UNIVERSITY
OF TURKU
Turku School of
Economics

Mean reversion in the Finnish stock market and its effect on risk

Accounting and Finance
Master's thesis

Author(s):
Otto Pajala

Supervisor(s):
Professor Luis Alvarez Esteban

29.09.2024
Turku

Master's thesis

Subject: Finance
Author(s): Otto Pajala
Title: Mean reversion in the Finnish stock market and its effect on risk
Supervisor(s): Professor Luis Alvarez Esteban
Number of pages: 51 pages
Date: 03.09.2024

The debate around the predictability of stock prices has been going on for decades and wealth managers have been scrambling to find any type of long-term edge over the market. This search inspired De Bondt and Thaler (1985) to compare returns between a portfolio of past losers and a portfolio of past winners. They found that past losers tended to have higher than market average returns in the following years. The phenomenon causing this is now known as reversal which has had a profound impact on how we think about the behaviour of stock prices and sparked numerous studies on the topic of mean reversion.

The aim of this thesis is to find out whether the Finnish stock market (OMXHPI-index) is mean reverting and if so, how long does it take for stock returns to revert to their mean (50-day moving average) and more importantly, how does this reversal affect risk. The presence of mean reversion is studied using a modified version of Lo and MacKinlay's (1988) variance ratio and the Hurst exponent and half-life. The risk measures chosen are value at risk and expected shortfall, otherwise known as conditional value at risk.

The results show that the Finnish stock market is mean reverting, and it takes 8 to 15 trading days for stock price changes to revert back to their mean. The Hurst exponent and half-life have a low positive correlation with the risk measures. Additionally, they are statistically significant methods of explaining the changes with the risk measures based on the robust standard errors test. This means that as the amount of mean reversion increases, the amount of risk in the market also increases. Changes in half-life had a stronger impact on risk. The findings of this thesis can be valuable for portfolio managers looking to hedge risk or opportunistic investors aiming to outperform the market.

Key words: mean reversion, value at risk, expected shortfall, variance ratio, Hurst exponent, half-life, risk management

Pro gradu -tutkielma

Oppiaine: Rahoitus
Tekijä(t): Otto Pajala
Otsikko: Keskiarvohakuisuus Suomen osakemarkkinoilla ja sen vaikutus riskiin
Ohjaaja(t): Professori Luis Alvarez Esteban
Sivumäärä: 51 sivua
Päivämäärä: 03.09.2024

Osakkeiden hinnanmuutoksien ennustettavuudesta on käyty keskustelua jo vuosikymmeniä, ja varainhoitajat pyrkivät löytämään pitkän aikavälin etulyöntiaseman markkinoihin nähden. Tämä innoitti De Bondt ja Thaleria (1985) vertailemaan aiempia häviäjiä sisältävän salkun ja aiempia voittajia sisältävän salkun tuottoja. He havaitsivat, että aiemmilla häviäjillä oli taipumus tuottaa markkinoita paremmin seuraavina vuosina. Tätä ilmiötä selitetään nykyään keskiarvohakuisuudella ja sillä on ollut suuri vaikutus siihen, mitä ajattelemme osakekurssien käyttäytymisestä.

Tämän tutkimuksen tavoitteena on selvittää, ovatko Suomen osakemarkkinat (OMXHPI-indeksi) keskiarvohakuisia, ja jos ovat, kuinka kauan kestää, että osakkeiden tuotot palautuvat keskiarvoonsa (50 päivän liukuva keskiarvo), sekä miten keskiarvohakuisuus vaikuttaa riskiin. Keskiarvohakuisuuden esiintymistä tutkitaan käyttämällä muunneltua versiota Lo ja MacKinlayn (1988) varianssisuhteesta sekä Hurstin eksponentista ja puoliintumisajasta. Riskimittareiksi valittiin value at risk ja odotettu alijäämä, joka tunnetaan myös nimellä ehdollinen value at risk.

Tulokset osoittavat, että Suomen osakemarkkinat ovat keskiarvohakuisia ja kestää 8–15 kaupankäyntipäivää ennen kuin osakekurssien muutokset palautuvat keskiarvoonsa. Hurstin eksponentilla ja puoliintumisajalla on matala positiivinen korrelaatio riskimittareiden kanssa. Lisäksi ne ovat tilastollisesti merkitseviä selittäjiä riskimittareiden muutoksille robustin keskivirhetestin perusteella. Kun keskiarvohakuisuus voimistuu, myös riskin määrä markkinoilla kasvaa. Puoliintumisajan muutoksilla oli voimakkaampi vaikutus riskiin. Tämän tutkielman tulokset voivat olla arvokkaita salkunhoitajille, jotka pyrkivät suojautumaan riskeiltä tai lisätuottoa hakeville opportunistisille sijoittajille.

Avainsanat: keskiarvohakuisuus, value at risk, odotettu alijäämä, varianssisuhde, Hurstin eksponentti, puoliintumisaika, riskienhallinta

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 Introduction

The debate on whether stock prices change in specific patterns has been going on for decades. If these predictable patterns would be found, they could be used in numerous trading strategies that could give investors lucrative returns. In financial theory it is often assumed that the prices of stocks represent all available information and the future changes in prices could not be predicted. This does not mean, for example, that directional predictability would be impossible. Additionally, it is generally accepted that, for investors to obtain meaningful returns, they would have to expose themselves to risk.

A new development in financial literature started to change the way people think about stock returns in the 1980s. Several separate studies indicated that historical prices of stocks could be, to some degree, used to forecast future prices. One of these studies was by De Bondt and Thaler (1985) that found that there is too much weight placed on new information and not enough on past information by investors. They called this phenomenon the Overraction Hypothesis. They found that past losers (stocks that underperformed the market) tended to outperform the overall market in the following 3- and 5-year period. They argued that this could be seen as evidence that prices of stocks were mean reverting with respect to the market, at least in longer time frames.

Jegadeesh and Titman (1993) found indications of the contrary happening in shorter time frames. In these shorter periods, they found that stock prices often persist in their current trend, which is what they called return persistence. By creating portfolios funded through the proceeds of selling underperforming stocks and investing in the outperforming ones, they demonstrated that this approach yields considerable gains over periods of three to twelve months. The research indicated that stocks with recent success are likely to maintain their positive trajectory for the subsequent three to twelve months, showing straightforward evidence of momentum within short-term price movements.

These two studies made researchers re-evaluate the true efficiency of stock markets. If an investor was able to detect these overreactions, they could incorporate this information into their own investment strategy. Such an approach could yield substantial returns while only using historical data. This would suggest a potential contradiction to the theory of weak form efficiency in the financial market.

1.1 Research questions

This thesis aims to study the existence of mean reversion in the Finnish stock market and if there is mean reversion, what is its effect on selected risk measures. In financial literature, mean reversion is often referred to as the tendency for equity prices to fluctuate around the mean price or trendline over time. The thesis will aim to find an answer for the following research questions:

- Does (relative) mean reversal occur in equity prices?
- If so, how long does it take for prices to reverse back to their mean?
- What is the relationship between mean reversion and risk? In this thesis, risk will be measured using value at risk and expected shortfall.

1.2 Structure of the thesis

The thesis is structured as follows. In the second chapter the efficient market hypothesis and random walk are presented. The theory behind them is explained to give the reader a basic understanding of the concepts. The third chapter explains mean reversion and some key findings related to it. In the fourth chapter the properties of financial time series are introduced. The methodology is described in the fifth chapter. This includes the description of the data used and tests performed. In chapter six, the results of the tests are presented. Finally, in chapter seven, conclusions of this thesis and suggestions for future works are made.

2 Theory

2.1 Efficient market hypothesis

The groundwork for the efficient market hypothesis (EMH) was laid by Samuelson (1965) and Mandelbrot (1966) when they studied the role of expected return models and their relationship with the efficient markets theory. Fama (1970) built upon this work to form the EMH as it is now known. The hypothesis has had a major influence on the practices and theory in the field of finance and has been extensively referenced across various studies and papers. It states that the stock markets are deemed efficient when share prices instantaneously and accurately incorporate all available information at any given moment, thus reflecting the true value of securities without any delay. For adequate market efficiency to be achieved, Fama (1970) lists three conditions:

1. No transaction costs when trading securities.
2. All available information is available to everyone in the market without cost.
3. All participants are homogeneous in what the implications of the current available information are to the stock's price and how the future prices will be distributed across each security.

Fama (1970) acknowledges that the real-world market does not align with the ideal of a frictionless market where information is uniformly accessible and interpretations of its impact on security prices are unanimous. However, these ideal conditions are not strictly necessary for market efficiency. Even in a market with significant transaction costs affecting trades, the market can still be efficient if enough participants are informed. Variations in how individuals assess the information's effect on prices do not necessarily undermine market efficiency. What is crucial is the absence of any systematic bias where some investors consistently outperform others in predicting the future prices based on the available information. This perspective maintains that markets can be efficient despite obvious practical imperfections and varied interpretations of data.

In reality, markets are influenced by transaction costs, taxes, and the resources required to acquire new information, knowing that time spent gathering data could be allocated elsewhere. Financial theory recognises these realities and concedes that market efficiency that market efficiency can exist alongside these imperfections. It does not necessitate that stock prices always precisely reflect the underlying value of assets. Instead, market efficiency is characterised by the notion that any discrepancies between the pricing of the market and the intrinsic value occur randomly. This randomness implies that stocks are equally likely to be priced above or below their value at any given moment, with these variations being inherently unpredictable. Consequently, no pattern should exist that allows for consistent prediction of overvalued or undervalued stocks, nor should it be possible to systematically exploit market anomalies for profit. This concept asserts that no investment strategy can reliably identify and capitalise in mispriced securities over time. (Knüpfer & Puttonen 2018.)

Market efficiency implies that stock returns follow a random walk, meaning that the returns from previous days have no bearing on future returns in an efficient market. This lack of correlation is due to the fact that stock prices only change in response to new information. Under the EMH model, when new information is made available, a stock's price adjusts instantly and appropriately in an efficiently functioning market. In contrast, in less efficient markets, investors might delay or overreact to new information, leading to a gradual correction of the stock price over time. Knüpfer & Puttonen (2018) illustrate this concept with two scenarios depicting stock price movements upon the release of new information.

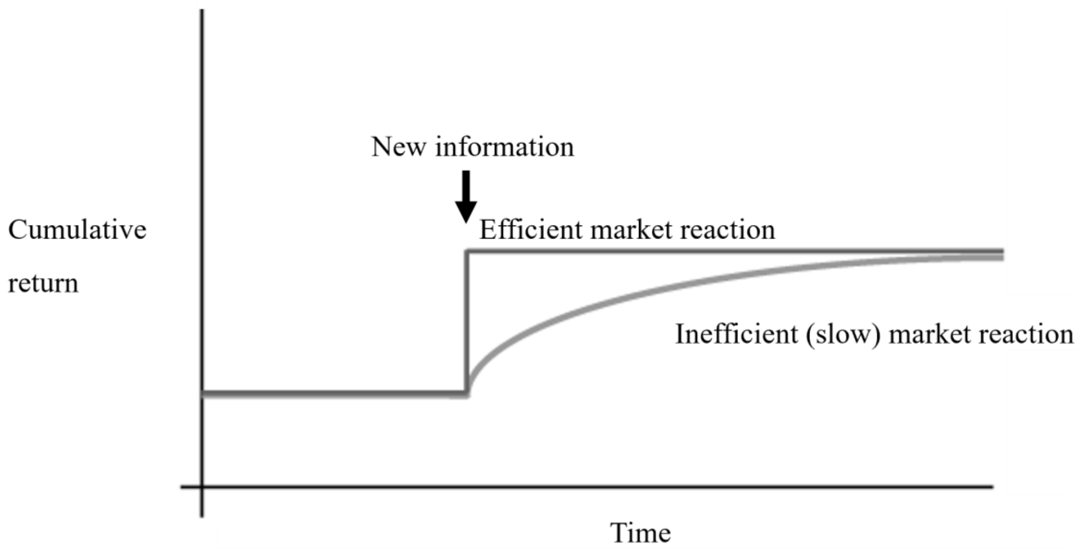

Figure 1 Stock price reaction to a release of new information (Knüpfer & Puttonen 2018)

In Figure 1, the efficient market is represented by a black line, where price adjustment is immediate and accurate, while the lighter line denotes the reaction of an inefficient market, characterised by a sluggish response to new information.

