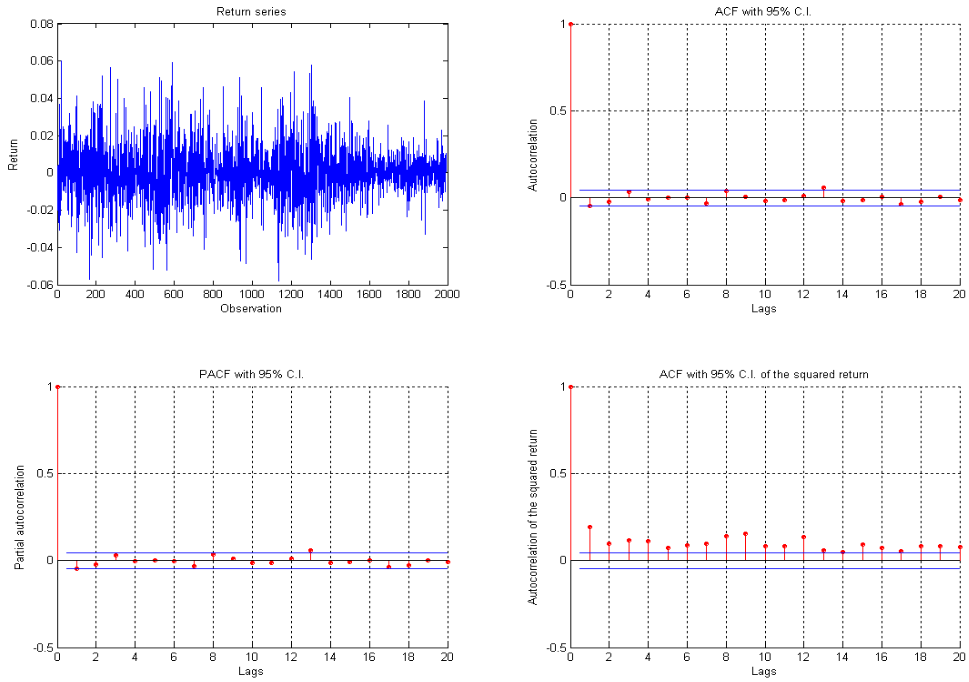

Barclays Bank

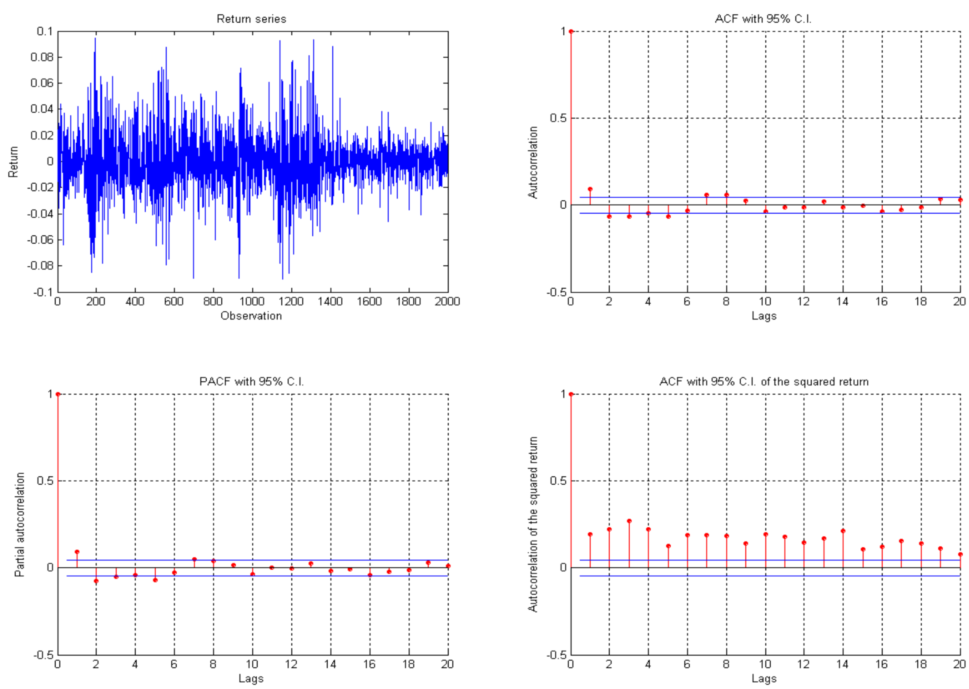

Reed International

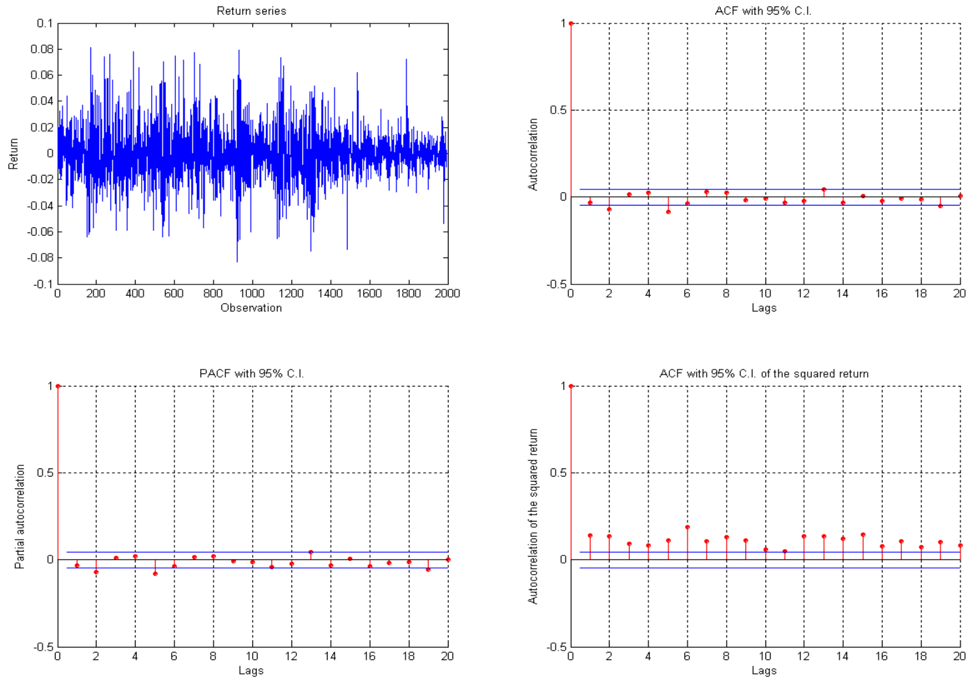

Rentokil

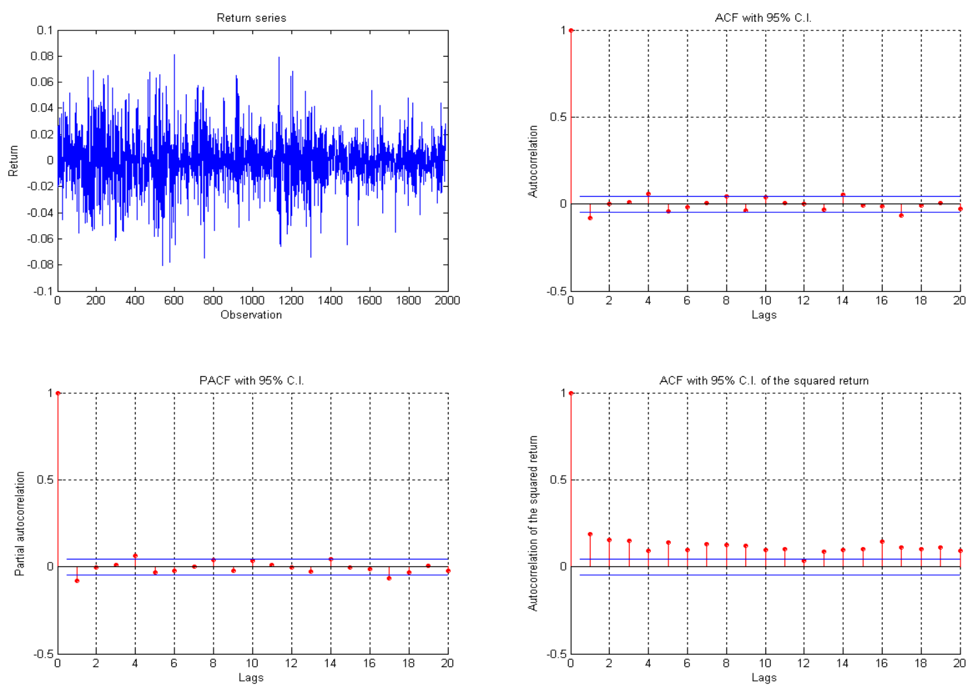

8. Results of the Engle's ARCH-test (5 percent level)

BAA

Lags	H	P-value	Arch-stat	Critical value
10	1	0	156,0	18,3
15	1	0	172,3	25,0
20	1	0	175,6	31,4

Barclays Bank

Lags	H	P-value	Arch-stat	Critical value
10	1	0	301,8	18,3
15	1	0	325,2	25,0
20	1	0	332,9	31,4

BOC

Lags	H	P-value	Arch-stat	Critical value
10	1	0	170,9	18,3
15	1	0	187,8	25,0
20	1	0	195,3	31,4

Cable & Wireless

Lags	H	P-value	Arch-stat	Critical value