Market efficiency, as defined in the EMH, is categorised into three levels: weak, semi-strong, and strong. These levels can be distinguished by how much information is reflected in the prices of securities. In the weakest form of efficiency, stock prices incorporate all past trading information (prices), meaning that historical data cannot be used to secure superior returns. Semi-strong form efficiency asserts that the prices of stocks reflect all publicly available information on top of the historical data, negating any advantage from trading on public news or data (such as annual reports). Finally, strong form efficiency goes a step further by suggesting that stock prices factor in all information, public and private (including insider information), making it practically impossible for any investor to achieve consistently higher returns than the general market. Essentially, as the market efficiency scales from weak to strong, the possibility of earning abnormal returns by exploiting specific information diminishes, with strong form efficiency representing a state where even inside information is immediately priced into market values. (Fama 1970, 388.)

The first studies around the EMH focused around testing the weak form efficiency of markets and therefore the past prices of stocks. Historical data is practically costless to obtain from various sources (e.g. *Yahoo Finance*) by investigating market-trading data (Fama 1970, 388). Therefore, achieving superior risk-adjusted returns should not be feasible based solely on prior knowledge of historical prices and returns (Shleifer 2006, 6). The majority of findings align with the random walk literature, which supports the weak form of the EMH. Consequently, these findings correspond with the notion that returns cannot be predicted solely on the basis of past returns. (Fama, 1970.)

As many studies supported the weak form of the EMH, studies shifted focus to the semi-strong form. This form states that stock prices should instantly incorporate all historical and publicly available information (Fama 1970). Essentially, the semi-strong form examines how swiftly prices adapt to new public information. Under this hypothesis, if markets are truly efficient, investors would be unable to achieve higher risk-adjusted returns because any public information would already be priced in (Shleifer 2000, 6).

However, the possibility of earning excess returns is not entirely dismissed; such profits might arise from access to non-public or insider information. The strong-form efficiency hypothesis expands on this, suggesting that all information, both public and private, potentially available in the future is already priced in (Shleifer 2000, 6). It particularly scrutinises whether any individuals or groups, such as fund managers, possess exclusive information that can influence price formation (Fama 1970, 388). While the strong form suggests a theoretical possibility of profit from insider information, in practice, it is challenging to capitalise on such data as markets tend to assimilate any leaked insider news rapidly and accurately into security prices (Shleifer 2000, 6).

Critique of the EMH often highlights misunderstandings regarding its implications for market behaviour, specifically in three areas: stock price deviations, the ability to outperform the market, and overall market rationality. (Clarke et al. 2001, 7–11; Knüpfer & Puttonen 2018.)

Firstly, critics note that stock prices do not constantly reflect their intrinsic values. EMH acknowledges this, suggesting that while prices may deviate from actual values, these deviations are random and unpredictable. Thus, market efficiency does not require that stock prices always mirror intrinsic value; instead, it contends that price variations are not systematically exploitable. (Knüpfer & Puttonen 2018.)

Secondly, EMH does not claim that no investor has the ability to beat the market; rather, it implies that consistent outperformance is more likely due to chance than skill (Clarke et al. 2001, 8). While investors can profit from new information that causes stock prises to rise, the hypothesis argues that over time, successful investing is not generally attributable to a particular strategy but to luck. In efficient markets, investment returns correspond to their risk level over the long term, but short-term disparities between actual and expected returns are common. (Knüpfer & Puttonen 2018.)

Lastly, EMH assumes that the market operates rationally, not that every individual market participant acts rationally. Individual irrational decisions do not necessarily indicate market inefficiency. The theory assumes that the collective actions of all market participants, including rational traders capitalising on the misjudgements of others, drive prices to reflect intrinsic value, thereby upholding market efficiency even in the presence of irrational behaviour. (Knüpfer & Puttonen 2018.)

2.2 Random walk

The random walk hypothesis is a theory that states that market prices fluctuate in an unpredictable manner, moving up and down without being influenced by previous price trends (Knüpfer & Puttonen 2018, 169). This randomness makes it difficult to accurately predict the future direction of the market at any given moment. This concept implies that a financial market's past behaviour does not provide any useful clues for forecasting future security prices. Essentially, the history of a stock's price offers no advantage in predicting its future trajectory, suggesting that investors cannot consistently predict security prices with more than a 50 per cent success rate. The random walk theory aligns with the principles of the EMH, operating under the premise that financial markets function efficiently and that all known information is already reflected in stock prices (Fama 1965b, 76).

Numerous empirical studies have investigated the random walk theory and the ability to predict stock prices. For instance, Odean (1998) analysed the performance of investors discount brokerage accounts, comparing their trading profits against their trading costs. He found that these investors generally performed worse than those who adopted a simple buy-and-hold strategy, frequently buying securities that underperformed compared to the ones they sold, not even managing to cover their trading costs. Similarly, Chitenderu et al. (2014) studied the Johannesburg Stock Exchange All Share Index between 2000 and 2011 using monthly data. Their research concluded that the stock prices in this market were not correlated and followed a random walk pattern, supporting the hypothesis that future stock prices in this market are unpredictable and move independently. An example of what a random walk would look like for stock returns is shown in Figure 2.

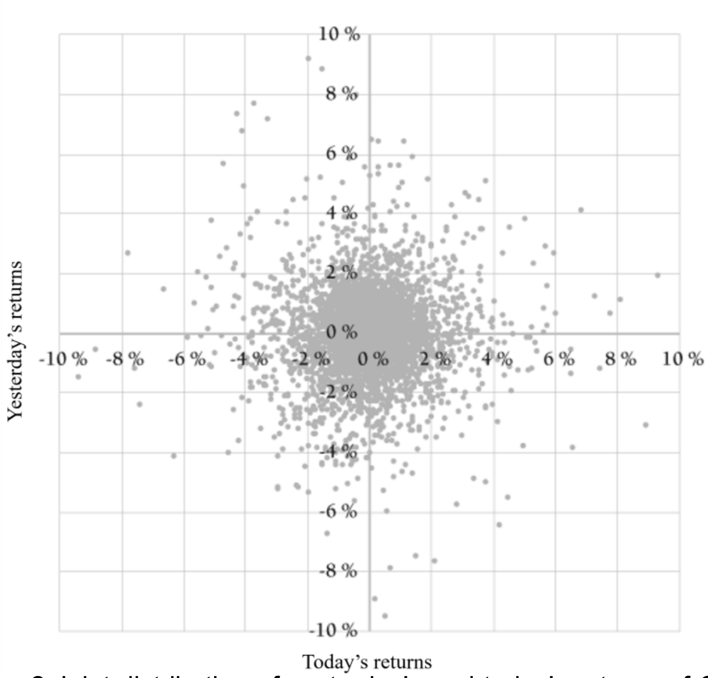

Figure 2 Joint distribution of yesterday's and today's returns of OMX Helsinki Cap -index 2002-2017 (Knüpfer & Puttonen 2018)

Each dot shows the return from a certain day and the day after it. Based on the figure it seems to be impossible to predict tomorrow's return based on what has happened today. If this was the case, the dots would be grouped in a way that could be interpreted as a line (Knüpfer & Puttonen 2018).

Conversely, Lo and MacKinlay (1988) presented findings that challenged the notion of random walk in stock prices. Analysing data from the United States stock market between 1962 and 1985, including various indices and size-sorted portfolios, they tested the random walk hypothesis specifically on weekly returns. Their research strongly refuted the random walk model across the whole sample period and all subperiods examined. They highlighted evidence suggesting that stock returns might be somewhat predictable, contradicting earlier studies that upheld the random walk theory.

Fama and French (1988) reported findings in line with those of Lo and MacKinlay (1988), noting that while autocorrelation is weak for short-term daily and weekly periods, it becomes more pronounced over longer durations. This suggests that there are times when stock prices deviate from the random walk model, particularly over the long term, implying a degree of predictability in stock prices. Kim et al. (2002) observed that stock price movements are not entirely random. Their research identified significant statistical correlations between the prices of certain stocks, suggesting that in some instances, the movement of one stock could predict the movement of another.

3 Mean reversion

3.1 Definitions and types of mean reversion

The concept of mean reversion means that asset prices will ultimately gravitate back to their historical averages or a relative measure like the MSCI World Index. This historical average can be, for example, a simple price-to-earnings multiple. Assets trading below their long-term average are typically viewed as buys, while those above are expected to decrease. The further a price strays from its mean, the more likely it is that it will revert, creating opportunities for investors and traders to plan their market entry and exit strategies. Mean reversion is a principle that can be utilised across various financial metrics like price, earnings, and ratios. (Chen 2023.)

A straightforward mathematical approach to mean reversion could be presented as a simple autoregressive process of order one with drift:

$$x_t = \theta_0 + \theta_1 x_{t-1} + \varepsilon_t,$$

where $\varepsilon_t \sim (0, \sigma)$ and $\theta \in (0,1)$. The unconditional mean of the process is

$$\mathbb{E}(x) = \frac{\theta_0}{1-\theta_1},$$

where the persistence parameter, θ_1, determines how quickly the process returns to this mean. Essentially, a shock ε_{t-1} impacts the variable x_t proportionally to θ_1; its effect diminishes over time, influencing x_{t+1} to a degree of θ_1^2, and so on. In other words, a portion θ_1 of the initial shock continues onward with each time unit, and conversely, a proportion $[1-\theta_1 \in (0,1)]$ of the shock dissipates each time unit. The inverse $(\frac{1}{1-\theta_1})$ represents the average duration it takes for a shock to dissipate fully, known as the mean reversion time. The lower the persistence parameter θ_1 is, the faster the process reverts to its mean. It should be noted that this process only models positive autocorrelations.

In economic literature, two distinct types of mean reversion are identified: absolute and relative. Absolute mean reversion suggests that stock prices tend to revert to a constant, unspecified mean value over time, typically evidenced by negative autocorrelation in stock market returns. On the other hand, relative mean reversion indicates a more dynamic relationship where stock prices revert to a mean that is directly related to fundamental indicators such as dividends and earnings, reflecting the intrinsic value of the stocks. Both concepts are crucial in understanding the patterns and long-term behaviours of stock prices, with absolute mean reversion focusing on the historical average and relative mean reversion emphasising the underlying economic fundamentals. (Spierdijk et al. 2012, 230; Arsalan et al., 2022.)

3.2 Rational of mean reversion

The notion of mean reversion conflicts with the EMH. For example, sudden positive news may cause the price of a stock to jump 30 per cent instead of 10 per cent causing it to go over its fundamental value, a value which it would have in an efficient market. After this, traders who pay attention to the fundamentals understand that the stock's price is over its fundamental value and start selling it, causing downward pressure on the price, and therefore causing the average price to fall in the following periods. Eventually, the value of the stock goes back to its fundamental value. (Engel & Morris 1991, 24–25.)

Due to this conflict, many explain mean reversion through human psychology. The Overreaction Hypothesis developed by De Bondt and Thaler (1985) could be used to explain the temporal dependence of why the return process is driven back to its mean at some time scale. The hypothesis suggests that investors typically respond sluggishly to favourable news about a stock, incorporating the information into prices more slowly than they should, which leads to continued positive returns. This underreaction is believed to cause positive autocorrelation in returns over certain periods. On the other hand, overreaction occurs when investors get carried away by a series of positive developments, assuming the trend will continue and thus pushing prices beyond what is justified. When negative news finally emerges, it causes a sharp downturn in price. This overreaction is often associated with a negative autocorrelation in returns.

Griffin and Tversky (1992) propose that both underreaction and overreaction can occur simultaneously and serve to explain market behaviours. They differentiate between the "strength" and "weight" of signals. For instance, a series of positive earnings announcements is noticeable and widely discussed (strong) but might not be as informative as it seems (light) because such sequences can occur randomly. Conversely, a single positive earnings announcement might not attract much attention (weak) but could be extremely informative (heavy). Investors tend to overemphasise the strength of signals while overlooking their weight, leading to underreaction to isolated pieces of good news and overreaction to continuous good news. This tendency results in positive autocorrelation of returns over shorter timescales (between one month and one year) and negative over longer timescales (between three and four years).

3.3 Evidence and measurement of mean reversal

The debate around mean reversion in stock returns was ignited by a study conducted by De Bondt and Thaler (1985), which examined stocks on NYSE (New York Stock Exchange) from early 1926 to late 1982. They classified stock into “winner” an “loser” portfolios based on their performance across sixteen separate three-year intervals from January 1930 to January 1975, using monthly returns as a measure. Stocks with top performance (top 35, top 50, or top decile) were assigned to the “winner” portfolio, while those with the lowest performance (bottom 35, bottom 50, or bottom decile) were placed in the “loser” portfolio. The study aimed to determine whether the stocks past performance could predict their future returns, specifically looking at the cumulative excess returns (returns over the average) for the subsequent years.

What De Bondt and Thaler (1985) discovered was counter to what market efficiency would predict: stocks that were previously “winners” significantly underperformed in the following years, whereas the “loser” stocks tended to yield much higher returns. This effect was more pronounced for “losers”, suggesting an asymmetry in the mean reversion effect. These findings were consistent regardless of the specific criteria for “winner” and “loser” categorization or the length of the periods examined. This meant that the past performance of a stock could be utilised in predicting its future performance. Further evidence of mean reversion can be found from the likes of Poterba and Summers (1988), Fama and French (1988), and Jegadeesh (1991).

Poterba and Summers (1988) used a variance ratio test derived from the model made by Summers (1986) to measure mean reversion. The variance ratio for monthly returns r_t was defined as

$$Variance\ ratio(k) = \frac{Var(r_{t,t+k})}{Var(r_{t,t+12})} \times \frac{12}{k},$$

where k-period returns are compared to the variance of annual returns. If the stock price were to follow a random walk, then the returns would act as white noise, meaning their variance increases linearly with k. In such a scenario, the variance of the returns over a

12-period horizon would be $Var(r_{t,t+12}) = \sigma_r^2 \times 12$, and for any k-period, it would be $Var(r_{t,t+12}) = \sigma_r^2 \times k$, making the variance ratio $VR(k) = 1$ for all k periods. However, if the returns are mean reverting, the variance of returns over longer periods will not increase as quickly as it would under a white noise process. As a result, the variance ratio for these longer periods will drop below one, indicating the presence of mean reversion rather than a random walk in stock price returns.

As an example, to illustrate (Figure 3) the difference in behaviour between a mean reversing stock and a stock in efficient markets: consider a stock with potential annual changes of a 20 per cent increase or a 10 per cent decrease. One way to examine this stock's volatility is by the range of its possible returns over a set time. For a one-year period, the maximum return is 20 per cent, and the minimum is a 10 per cent loss, making the volatility 30 per cent. In a two-year span, the best scenario yields a 40 per cent return (20 per cent each year), and the worst case is a 20 per cent loss, resulting in a volatility of 60 per cent[1]. Therefore, for an efficient market, the volatility of a two-year investment is double that of a one-year investment. (Engel & Morris 1991, 26–27.)

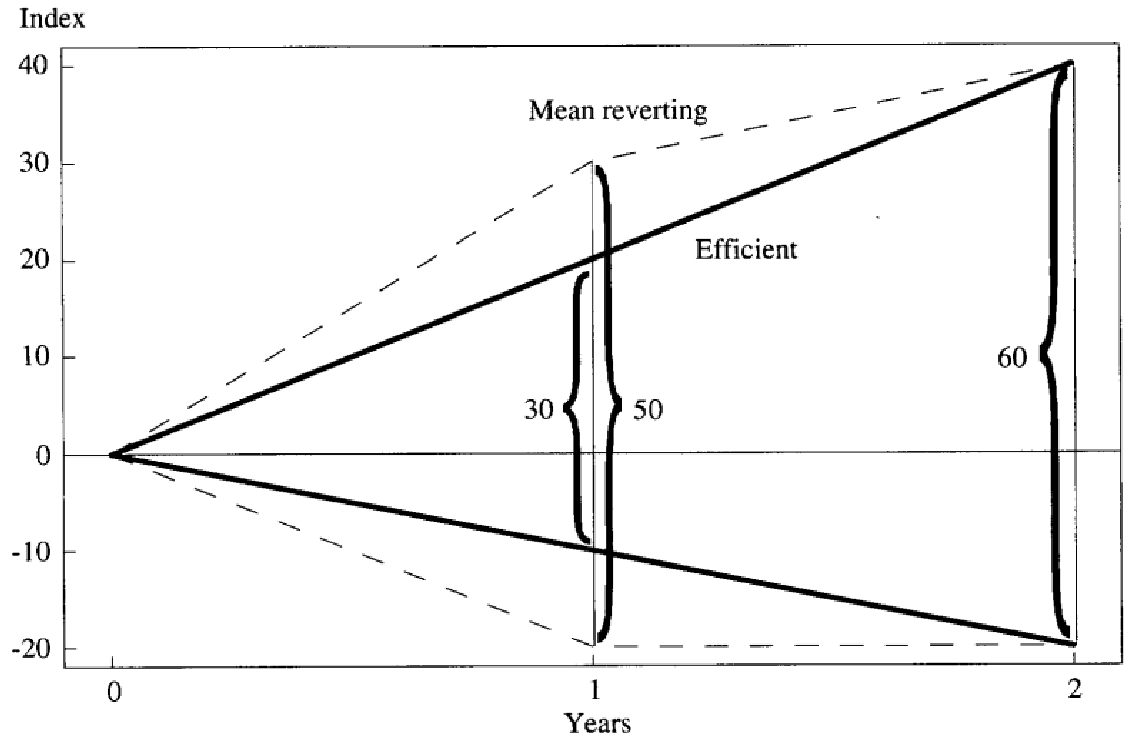

Figure 3 Volatility of a mean reversing stock and a stock in efficient market (Engel & Morris 1991, 27)

[1] When not taking into account for interest on interest -effect.

However, if the stock prices are mean reverting, the volatility of a longer-term investment will not scale linearly (Figure 3, dashed line). In the short term, prices might deviate significantly from the fundamental value (Figure 3, solid line) – say, fluctuating between a 30 per cent increase and a 20 per cent decrease in the first year, giving a 50 per cent volatility. But if the price reverts to its fundamental value in the following period, the volatility over two years would align with the efficient market scenario at 60 per cent, indicating that the long-term volatility is significantly less than double the one-year volatility. This reduced scaling reflects the dampening effect of mean reversion on price fluctuations over extended periods. (Engel & Morris 1991, 26–27.)

Further methods for detecting mean reversion utilize the Hurst exponent and the half-life of stock returns. Both methods can be used to evaluate the magnitude of mean reversion. Serletis and Rosenberg (2007, 2009) used the Hurst exponent to measure the mean reversion in energy futures prices and the United States stock market. They discovered that both display mean reversing behaviour.

The half life is a commonly used metric for assessing mean reversion. It represents the duration needed for a time series' reaction to a single shock to reduce by one half (Kim & Ji 2011, 1959). Chaudhuri and Wu (2003) found mean reversion in emerging market stock indices while Kim and Ji (2011) used half-life to study mean reversion in global real interest rates. These studies show how mean reversion can be detected with multiple methods, all of which will be used in this thesis. Further explanation to the methods (variance ratio, Hurst exponent, and half-life) will be given in Chapter 5.

3.4 Mean in mean reversion

The absence of proof supporting mean reversion is frequently linked to using too small a sample size and the use of statistical tests that are inherently not powerful enough to capture the phenomenon. By clearly defining the fundamental value process to which stock prices revert, a significant enhancement in the accuracy of estimations can be realised. A critical issue that arises when looking at relative mean reversion is the method of approximating the fundamental value process, which is by nature not directly observable. (Spierdijk et al. 2012, 231.)

Proxies for the fundamental value could be dividends, earnings, or valuation metrics, enterprise value or price-to-earnings (P/E) ratios (Spierdijk et al 2012, 231). From a theoretical standpoint, these ratios should exhibit mean reverting behaviour, given that fundamentals drive stock prices. For example, should the stock price be elevated relative to the fundamentals of a company, an adjustment in either the stock's price or its fundamentals are expected to occur. Campbell and Shiller (2005) investigated the tendency of the dividend yield and P/E-ratio to revert to their means over time. They found that more often than not, the price of a stock was the key contributor to the fundamental ratio reverting towards its mean.

4 Stylized features of financial time series

Research into the statistical characteristics of financial time series has uncovered numerous intriguing and standardized observations that appear to be consistently present across various markets, instruments, and time frames. Cont (2001) gives a comprehensive list of stylized facts for financial time series and this thesis will present some of them in more detail:

- *Excess volatility*: Cutler et al. (1989) showed that the occurrence of significant positive or negative returns often cannot be fully accounted for by the emergence of new market information. Empirical investigations have highlighted the challenges in explaining the observed volatility of asset returns solely based on changes in fundamental economic variables.

- *Heavy tails*: the distribution of returns, without any given conditions, shows a pronounced heavy tail, indicating a higher peak and fatter tails than the normal distribution, known as positive excess kurtosis (Ghose & Kroner 1995; Mandelbrot 1963).

- *Absence of autocorrelations*: asset returns do not show significant level of autocorrelation besides on short intraday time scales (Cont 2001).

- *Volatility clustering*: Mandelbrot (1963) noted that significant market movements are typically succeeded by other significant movements, which can be either upward or downward, while minor movements are likely to be followed by similarly minor movements. This phenomenon is reflected in the behaviour of absolute returns $|r_t|$ or their squares r_t^2. Although the returns might be uncorrelated (see Chapter 4.2), the squared and absolute returns exhibit a positive, significant, and slowly diminishing autocorrelation (Ding et al. 1993, 87).

- *Volume and volatility correlation*: trading volume and market volatility are positively correlated (Cont 2001). Additionally, both volatility and trading volume exhibit a similar pattern of long memory, indicating that past values can influence future values over extended periods (Lobato & Velasco 2000).

4.1 Excess volatility

The findings of Cutler et al. (1989) suggest a significant challenge in attributing even half of the aggregate stock price volatility to public news related to fundamental values, echoing Roll's (1984) conclusion that new information cannot fully account for the variations in individual stock or futures returns. While it is plausible that there is some form of crucial news being overlooked that would significantly impact asset prices volatility, scepticism remains. This scepticism is due to the belief that news impactful enough to cause major shifts in stock demand should be detectable in economic indicators or market reports.

Moreover, the difficulty in correlating price changes with fundamental values extends beyond the general stock market. Research in areas with directly measurable fundamental values also struggles to explain price behaviour. Most notably with stock return anomalies, like those occurring around holidays and various calendar periods documented by Thaler (1987a, 1978b) are similarly challenging to link to news about fundamentals, as these values typically do not exhibit systematic changes during such times. This persistent issue highlights the complexity and potential gaps in understanding the factors that drive asset prices.

Excess volatility can be witnessed in the bond market as well. Bao and Pan (2013) used a Merton (1974) model with the primary inputs being stochastic interest rates and the volatility of equities and found that the Merton model gave significantly lower bond volatilities compared to their actual volatilities, indicating that there is a notably higher level of excess volatility in the corporate bond market. They proposed two potential explanations for this excess volatility. The first is that the excess volatility might be due to the Merton model's inability to accurately capture the underlying fundamentals that link equity and credit markets. The second possibility they suggest is that the volatility could be a result of fluctuating levels of illiquidity in the credit markets. (Bao & Pan, 2013, 3095.)

4.2 Heavy tails

Mandelbrot (1963) showed the insufficient nature of the Gaussian distribution in modelling the marginal distribution of returns for assets, particularly noting their heavy tails. This deviation from normality is prominently observed in two statistical phenomena:

extreme events happen more frequently than what should be expected in a normal distribution (Mandelbrot 1963; Fama 1963; Kon 1984), and significant market downswings tend to occur more often than meaningful market upswings (Fama 1965a; Arditti 1971; Singleton & Wingender 1986). Following this, the non-Gaussian properties of price change distributions have been consistently observed across multiple studies (see Fama 1965a; Guillaume et al. 1997; Gopikrishnan et al. 2000).