10	1	0	189,8	18,3
15	1	0	199,4	25,0
20	1	0	208,5	31,4

Centrica

Lags	H	P-value	Arch-stat	Critical value
10	1	0	162,3	18,3
15	1	0	174,8	25,0
20	1	0	191,0	31,4

Diageo

Lags	H	P-value	Arch-stat	Critical value
10	1	0	173,4	18,3
15	1	0	201,6	25,0
20	1	0	219,9	31,4

HSBC

Lags	H	P-value	Arch-stat	Critical value
10	1	0	192,1	18,3
15	1	0	215,0	25,0
20	1	0	219,3	31,4

Marks & Spencer

Lags	H	P-value	Arch-stat	Critical value
10	1	0	144,8	18,3
15	1	0	172,1	25,0
20	1	0	192,2	31,4

National Grid

Lags	H	P-value	Arch-stat	Critical value
10	1	0	150,2	18,3
15	1	0	166,4	25,0
20	1	0	169,9	31,4

Reed International

Lags	H	P-value	Arch-stat	Critical value
10	1	0	153,0	18,3
15	1	0	189,7	25,0
20	1	0	195,0	31,4

Rentokil

Lags	H	P-value	Arch-stat	Critical value
10	1	0	164,8	18,3
15	1	0	178,5	25,0
20	1	0	195,6	31,4

Sainsbury J

Lags	H	P-value	Arch-stat	Critical value
10	1	0	145,7	18,3
15	1	0	152,9	25,0
20	1	0	168,4	31,4

Tesco

Lags	H	P-value	Arch-stat	Critical value
10	1	0	146,3	18,3
15	1	0	164,1	25,0
20	1	0	169,7	31,4

Vodafone

Lags	H	P-value	Arch-stat	Critical value
10	1	0	234,6	18,3
15	1	0	253,4	25,0
20	1	0	268,5	31,4

All stocks, though not disclosed here, reject the null-hypothesis, which means that an ARCH-effect is present.