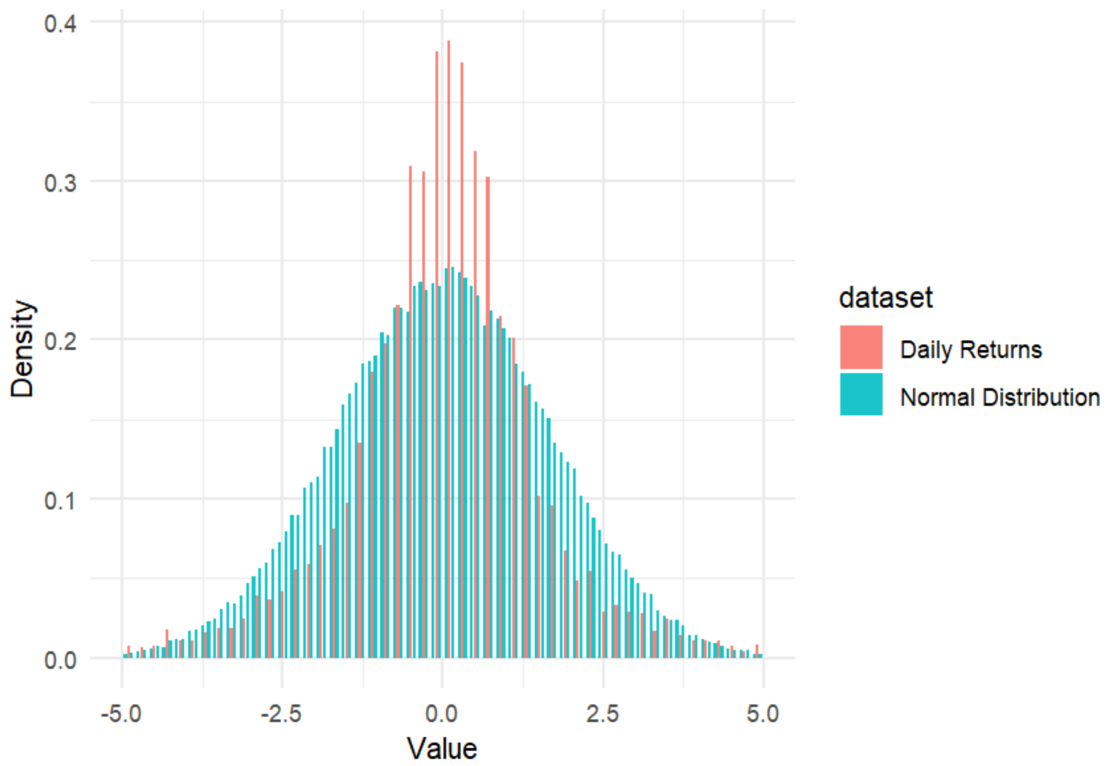

Figure 4 Density of daily natural logarithmic of the OMXH25-index and a normal distribution

The non-gaussian property of asset returns is depicted in Figure 4. The red histogram shows the density of the daily logarithmic returns of the OMXH25-index, and the green histogram is what the normal distribution, with an equal mean and variance to the sample, would be. The histogram of the OMXHPI-index includes 6273 observations, while the normal distribution was simulated using 100 000 observations. The distribution of asset returns appears to be sharp-edged with a higher peak and heavy tailed, similar to the findings of Ghose and Kroner (1995) and Mandelbrot (1963). These properties become more pronounced when the frequency of data increases, meaning that the distribution gets further from a normal distribution as the frequency increases, leading to a phenomenon Cont describes as "aggregated Gaussianity". Because of these features, the distribution of returns cannot be fully determined using one single model. (Cont 2001, 226.)

Consequently, to bridge this gap, there has been numerous parametric models that have been presented in economic literature besides normal distribution. Only naming some of the models: exponentially truncated stable distribution (Bouchaud and Potters 1997), stable distribution (Mandelbrot 1963), hyperbolic distribution (Prause 1998), Student distribution (Kon 1984), and normal inverse Gaussian distribution (Barndorff-Nielsen 1997).

4.3 Absence of autocorrelation

In liquid markets, price movements do not show substantial autocorrelation. The autocorrelation function of price changes quickly diminishes to zero in a matter of minutes. For time periods longer than 15 minutes, the autocorrelation is effectively negligible and can be considered zero. (Cont 2001, 229.)

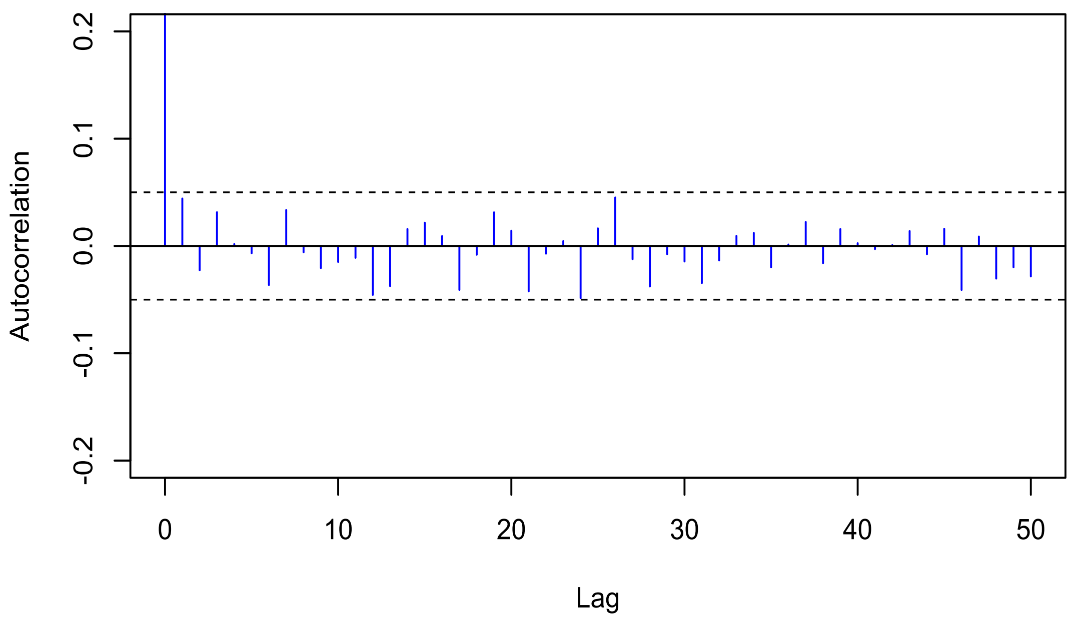

Figure 5 OMXH25 total return index daily logarithm returns autocorrelation function

In Figure 4 is shown the autocorrelation function of daily logarithmic returns from OMXH25 stock index. The black dotted line depicts the five per cent significance level. No lag, besides 0, goes over this level, meaning that there is no statistically significant autocorrelation up to lag 50 which is consistent with the remarks made by Cont (2001, 224). Additionally, the ACF shows that there is no trend or seasonality for 50-day periods in the chosen index. One explanation to the lack of correlation can be explained by arbitrage. If price changes were significantly correlated, traders could exploit this relationship through simple strategies aimed at gaining positive expected returns. These strategies would naturally work to diminish any correlations, except over very brief periods representing the market's response time to new information. This reaction time is generally just a few minutes in organised options markets and even shorter in foreign exchange markets. (Cont 2001, 229.)

However, this absence of autocorrelation does not consistently apply to longer time scales. When examining weekly or monthly returns, some degree of autocorrelation does appear. But, as the scope of the time scale increases, the amount of available data decreases, making the statistical evidence less definitive and more varied across different data sets. (Cont 2001, 230.)

4.4 Volatility clustering

A well recognised characteristic of financial asset returns is the phenomenon known as volatility clustering. Essentially, periods of high volatility, marked by large swings in returns, are often followed by similar periods of intense fluctuation. Conversely, times of low volatility, where returns change minimally, typically precede calmer market conditions. This tendency leads to observable patterns of clustered high and low volatility. Understanding and modelling these volatility clusters is crucial in the financial market, as asset return volatility significantly affects option pricing and the risk levels of stocks and portfolios. (Ning et al. 2015, 62.)

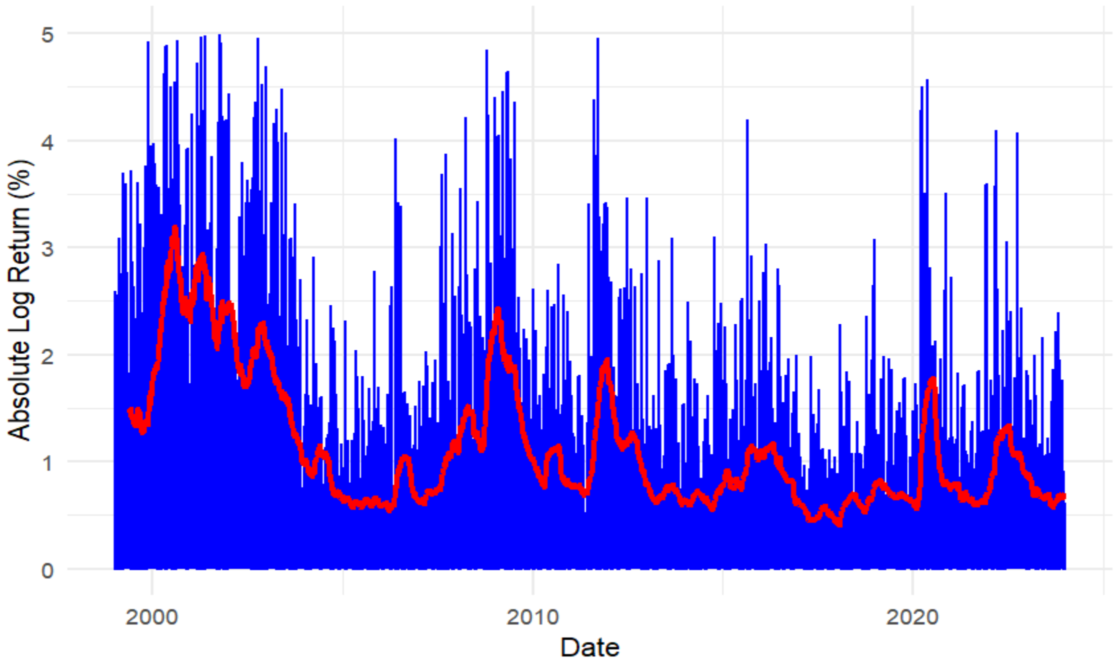

Figure 6 Daily logarithmic returns of the OMXH25 and rolling 100-day mean return in absolute values

Evidence of volatility clustering can be seen in Figure 6 where there has been plotted the daily logarithmic changes in the OMXH25-index from early 1999 to the end of 2023. The blue bars show daily logarithmic returns while the red line showcases the 100-day rolling mean return which is achieved by calculating the average of the previous 100 data points starting from the 100th index and repeating this for each subsequent index. Times of higher volatility (larger changes in price) tend to be close together and the same can be said about the times when volatility is lower. Notable periods of high volatility are around 2001–2002 and 2008. This visual pattern is distinctive to volatility clustering. Ning et al. (2015) discovered that in series of high returns, clusters of higher volatility occur more frequently than those of lower volatility, indicating an asymmetry in how volatility aggregates. This pattern was consistently observed over various time frames before, during, and after the financial crisis of 2008. Additionally, Ning et al. (2015) noted that these clusters exhibit strong persistence and longevity, remaining evident for extended periods of over a month and longer.

There have been multiple studies on what the reasons behind volatility clustering are. Ghashghaie et al. (1996, 769) argued that interruptions in the trading process, especially longer ones for holidays or weekends, were one reason behind the clustering and changes in the shape of probability densities. Clark (1973) reasons that irregularly arriving information causes volatility clustering and non-gaussianity.

5 Methodology

The research methodology of this thesis is to collect historical data from a stock index, calculate the daily logarithmic returns and study these returns for possible mean reverting behaviour. Following this, the daily value at risk and expected shortfall (also referred to generally as risk measures) values will be calculated in a rolling fashion to examine whether there is visible correlation between the possible mean reversion and the risk measures. For a more analytical approach, this thesis will conduct several parametric models. The data manipulation and analysis will be conducted using the programming language R.

5.1 Data

The data that will be used in this thesis is the Finnish OMXH price index (OMXHPI), which includes the 139 largest companies listed in Finland by market capitalisation as of March 1st 2024. The data was downloaded online from Nasdaq OMX Nordic. The total amount of observations amounts to 6273, from January 1999 to December 2023, further descriptive statistics can be seen from Table 1. Lo and MacKinlay (1988) argued that the volatility of inflation is insignificant compared to the volatility of the return of stocks and thus can be disregarded when doing the variance ratio tests. Therefore, the price index is not adjusted for inflation.

Table 1 Descriptive statistics

Descriptive Statistics	Value
Minimum	-17.42461%
1st Quartile	-1.05717%
Median	0.09551%
Mean	0.02575%
3rd Quartile	1.06907%
Maximum	14.56310%
Skewness	-0.42417
Kurtosis	11.91586
Standard Deviation	1.66506%
Observations	6273

Based on the histogram (Figure 4) of the daily returns, the daily returns seem to follow a normal distribution. However, based on the Shapiro-Wilk test conducted in R, the null hypothesis of normality must be rejected. This finding is consistent with that of Fama's (1965a, 89) that daily returns are not normally distributed.

5.2 Variance ratio

The Lo and MacKinley (1988) (LMK) variance ratio statistic is based on the characteristic that the variance of random walk's (X_t) changes increases proportionally with the length of the time interval (Figure 3, solid line). This implies that the variance of the difference of X_t and X_{t-q} is equal to q times the variance of the difference between X_t and X_{t-1}. In other words, to test the random walk hypothesis, one could calculate whether

$$\frac{1}{q} * Var\left(X_t - X_{t-q}\right) = Var(X_t - X_{t-1})$$

holds true. Letting P_t represent the rate of exchange at time t and X_t be the natural logarithm of $P_t (X_t = lnP_t)$, the variance ratio $VR(q)$ can be calculated as

$$VR(q) = \frac{\sigma^2(q)}{\sigma^2(1)},$$

where

$$\sigma^2(q) = \frac{1}{m}\sum_{t=q}^{nq}(X_t - X_{t-q} - q\bar{X})^2,$$

where

$$m = q(nq - q + 1)\left(1 - \frac{q}{nq}\right),$$

and

$$\sigma^2(1) = \frac{1}{nq-1}\sum_{t=1}^{nq}(X_t - X_{t-1} - \bar{X})^2.$$

$\bar{X}$ is the mean of $(X_t - X_{t-1})$ and X_{nq} is the final observation in the time series. The assumption, or null hypothesis, posits that *VR(q)* is not significantly different from one.

Poterba and Summers (1988) compared the two methods of LMK (1988) and Fama and French (1988) and concluded that the variance ratio test used by LMK was more powerful when it comes to detecting mean reversion. Chow and Denning (1993, 389) argue that the LMK method is suitable for evaluating a variance ratio tied to a particular level, q, by comparing the test statistic, $Z_1(q)$ or $Z_2(q)$, against the standard critical value. Yet, as the random walk hypothesis necessitates that variance ratios for all chosen aggregation intervals, q, should equal one, an effective approach to test this hypothesis is by conducting a multiple comparison of all the chosen variance ratio estimates against the number one.

Chow and Denning (1993) modified LMK's methodology into a method they called the "multiple comparison test" to test multiple variance ratios at once. Examining a collection of variance ratio estimates, denoted as $\{VR(q_i) \mid i = 1,2, \dots, L\}$ which align with a predetermined series of lags $\{q_i \mid i = 1,2, \dots, L\}$. Under the assumption of a random walk, we examine a series of smaller hypotheses, $H_{0i}: VR(q_i) = 1$, for each $i = 1,2, \dots, L$. As any single disproof of H_{0i} implies the failure of the random walk hypothesis, we consider the most extreme (largest value in absolute terms) of the test statistics as significant. The values of the test statistics are

$$Z_1(q) = \max_{1 \le i \le L} |Z(q_i)|,$$

and

$$Z_2(q) = \max_{1 \le i \le L} |Z^*(q_i)|,$$

where $Z(q_i)$ and $Z^*(q_i)$ are calculated as

$$Z(q) = \frac{(VR(q) - 1)}{\sqrt{\theta(q)}} \sim N(0,1),$$

where

$$\theta(q) = \frac{2(2q-1)(q-1)}{3q(nq)},$$

and

$$Z^*(q) = \frac{(VR(q) - 1)}{\sqrt{\theta^*(q)}} \sim N(0,1),$$

where

$$\theta^*(q) = \sum_{j=1}^{q-1} \left[\frac{2(q-j)}{q}\right]^2 \delta(j),$$

where

$$\delta(j) = \frac{\sum_{t=j+1}^{nq} (X_t - X_{t-1} - \bar{X})^2 (X_t - X_{t-j-1} - \bar{X})^2}{\left[\sum_{t=1}^{nq} (X_t - X_{t-1} - \bar{X})^2\right]^2}.$$

Rejecting the null hypothesis relies on the maximum absolute value among the individual variance ratio test statistics. This test statistic follows the studentised maximum modulus (SMM) distribution, characterised by L and T (the sample size) degrees of freedom, with critical values accessible from Stoline and Ury (1979). In instances where the sample size T is considerable, the null hypothesis is rejected at the α significance level if $Z_1(q)$ [or $Z_2(q)$] exceeds the $\left[1 - \frac{\alpha^*}{2}\right]$th percentile of a standard Gaussian distribution, where $\alpha^* = 1 - (1-\alpha)^{\frac{1}{L}}$. Both $Z_1(q)$ and $Z_2(q)$ adhere to identical critical values. For example, for a large T and at significance levels of 10%, 5%, and 1%, the SMM critical values for $L = 4$ are 2.23, 2.49, and 3.03, respectively.

This thesis will use the multiple comparison test of Chow and Denning (1993) instead of Lo and MacKinlay's (1988) individual variance ratio test. One reason for this is because the multiple comparison test notably lowers the probability of a Type I error (false positive) (Charles & Darné 2009). Although this thesis will not be using or going over them, additional versions of the variance ratio test can be seen from Malliaropulos and Priestley (1999), Wright (2000), and Kim (2006).

5.3 Hurst exponent

The Hurst exponent (HE) is formed from the detrending moving average (DMA) and detrended fluctuation analysis (DFA). DFA was first introduced by Peng et al. (1994) and involves segmenting the series $y(t)$ into non-overlapping segments of uniform length n. Within each segment of size n, a least squares line $y_{nm}(t)$ is fitted to represent a local linear trend. The series $y(t)$ is then detrended by removing this linear trend, $y_{nm}(t)$. Following this, the fluctuation of the root mean square of this detrended series is computed. This process is repeated across various segment sizes to distinguish the relationship between the time scale n and the average detrended fluctuation. For the mathematical representation of the process see Weron (2002, 288).

The DMA process represents an advancement over the DFA because it does not necessitate segmenting the time series $y(t)$ into distinct, non-overlapping sections. Instead, DMA detrends the time series by subtracting a moving average, a continuous function, from the series. This method is shown to be more precise, as the moving average (MA) functions as a more effective low-pass filter when comparing it to the polynomial filter employed in DFA. To explain the DMA process, let us consider a time series $g(t)$ where t ranges from 1 to N. The nth order MA of $g(t)$ is represented as

$$\bar{g}_n(t) = \frac{1}{n}\sum_{k=0}^{n-1} y(t-k).$$

In this process, the series $g(t)$ is detrended by removing the MA, $\bar{g}_n(t)$. The standard deviation of $g(t)$ relative to its MA, $\bar{g}_n(t)$, is then determined with

$$\sigma_{DMA} = \sqrt{\frac{1}{N-n_{max}}\sum_{t=n_{max}}^{N} [g(t)-\bar{g}_n(t)]^2},$$

where n_{max} is the maximum value for n. (Serletis & Rosenberg 2007, 326.)

The Hurst exponent is determined by plotting σ_{DMA} against n on a log-log scale and calculating the slope of the resulting line. The value of the slope, or Hurst exponent, is indicative of the type and strength of correlation in the time series and is explained in Table 2.

Table 2 Explanations to Hurst Exponent values (Serletis & Rosenberg 2007, 327)

Hurst Exponent (H) Value	**Correlation Type**	**Description**
0 < H < 0.5	Negative (Anti-persistent, Mean reverting)	Indicates anti-persistence or mean reversion in the data. Trends are more likely to reverse in the future.
H = 0.5	None (Random or Brownian)	Suggests a completely random process, similar to a Brownian motion, with no discernible trend (white noise).
0.5 < H < 1	Positive (Persistent)	Implies persistence. Trends are likely to continue.

5.4 Half-life

A random walk process means that any shock to the stock price is lasting, with no inclination for the price to revert to its historical trend over time (Chaudhuri & Wu 2003, 24). Conversely, half-life, a widely used metric for assessing mean reversion, is characterised by the duration (in periods) needed for the impact of a single shock on a time series to reduce by half (Kim & Ji 2011, 1959). The half-life is frequently used to gauge the mean-reversion characteristics of economic time series, especially when evaluating the legitimacy of purchasing power parity (PPP) conditions in international economics. A prime example is the mean reversion of real FX rates, which is crucial for PPP's validity (Rogoff 1996). Chaudhuri and Wu (2003) used the half-life method to study and discover mean reverting behaviour in emerging market equity prices. Balvers et al. (2000) studied the relative mean reversion in stock price indices from eighteen different countries and found evidence of mean reversion with a half-life of 3.5 years when using annual data.

In the context of the univariate AR(1) model with a slope coefficient α_1, the half-life is determined by the formula

$$\frac{\log(0.5)}{\log(\alpha_1)}.$$

When $\alpha_1 = 0.5$, the AR(1) model is in a stationary state, and the half-life is 1, signifying strong mean reversion. On the other hand, when $\alpha_1 = 1$, the model is non-stationary, and the half-life becomes infinite, reflecting a tendency towards mean aversion. (Kim & Ji, 2011.)

The half-life is derived from an autoregressive (AR) model represented as

$$Y_t = \mu + bt + \alpha_1 Y_{t-1} + \cdots + \alpha_p Y_t + u_t,$$

where the error term $u_t \sim (0, \sigma)$ and is assumed to be i.i.d. However, the method for estimating the half-life can also manage error terms that are (conditionally) heteroskedastic. This AR model can be reformulated as an MA(∞) model with coefficients $\{\beta_i\}_{i=0}^{\infty}$, where $\beta_0 = 0$ and β_i indicates the impulse response of Y_{t+i} to a unit shock in u_t at time t. The graph of $\{\beta_i\}_{i=0}^{\infty}$ against i, for a sufficiently large m, illustrates the impulse response function of Y, showing how the time series reacts to a shock in the error term within a period m. For an AR(p) model with $p > 1$, h is derived from $\{\beta_i\}_{i=0}^{\infty}$. When j falls between $i - 1$ and i, linear interpolation can be used to establish the value of h.

5.5 Value at risk

Before the 1987 market crash, standard deviation was commonly used to gauge investment risk (Tsay 2010, 325). However, the crash shifted focus towards measuring extreme risk or tail risk, garnering substantial interest from investors, industry practitioners, and academics. In response to this need for more comprehensive risk assessment tools, value at risk (VaR) emerged around the end of the 1980s at J.P. Morgan and quickly became a market risk measurement standard (Tsay, 2010, 325; Hull 2018, 269). J.P. Morgan’s leadership sought a single metric to encapsulate the risk across their entire investment portfolio (Miller 2019, 51). The changes to banking regulations in the Basel III Accord require banks to regularly disclose their VaR which further pushed VaR to be the standard risk measurement metric of the financial markets (Patton et al. 2019, 389)

As Tsay (2010, 327) explains, VaR provides an estimate of potential monetary loss at a specific time, quantified to a certain confidence level, and can be depicted as

$$VaR_\alpha = inf\{x | F_t(x) \geq \alpha\},$$

where $1 - \alpha$ represents the confidence level, $F_t(x)$ is the cumulative distribution function over the time horizon *t,* and *inf* is the smallest $x \in \mathbb{R}$ that fulfills $F_t(x) \geq \alpha$.

The specifics of VaR depend on its intended application. Traders, who typically engage in higher turnover with more liquid assets, opt for detailed short-term models that offer higher precision, allowing them to mitigate risk by offloading assets before potential losses materialise. Conversely, investment managers might favour longer time horizons and more general modelling due to their extended duration of risk exposure. The implementation of VaR and other risk metrics varies significantly among different financial professions due to these varying needs and practices. (Alexander 2008, 4, 14.)

VaR has become widely favoured for several reasons, notably its simplicity and the way it consolidates various risk factors into a single, intuitive statistic, making tracking risk over time quite practical. Unlike standard deviation or variance, which consider returns symmetrically around the mean, VaR specifically addresses negative returns, aligning with risk management's focus on downside risk. Its adaptability allows for consolidation of risks across distinct types of securities into one aggregate figure. Additionally, its robustness against extreme data values enables risk managers to concentrate on significant but not excessively rare losses, which some may find advantageous in practical risk assessment. (Miller 2019, 54.)

Yamai and Yoshiba (2005) point out some critical limitations to VaR, first of which is its lack of subadditivity, meaning it does not fully account for the benefits of diversification. Consequently, the estimated potential loss for a portfolio could exceed the sum of the individual VaRs for each asset within the portfolio. Another critical issue is VaR's disregard for negative returns beyond its set confidence level, where the most impactful losses might occur. This characteristic can lead to an underestimation of risk, especially under conditions of market stress. A more comprehensive breakdown of the strengths and weaknesses of both VaR and CvaR are in Table 3.