9. Results of the AIC-test

BAA

P,Q	1	2	3
1	-11476	-11478	-11479
2	-11474	-11477	-11479
3	-11472	-11474	-11477

Barclays Bank

P,Q	1	2	3
1	-10090	-10087	-10077
2	-10103	-10101	-10100
3	-10101	-10099	-10098

BOC

P,Q	1	2	3
1	-11211	-11210	-11210
2	-11209	-11208	-11208
3	-11208	-11207	-11206

Cable & Wireless

P,Q	1	2	3
1	-10946	-10833	-10983
2	-10684	-11454	-11073
3	-10736	-10916	-10781

Centrica

P,Q	1	2	3
1	-10465	-10466	-10466
2	-10464	-10465	-10464
3	-10466	-10467	-10467

Diageo

P,Q	1	2	3
1	-10888	-10888	-10888
2	-10888	-10888	-10886
3	-10887	-10886	-10885

HSBC

P,Q	1	2	3
1	-11030	-11028	-11026
2	-10994	-11028	-11026
3	-11030	-11028	-11026

Marks & Spencer

P,Q	1	2	3
1	-10461	-10463	-10466
2	-10495	-10498	-10500
3	-10494	-10497	-10499

National Grid

P,Q	1	2	3
1	-11329	-11324	-11333
2	-11327	-11326	-11331
3	-11324	-11322	-11327

Reed International

P,Q	1	2	3
1	-10340	-10338	-10341
2	-10326	-10337	-10329
3	-10337	-10335	-10327

Rentokil

P,Q	1	2	3
1	-10393	-10409	-10407
2	-10392	-10408	-10406
3	-10391	-10410	-10404

Sainsbury J

P,Q	1	2	3
1	-10564	-10565	-10563
2	-10563	-10563	-10561
3	-10563	-10563	-10570

Tesco

P,Q	1	2	3
1	-10891	-10890	-10892
2	-10890	-10888	-10890
3	-10892	-10890	-10892

Vodafone

P,Q	1	2	3
1	-9533,4	-9531,7	-9533
2	-9546,4	-9544,7	-9542,2
3	-9544,4	-9542,7	-9540,2

Glossary

ARMA-model = Linear time series-model which is suitable if { X_t } is stationary and satisfies the condition that: $X_t - \phi X_{t-1} = \mu_t + \theta\mu_{t-1}$, where $\{\mu_t\} \in WN(0, \sigma^2)$ and $\phi + \theta \neq \theta$.

Autocovariance = Covariance of one element with itself over time.

Economic significance = A situation with statistically insignificant results where the results are still of interest for an investor.

Eigen-value = The sum of the square of the columns from a factor matrix. In matrix algebra, the eigen-values of a correlation matrix are equal to the square root of the "characteristic" equation.

Non stationary = A non stationary time series is characterised by the following: a changing mean and covariance structure, periodic behaviour of trends and finally, time varying parameters and structure.

IID Noise = A time series without a trend- and season component where the observations are independent and identically distributed with the mean zero. The time series is also characterised by: $\{X_t\} = \rho_{X^2}(h) = 0$ for all $h \neq 0$, something that doesn't necessarily have to be the case with white noise.

Linear process = The time series { X_t } is a linear process if:

$$X_t = \sum_{j=-\infty}^{\infty} \psi_j Z_{t-j}, \text{ where } \{ Z_t \} \in WN(0, \sigma^2)$$

Stationary = A time series, X_t, is said to be weakly stationary if the mean is independent over time and if the autocovariance function is independent of *t* for each movement in time ($t+h$).

Unit circle = Popular educational tool. It consists of a circle with the radius 1, placed with its center in the origin of a two dimensional coordinate system.

White noise = A time series of uncorrelated random variables which follows a distribution with mean 0 and variance σ^2. This implies that { X_t } is stationary with zero autocorrelation.

Yule-Walker estimation = Equation for estimation of the preliminary AR-parameter, ϕ, the variance of white noise. Other models with the same purpose are: Burg's algorithm, the innovations process and the Hannan-Rissaneb algorithm.

Notation

$\{X_t\}$ = Time series X at time *t*.

$\gamma_x(a,b)$ = The covariance function of (X_t) between time *a* and *b*.
= $E[(X_a - \mu_x(r))(X_b - \mu_x(s))]$

$\hat{\gamma}_x(h)$ = Estimation of the covariance function for (X_t)
= $n^{-1}\sum_{t=1}^{n-|h|}(x_{t+|h|} - \bar{x})(x_t - \bar{x}), \quad -n < h < n$

$\rho_x(h)$ = Autocorrelation function of (X_t)
≡ $\frac{\gamma_x(h)}{\gamma_x(0)}$

$\hat{\rho}_x(h)$ = Estimation of the autocorrelation function for (X_t)
= $\frac{\hat{\gamma}_x(h)}{\hat{\gamma}_x(0)}, \quad -n < h < n$

θ = Parameter in the MA-model of which size, measured in absolutes, is less than 1 (In order for ARMA(1,1) to be invertible).

ϕ = Parameter in the AR-model.
$|\phi| < 1$

ψ_j = Constant in a linear process, which in a sequence of constants fulfils the condition:
$\sum_{j=-\infty}^{\infty}\psi < 0$

Γ_T = Gamma function