Since VaR is intended to approximate potential losses over a given time frame, its calculation relies on a predictive distribution of future returns, which is a complex task due to the need to account for the unique behaviour of asset returns discussed in Chapter 4 (Tsay 2010, 328). Given these complexities, various estimation techniques have been developed. Nadarajah and Chan (2016, 289) categorise these approaches into three broad types: conditional (parametric) methods, which presuppose a specific distribution for the calculations; unconditional (nonparametric) methods, which makes no such distributional assumptions; and semiparametric methods, which are a hybrid of the two.

Table 3 Pros and cons of VaR and CVaR (Sarykalin et al. 2014, 282–284)

Feature	VaR Pros	VaR Cons	CVaR Pros	CVaR Cons
Concept clarity	VaR is simple and has a clear interpretation based on a specific confidence level.	VaR does not account for the tail beyond the specified confidence level, missing potential extreme losses.	CVaR provides a clear measure of the worst-case average losses beyond a specified confidence level, focusing on tail risks.	CVaR's effectiveness is highly dependent on accurate tail modelling; without a good model, CVaR estimations can be misleading.
Risk measurement	Measures potential losses at a specified confidence level.	VaR may underestimate risk in distributions with fat tails or skewness, especially in extreme market conditions.	Measures the average of the worst losses, providing a more comprehensive risk assessment.	More sensitive to estimation errors than VaR, particularly in the tail.
Application	Widely used in various industries for its straightforwardness.	Optimization can be challenging as VaR is nonconvex and discontinuous for discrete distributions, making computational problems complex.	CVaR optimization is generally easier due to its convex nature, allowing for reduction to convex or even linear programming in some cases.	-
Mathematical properties	Comparatively stable estimation procedures, less affected by tail losses.	Discontinuous and nonconvex nature makes it computationally difficult to handle in optimization scenarios.	CVaR is a consistent risk measure and is continuous with respect to confidence levels, providing mathematical consistency.	Requires robust historical data for accurate modelling, especially in financial contexts where historical data may be limited.
Use case(s)	Can rank several different distributions by comparing their VaRs at the same confidence level.	High risks in the tail of the distribution can be overlooked, leading to potentially high unexpected losses.	Allows for the shaping of portfolio distributions with multiple CVaR constraints, enhancing risk management capabilities.	Equally weighted strategies may outperform CVaR-optimized ones out of sample, particularly if the historical data are skewed.

5.6 Expected shortfall

In response to the limitations of VaR, an alternative market risk metric called expected shortfall (otherwise known as conditional value at risk, CVaR) was developed to address the tail risk overlooked by VaR. Hull (2018, 274) describes the difference between VaR and ES followingly:

> *Whereas VaR asks the question "How bad can things get?" ES asks: "If things do get bad, what is the expected loss?"*

ES is defined by Yamai and Yoshiba (2005) in a more mathematical way as the expected negative return for returns that exceed the VaR threshold at a given confidence level

$$ES_\alpha = \mathbb{E}\{x|x \geq VaR_\alpha(x)\}.$$

ES provides a broader view of market risk, especially under extreme conditions by considering the severity of losses beyond the VaR threshold. However, accurately estimating ES involves larger sample sizes and more complex calculations, which can lead to greater estimation errors and increased computation time (Yamai & Yoshiba 2005).

6 Results

This chapter will start by presenting the results from the tests regarding mean reversion. Later, the rolling risk measures will be presented. Finally, the findings will be combined and analysed. All the parameters were presented thoroughly in Chapter 5. This thesis uses a 50-day rolling window to calculate the Hurst exponent, half-life, VaR, and CVaR. Therefore, the time span of mean reversion that will be studied is only 50 days or in other words, the time it takes for stock prices to the reverse to their 50-day moving average.

The variance ratio (VR) test follows the theory presented by Chow and Denning's (1993) multiple comparison test. One change to the final value has been made for this thesis. The figure calculated by the multiple comparison test has been subtracted by one to make the difference between mean reverting and non-mean reverting behaviour more apparent. Now mean reverting behaviour is indicated by negative Z-statistic values and the opposite is true for positive values. Figure 7 displays the results of the multiple comparison test.

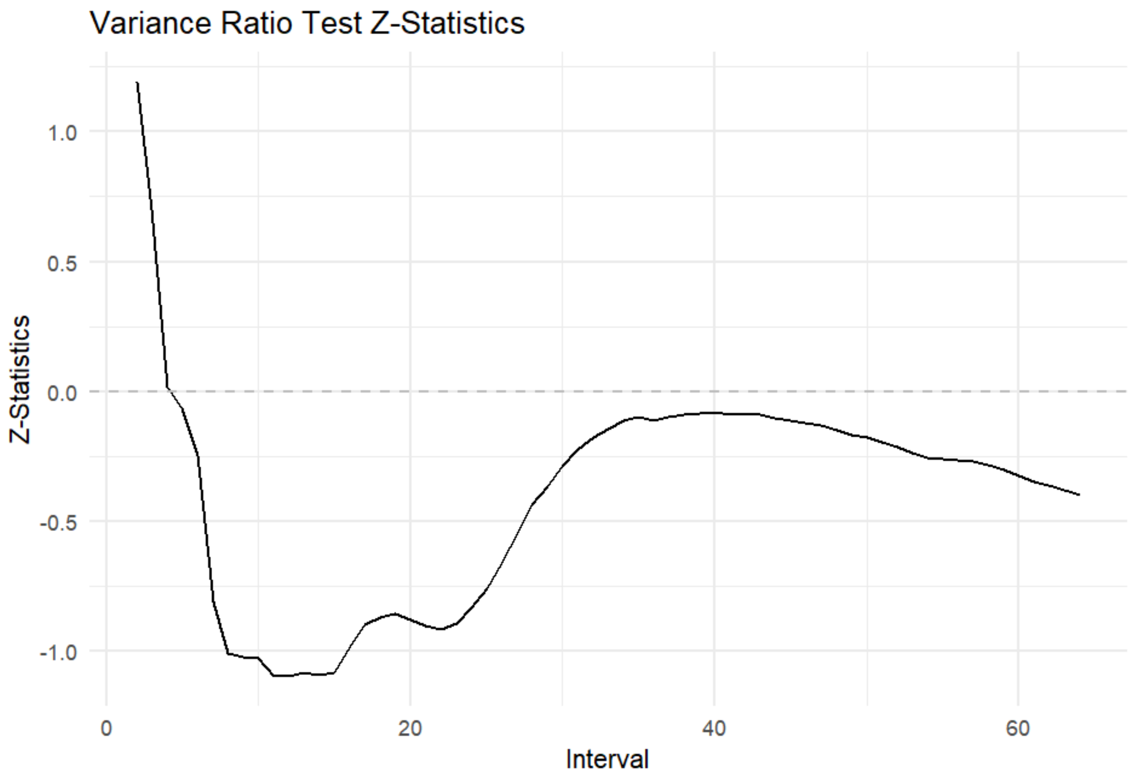

Figure 7 Variance ratio test

The presence of mean reversion in the Finnish stock market is evident based on the VR test (Figure 7). This can be concluded from the fact that the value of the Z-statistic becomes negative. Mean reversion is strongest between the intervals 8-15, because there the Z-statistic achieves its lowest values and afterwards the absolute Z-statistic value decreases. This means that it usually takes 8 to 15 trading days for price changes to go back to their mean. The Z-statistic does not go extremely negative and on longer intervals it approaches zero, indicating that the mean reversion in the Finnish stock market is not as strong on periods longer than 30 days. One interpretation of this could be that due to the strength of mean reversion when it reaches its lowest Z-statistic value of approximately -1 (around intervals 8 to 15), the returns have reached their relative mean (or close to it) and therefore no meaningful mean reversion can be achieved, which is indicated by the Z-statistic returning to around zero after the strongest period of mean reversion.

Like in the case of Lo and MacKinlay (1988), the random walk hypothesis can be rejected based on the VR test. Additionally, the VR test gives the first sign of mean reverting behaviour in the Finnish stock market. Poterba and Summers (1988) studied mean reversion in the New York Stock Exchange using yearly and monthly returns between 1926-1985. Additionally, they had 17 indices from other countries (not including Finland). Using VR, they found consistent positive autocorrelation in short periods and negative autocorrelation in longer periods, meaning that the VR can be used to detect mean reversion using daily, monthly, and yearly data.

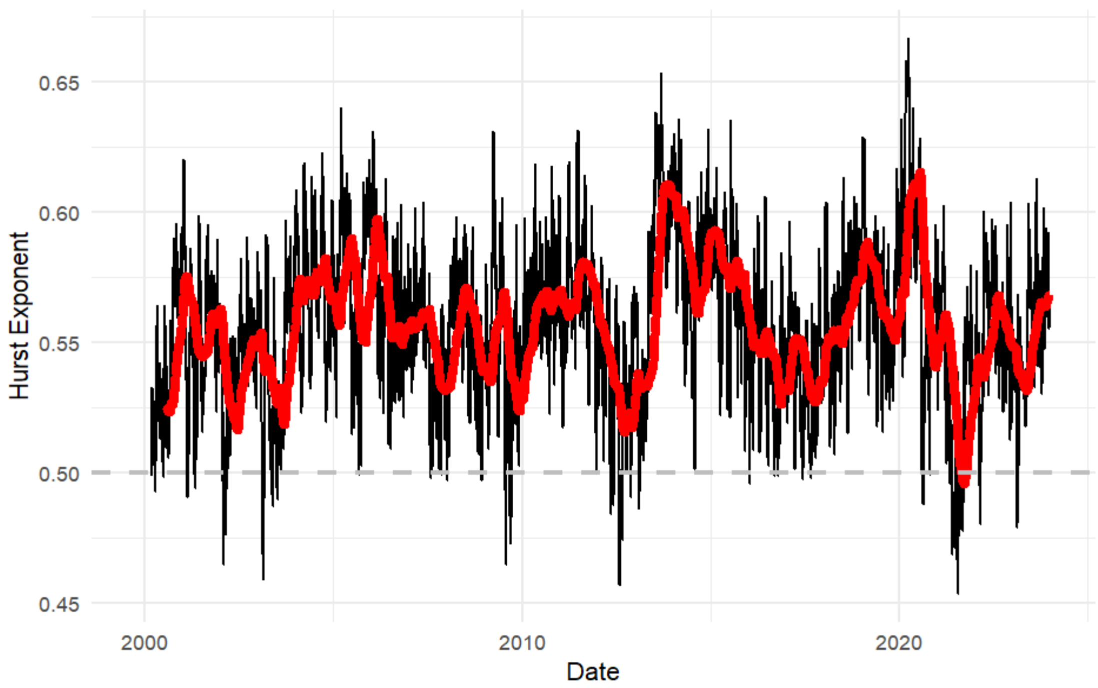

Figure 8 Hurst exponent

The Hurst exponent (HE) further validates the presence of mean reversion in the Finnish stock market. The black line in Figure 8 is the rolling value HE and was calculated using a 50-day rolling window. The red line depicts the average of the previous 100 values, therefore showcasing the movements of the HE in a clearer way. The rolling value HE was calculated as the mean of the previous 100 observations, starting with the 101st index and continuing it in a rolling fashion. As a reminder, the sign of mean reverting behaviour is a Hurst exponent value below 0.5, which in Figure 8 is the dotted grey line. Based on this, several periods of mean reversion can be identified, the most notable of which are around 2009, 2012, and 2021. These periods are when notable economic downturns occurred and out of these three, the most powerful mean reversion occurred in 2021, this being the only time in the chosen period where even the 100-day average HE fell below 0.5. The lengthy period from 2013 to 2020 of no mean reversion (even the rolling value HE depicted by the black line staying above 0.5) could be argued as being one reason behind why the reversal was so strong in 2021.

The results from the Hurst exponent data imply that mean reversion should not be thought as a phenomenon that is constantly, or even commonly, affecting the market. Rather, it mostly occurs during large market swings. Most of the time prices are persistent following a certain trend, indicated by the Hurst exponent staying above the grey dotted line. Serletis and Rosenberg's (2009) analysis reached a similar conclusion using for four different U.S. stock indices. They found the U.S. stock market strongly mean reverting.

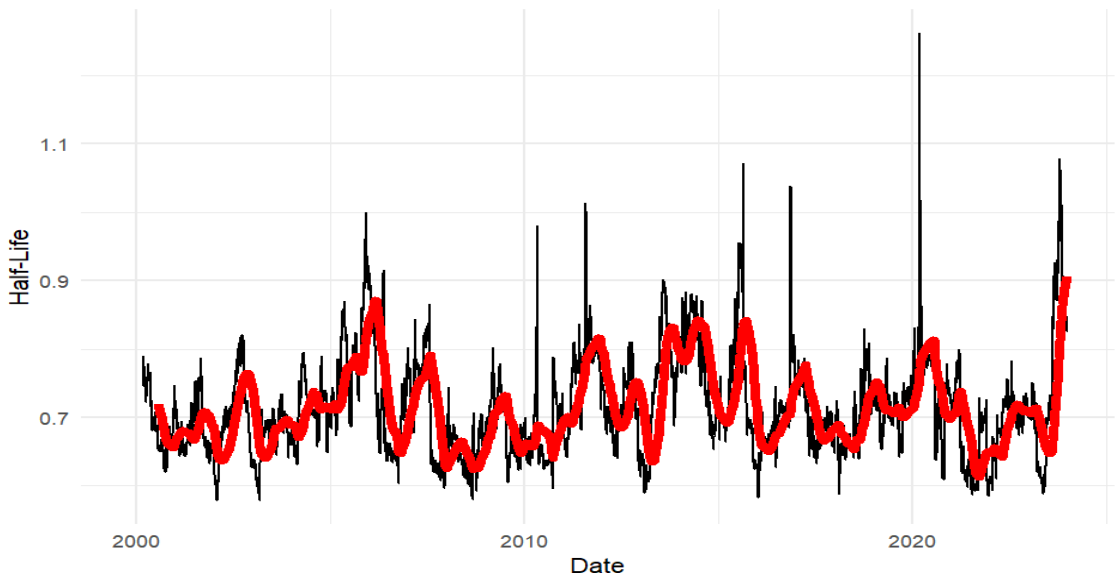

Figure 9 Half-life

As we already observed that there is mean reversion in the Finnish stock market, the half-life can be used to assess the strength of mean reversion. In Figure 9, the black line is the rolling half-life value and was calculated using a 50-day rolling window. the red line is the rolling 100-day average, like what was shown in Figure 8 with the Hurst exponent. Essentially, the lower the half-life is, the stronger the mean reversion as it takes a shorter period for the returns to come back to their mean. The periods where the half-life is the shortest, and therefore the mean reversion at its strongest, are the same as with the Hurst exponent (most notably 2009, 2012, 2021). As with the Hurst exponent, the strongest mean reversion period occurred during 2021. It should be noted that the correlation between the Hurst exponent and half-life is 0.615 and statistically significant (p-value < 0.001), indicating a strong positive correlation.

Chaudhuri and Wu (2003) and Spierdijk et al. (2012) studied mean reversion with half-life using monthly emerging market data and yearly OECD country data, respectively. Both studies found mean reverting behaviour in their data. Additionally, Spierdijk et al. (2012, 3) observed large fluctuations in the rate of mean reversion over time. While there is notable fluctuation in the daily data's half-life, it is not as profound as Spierdijk et al. (2012, 3) found using yearly data. They found half-lives ranging from 2.1 years to 23.8 years. Again, as with the Hurst exponent, half-life can be used to find mean reverting behaviour for yearly, monthly, and daily data.

A key observation from Figures 8 and 9 is that the strength of mean reversion changes in time, which is homogenous with the findings of Kim et al. (1991) and Spierdijk et al. (2012). Reasons for this can be numerous but can most likely be explained by the political and economic environment at a specific period (Fama & French 1988, 264; Spierdijk et al. 2012, 228). Another observation is that mean reversion is often at its strongest when stock prices are going down. This indicates that stock market downturns are more profound than positive stock market movements and that when stock prices return their mean, it is usually because they have been above, rather than below, their mean. In other words, the stock market has been above its fundamental value (if you consider the mean to depict fundamental value) before reverting.

In Figure 10, the daily VaR (blue line) and CVaR (red line) are shown. The risk measures were calculated in a rolling fashion by using data from the previous 100 days. The highest values are captured in 2001, 2009, and 2020 which are similar to the periods where the strongest mean reversions were observed.

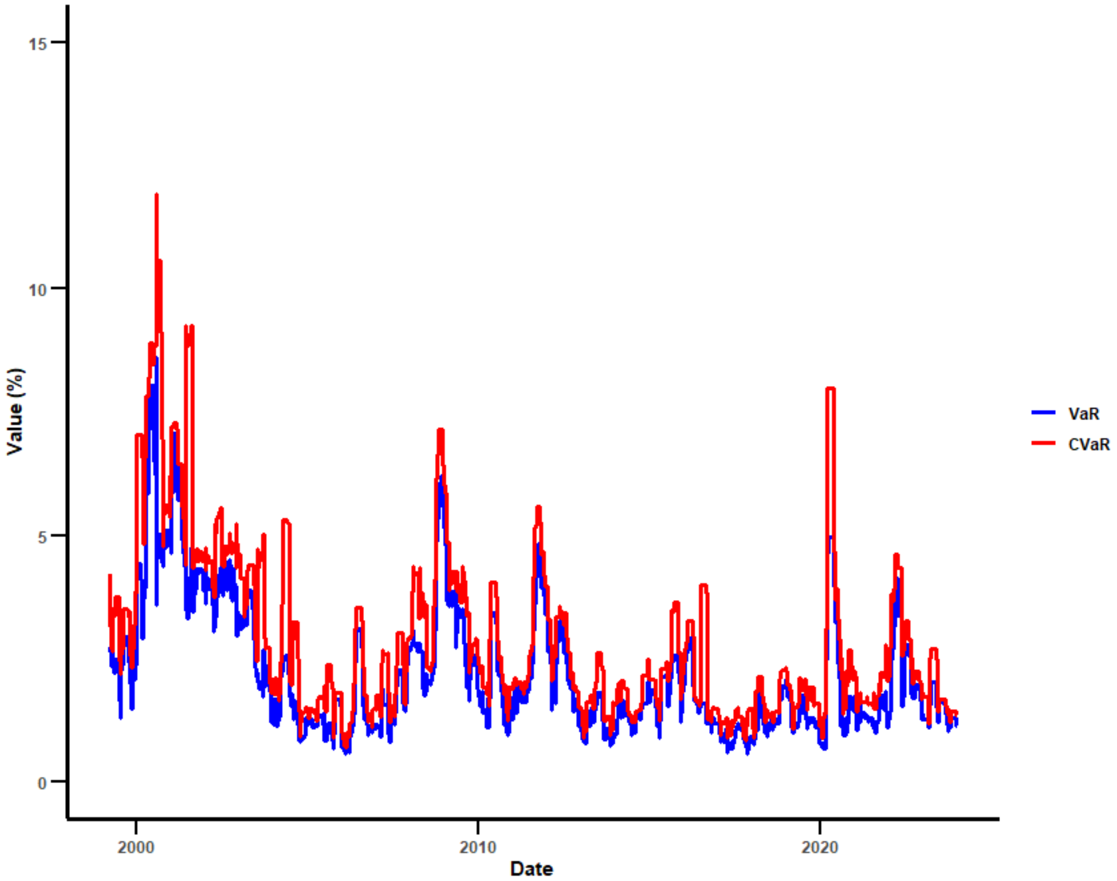

Figure 10 Daily VaR and CVaR calculated using a 100-day window

Visually, it appears that the risk measures become more pronounced (the value at risk increases) when mean reversion is stronger, indicating a positive correlation with the amount of risk and strength of mean reversion in the Finnish stock market. The correlations between risk and mean reversion measures are presented in Table 3.

Table 4 Correlation between risk and mean reversion measures

	Hurst exponent	Half-life
VaR	0.208755	0.244653
CVaR	0.144425	0.226792

For the correlation analysis in Table 3, the 100-day rolling averages for VaR and CvaR were calculated to make them comparable to the Hurst exponent and half-life measures. All the correlations indicate weak positive correlation. VaR had a stronger correlation with both mean reversion measures, while half-life was stronger correlated with the risk measures. All the correlations are statistically significant at a significance level of 0,1%. Therefore, half-life can be considered a superior method of assessing the affect of mean reversion on risk. This observation is further confirmed by parametric models used to evaluate the relationship between mean reversion and the chosen risk measures.

Further analysing the possible causes behind mean reversion gives at least two options, one where market efficiency rules and one where it does not. Assuming market efficiency, the price of a stock (or index) is determined by its future returns per share. As shown by Summers (1986), mean reversion occurs when these expected returns are mean reverting. Combining this finding with the results of this thesis, during high periods of high economic uncertainty, expected returns most likely deviate from their long-term values and quickly revert back to these levels, much faster than during more stable periods. When economic uncertainty dissipates, expected returns tend to rise significantly over a short period, contributing to the high speed of mean reversion. Actions by financial and governmental institutions to restore stability may also accelerate this adjustment process, as was the case during COVID-19.

Outside the efficient market framework, mean reversion could also be driven by the irrational behaviour of speculative short-term traders, leading to stock prices that significantly diverge from the fundamental values. Such irrational pricing behaviour can be due to overreactions to financial news as shown by De Bondt and Thaler (1985). This can manifest itself in two ways: at the onset of uncertainty, overreaction to negative news can drive stock prices well below their intrinsic values, and during recovery, overreaction to positive news can push stock prices far above their intrinsic values. In both scenarios, significant price swing over a short period lead to rapid mean reversion.

This thesis checks for heteroskedasticity in two separate regression models where the dependent variables separately are VaR and CVaR and in both regression models the independent variables are the Hurst exponent and the rolling half-life. The Breusch-Pagan (BP) test is used to examine the presence of heteroskedasticity in the data (Breusch & Pagan 1979). The BP test's null hypothesis of homoskedasticity is rejected for both regression models. Following this result, this thesis will implement the heteroscedasticity-consistent covariance matrix estimators popularised by Halbert (1980), otherwise known and later referred to as robust standard errors (RSE), to compute the relationship between the variables.

Table 5 Results for the RSE method

1st model (VaR)	Estimate	Std. Error	t-value	Significance
VaR (dependent)	0.0333542	0.00099934	33.3762	***
Hurst exponent	-0.00225947	0.00082741	-2.7308	**
Half-life	-0.01394741	0.00124579	-11.1956	***
2nd model (CVaR)	Estimate	Std. Error	t-value	Significance
CVaR (dependent)	0.04608153	0.00099934	46.1119	***
Hurst exponent	-0.00227397	0.00082741	-2.7483	**
Half-life	-0.02185762	0.00124579	-17.5451	***

Significance codes: '***' 0.001, '**' 0.01, '*' 0.05, '.' 0.1, ' ' 1

The results for the RSE method are in Table 5. The first model used VaR as the dependent variable and Hurst exponent and half-life as the independent variables. Both independent variables are statistically significant at the one per cent level. For each one unit increase in Hurst exponent the VaR decreases by approximately 0.0023. This indicates that more persistent movements, and therefore a higher Hurst exponent, results in lower value at risk. Only when prices start to become mean reverting (Hurst exponent decreases) the VaR goes up. The same could be said for the half-life measure, although the effect of changes are more pronounced, as a one unit increase in half-life results in a 0.0139 decrease in VaR. Therefore, changes in half-life are more impactful than changes in Hurst exponent.

The results are practically identical for VCaR, although slightly more noticeable. Both independent variables are statistically significant at the one per cent level. An increase of one unit in Hurst exponent results in an identical decrease of 0.0023 for CVaR as with VaR. The more sizeable difference is with how CVaR behaves to changes on half-life. One unit of increase in half-life results in a decrease of 0.021 for CVaR.

7 Conclusions

The purpose of this thesis was to evaluate whether the Finnish stock market (FSM) was mean reversing and if it was, what was its effect on chosen risk measures, value at risk (VaR) and expected shortall (ES or CVaR). A quarter of a century of daily data spanning from January 1999 to December 2023 from the OMXHPI was used to first find out whether the FSM was mean reverting and if it was, how would it affect the aforementioned risk measures. There were three separate methods used to evaluate whether the FSM was mean reverting or not: variance ratio, Hurst exponent, and half-life.

Based on the methods used, the FSM is clearly mean reverting. Although that is true, that does not mean that mean reversion is a constant phenomenon affecting the FSM. While there were periods of notable mean reversion (2009, 2012, 2021), there were also lengthy periods of up to eight years (2013–2020) where one could argue no mean reversion took place. This could also be seen as a reason why the reversal to the mean was so strong in 2021. Mean reversion was at its strongest around the times of significant economic uncertainty (global financial crisis around 2009 and COVID-19 around 2021) and during times of low economic uncertainty, mean reversion appears to be almost non-existent (trending behaviour).

The affect of mean reversion on risk was evaluated using the robust standard errors method that takes into account the heteroskedastic nature of stock returns. Both VaR and CVaR are affected by the persistence of price movements and the mean reverting nature of of stock prices. The findings indicate that as market behaviour shifts from persistent trends to mean reverting patterns, risk measures like VaR and CVaR tend to rise. Additionally, half-life appears to be the more impactful of the variables used (Hurst exponent and half-life) in explaining the changes of the risk measures. CVaR shows greater sensitivity to changes in half-life compared to VaR, indicating that CVaR may provide a more responsive assessment of risk under different market conditions.

The results suggest that the FSM is mean reverting which is consistent with previous studies on the subject of mean reversion conducted for stock indices in other regions (see Poterba and Summers 1988). Possible explanations for this mean reversion could be numerous, but this thesis raises two: expected returns reverting to their mean and irrational behaviour by traders caused by overreactions to new information.

This thesis used a value-weighted index for its analysis. In future research, an equally weighted index could be used and compared, whether the weightings affect the amount of mean reversion or not. Additionally, more complicated methods of detecting mean reversion, such as the Ornstein-Uhlenbeck process or the Vasicek model, could be used. This thesis argued that the length of the persistent period could affect the strength of the following mean reverting period but could not claim this as a certainty. Therefore, this could be another angle of future research on mean reversion. Another topic of research around mean reversion that could bring value would be if mean reversion could be predicted and what are the actual reasons behind mean reversion. Lastly, as this thesis used daily data, the affect of different data intervals (weekly, monthly) on risk measures could be studied, whether there are any differences in the results compared to daily data.

References

Alexander, Carol (2008). *Market Risk Analysis Volume IV: Value-at-Risk Models*. 1st edition. Chichester, England. Wiley.

Arditti, Fred (1971). Another Look at Mutual Fund Performance. *Journal of financial and quantitative analysis*. 6 (3). 909–912.

Arsalan, Tazeen – Chishty, Bilal – Ghouri, Shagufta – Ansari, Nayeem (2022). Comparison of volatility and mean reversion among developed, developing and emerging countries. *Journal of economic and administrative sciences*.

Balvers, Ronald – Wu, Yangru – Gilliland, Erik (2000). Mean reversion across national stock markets and parametric contrarian investment strategies. *Journal of Finance*. 55 (2). 745–772.

Bao, Jack – Pan, Jun (2013). Bond Illiquidity and Excess Volatility. *The Review of financial studies*. 26 (12). 3068–3103.

Breusch, Trevor – Pagan, Adrian (1979). A Simple Test for Heteroscedasticity and Random Coefficient Variation. *Econometrica*. 47 (5). 1287–1294.

Campbell, John – Shiller, Robert (2005). *Advances in Behavioural Finance*. 2nd edition. Princeton University Press.

Charles, Amélie – Darné, Olivier (2009). Variance-ratio Tests of Random Walk: An Overview. *Journal of economic surveys*. 23 (3). 503–527.

Chaudhuri, Kausik – Wu, Yangru (2003). Mean reversion in stock prices: evidence from emerging markets. *Managerial finance*. 29 (10). 22–37.

Chen, James (2023). What is Mean Reversion, and How Do Investors Use It? <https://www.investopedia.com/terms/m/meanreversion.asp>, read 2.1.2024.

Chitenderu, Tafadzwa – Maredza, Andrew – Sibanda, Kin (2014). The Random Walk Theory and Stock Prices: Evidence from Johannesburg Stock Exchange. *The international business & economics research journal*. 13 (6). 1241–1249.

Chow, Victor – Denning, Karen (1993). A simple multiple variance ratio test. *Journal of econometrics*. 58 (3). 385–401.

Clark, Peter (1973). A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica*. 41 (1). 135–155.

Clarke, Jonathan – Jandik, Tomas – Mandelker, Gershon (2001). The Efficient Markets Hypothesis. *Expert financial planning: Advice from industry leaders*. 7 (3). 126–141.

Cont, Rama (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative finance*. 1 (2). 1–14.

Cont, Rama (2007). Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models. *Long Memory in Economics*. Berlin, Heidelberg: Springer Berlin Heidelberg. 289–309.

Cutler, David – Poterba, James – Summer, Lawrence (1989). What Moves Stock Prices. *Journal of portfolio management*. 15 (3). 4–12.

De Bondt, Werner – Thaler, Richard (1985). Does the Stock Market Overreact? *The Journal of finance (New York)*. 40 (3). 793–805.

Ding, Zhuanxin – Granger, Clive – Engle, Robert (2001). A Long Memory Property of Stock Market Returns and a New Model. *Journal of Empirical Finance*. 1. 83–116.

Engel, Charles – Morris, Charles (1991). Challenges to Stock Market Efficiency: Evidence from Mean Reversion Studies. *Economic review (Kansas City)*. 76 (5). 21–45.

Fama, Eugene (1963). Mandelbrot and the Stable Paretian Hypothesis. *The Journal of business*. 36 (4). 420–429.

Fama, Eugene (1965). The Behaviour of Stock-Market Prices. *The Journal of business*. 38 (1). 34–105.

Fama, Eugene (1965). Random Walks in Stock Market Prices. *Financial analysts journal*. 51 (1). 75–80.

Fama, Eugene (1970). Efficient Capital Markets: A Review of Empirical Work. *Journal of Finance*. 25 (2). 383–417.

Fama, Eugene – French, Kenneth (1988). Permanent and temporary components of stock prices. *Journal of Political Economy*. 96 (2). 246–273.

Ghashghaie, S. – Breymann, W. – Peinke, J. – Talkner, P. – Dodge, Y. (1996). Turbulent cascades in foreign exchange markets. *Nature*. 381 (6585). 767–770.

Ghose, Devajyoti – Kroner, Kenneth (1995). The relationship between GARCH and symmetric stable processes: Finding the source of fat tails in financial data. *Journal of empirical finance*. 2 (3). 225–251.

Gopikrishnan, P. – Plerou, V. – Liu, Y. – Amaral, L. – Gabaix, X. – Stanley, H. (2000). Scaling and correlation in financial time series. *Physica A*. 287 (3). 362–373.

Griffin, Dale – Tversky, Amos (1992). The Weighing and the Determinants of Confidence. *Cognitive Psychology*. 24 (1). 411–435.

Guillaume, Dominique – Dacorogna, Michel – Davé, Rakhal – Müller, Ulrich – Olsen, Richard – Pictet, Olivier (1997). From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and stochastics*. 1 (2), 95–129.

Hull, John (2018). *Risk management and financial institutions*. 5th edition. New Jersey. Wiley.

Jegadeesh, Narasimhan (1991). Seasonality in Stock Price Mean Reversion: Evidence from the U.S. and the U.K. *The Journal of finance*. 46 (4). 1427–1444.

Jegadeesh, Narasimhan ─ Titman, Sheridan (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of finance (New York)*. 48 (1). 65–91.

Kim, Myung – Nelson, Charles – Startz, Richard (1991). Mean Reversion in Stock Prices? A Reappraisal of the Empirical Evidence. *The Review of economic studies*. 58 (3). 515–528.

Kim, Hyon-Joo – Lee, Youngki ─ Kahng, Byungnam ─ Kim, In-mook (2002). Weighted scale-free network in financial correlations. *Journal of the Physical Society of Japan*. 71 (9). 2133–2136.

Kim, Jae (2006). Wild bootstrapping variance ratio tests. *Economics letters*. 92 (1). 38–43.

Kim, Jae – Ji, Philip (2011). Mean-reversion in international real interest rates. *Economic modelling*. 28 (4). 1959–1966.

Knüpfer, Samuli – Puttonen, Vesa (2018). *Moderni rahoitus*. 10th edition. Alma Talent, Helsinki.

Kon, Stanley (1984). Models of Stock Returns – A Comparison. *The Journal of finance*. 39 (1) 147–165.

Lo, Andrew – MacKinlay, A. (1988) Stock market prices do not follow random walks: evidence from a simple specification test. *Review of Financial Studies*. 1 (1). 41–66.

Malliaropulos, Dimitrios – Priestley, Richard (1999). Mean reversion in Southeast Asian stock markets. *Journal of empirical finance*. 6 (4). 355–384.

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

Mandelbrot, Benoit (1966). Forecasts of Future Prices, Unbiased Markets, and "Martingale" Models. *The Journal of Business*. 39 (1). 242–255.

Merton, Robert (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *The Journal of finance*. 29 (2). 449–470.

Miller, Michael (2019). *Quantitative financial risk management*. 5th edition. New Jersey. Wiley.

Nasdaq OMX Nordic. Data for OMXHPI-index. Online. Read 1.3.2024. https://www.nasdaqomxnordic.com/indeksit/index_info?Instrument=FI0008900006.

Ning, Cathy – Xu, Dinghai – Wirjanto, Tony (2015). Is volatility clustering of asset returns asymmetric? *Journal of banking & finance*. 52 (1). 62–76.

Odean, Terrance (1998). Volume, Volatility, Price, and Profit When All Traders Are Above Average. *The Journal of finance (New York)*. 53 (6), 1887–1934.

Pagan, Adrian (1996). The econometrics of financial markets. *Journal of empirical finance*. 3 (1). 15–102.

Patton, Andrew – Ziegel, Johanna – Chen, Rui (2019). Dynamic semiparametric models for expected shortfall (and Value-at-Risk). *Journal of econometrics*. 211 (2). 388–413.

Peng, C. – Buldyrev, S. – Havlin, S. – Simmons, M. – Stanley, H. – Goldberger, A. (1994). Mosaic organization of DNA nucleotides. *Physical Review E*. 49 (1). 1685–1689.

Poterba, James – Summers, Lawrence (1988). Mean reversion in stock prices. Evidence and Implications. *Journal of financial economics*. 22 (1). 27–59.

Rogoff, Kenneth (1996). The Purchasing Power Parity Puzzle. *Journal of economic literature*. 34 (2). 647–668.

Roll, Richard (1984). Orange juice and weather. *The American economic review*. 74 (5). 861–880.

Samuelson, Paul (1965). Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial management review*. 6 (2). 41–49.

Sarykalin, Sergey – Serraino, Gaia – Uryasev, Stan (2014). Value-at-Risk vs. Conditional Value-at-Risk in Risk Management and Optimization. *Tutorials in Optimization Research*. INFORMS. Maryland, USA. 270–294.

Serletis, Apostolos – Rosenberg, Aryeh (2007). The Hurst exponent in energy futures prices. *Physica A*. 380 (1–2). 325–332.

Serletis, Apostolos – Rosenberg, Aryeh (2009). Mean reversion in the US stock market. *Chaos, solitons and fractals*. 40 (4). 2007–2015.

Shleifer, Andrei (2000). *Inefficient markets: an introduction to behavioural finance*. Oxford University Press, Oxford.

Singleton, Clay – Wingender, John (1986). Skewness Persistence in Common Stock Returns. *Journal of financial and quantitative analysis*. 21 (3). 335–341.

Spierdijk, Laura – Bikker, Jacob – Hoek, Pieter van den (2012). Mean reversion in international stock markets: An empirical analysis of the 20th century. *Journal of international money and finance*. 31 (2). 228–249.

Stoline, Michael – Ury, Hans (1979). Tables of the Studentized Maximum Modulus Distribution and an Application to Multiple Comparisons Among Means. *Technometrics*. 21 (1). 87–93.

Summers, Lawrence (1986). Does the Stock-Market Rationally Reflect Fundamental Values. *The Journal of finance (New York)*. 41 (3). 591–601.

Thaler, Richard (1987). Anomalies The January Effect: Why a Feature on Anomalies? *The Journal of economic perspectives*. 1 (1). 197–201.

Thaler, Richard (1987). Anomalies Seasonal Movements in Security Prices II: Weekend, Holiday, Turn of the Month, and Intraday Effects. *The Journal of Economic Perspectives.* 1 (2). 169–177.

Tsay, Ruey (2010). *Analysis of financial time series*. 3rd edition. Wiley.

Weron, Rafał (2002). Estimating long-range dependence: finite sample properties and confidence intervals. *Physica A*. 312 (1). 285–299.

White, Halbert (1980). A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica*. 48 (4). 817–838.

Wright, Jonathan (2000). Alternative Variance-Ratio Tests Using Ranks and Signs. *Journal of business & economic statistics*. 18 (1). 1–9.

Yamai, Yasuhiro – Yoshiba, Toshinao (2005). Value-at-risk versus expected shortfall: A practical perspective. *Journal of banking & finance*. 29 (4). 997–1015.