ICHEC BRUSSELS MANAGEMENT SCHOOL

UCLouvain

An empirical investigation of weak form market efficiency in the Belgian stock market

Thesis presented by:
Kaya Dogan

Supervisor :
Christophe Desagre

Master of Business Management
Master of Sciences in Management

ACADEMIC YEAR 2023-2024

19/08/2024

Acknowledgements

I would like to express my most sincere thanks to my supervisor, Christophe Desagre, for his guidance, support, and encouragement throughout the process of this thesis. His expertise and insights were instrumental in shaping the direction and quality of my research.

I also extend my gratitude to the teaching staff of ICHEC and UCL for the quality of the courses and the support provided throughout this dual master's program.

Finally, my deepest appreciation goes to my close circle.

Anti-plagiarism commitment

I, the undersigned, DOGAN, Kaya, in my final year of master, hereby declare that the attached work complies with the referencing rules as outlined in the study regulations signed upon my enrollment at ICHEC (adherence to the APA standard for in-text referencing, bibliography, etc.); that this work is the result of an entirely personal effort; that it does not contain any content produced by artificial intelligence without explicitly referencing it. By my signature, I certify on my honor that I have taken note of the aforementioned documents and that the submitted work is original and free from any uncredited borrowing from a third party.

19/08/2024

I, the undersigned, DOGAN, Kaya, 190332, solemnly declare the following regarding the use of artificial intelligence (AI) in my thesis:

Type d'assistance		Case à cocher
Aucune assistance	J'ai rédigé l'intégralité de mon travail sans avoir eu recours à un outil d'IA générative.	
Assistance avant la rédaction	J'ai utilisé l'IA comme un outil (ou moteur) de recherche afin d'explorer une thématique et de repérer des sources et contenus pertinents.	
Assistance à l'élaboration d'un texte	J'ai créé un contenu que j'ai ensuite soumis à une IA, qui m'a aidé à formuler et à développer mon texte en me fournissant des suggestions.	
	J'ai généré du contenu à l'aide d'une IA, que j'ai ensuite retravaillé et intégré à mon travail.	
	Certains parties ou passages de mon travail/mémoire ont été entièrement été générés par une IA, sans contribution originale de ma part.	
Assistance pour la révision du texte	J'ai utilisé un outil d'IA générative pour corriger l'orthographe, la grammaire et la syntaxe de mon texte.	X
	J'ai utilisé l'IA pour reformuler ou réécrire des parties de mon texte.	X
Assistance à la traduction	J'ai utilisé l'IA à des fins de traduction pour un texte que je n'ai pas inclus dans mon travail.	
	J'ai également sollicité l'IA pour traduire un texte que j'ai intégré dans mon mémoire.	
Assistance à la réalisation de visuels	J'ai utilisé une IA afin d'élaborer des visuel, graphiques ou images.	
Autres usages		X

I commit to adhering to these declarations and to providing any additional information required regarding the use of AI in my thesis, specifically: I have included in the appendix the questions asked to the AI, and I am able to provide the questions posed and the responses obtained from the AI. I can also explain the type of assistance I used and for what purpose.

Done at Brussels, on 19/08/2024

Kaya Dogan 190332

Table of contents

1. Introduction

Market efficiency is a fundamental concept in finance, which stipulates that financial asset prices reflect all available information at any given time. This concept was widely explored and popularized by Eugene Fama in the 1970s. According to Fama (1970), a market is efficient if "prices fully reflect all available information at any time." In other words, the stock market operates like a vast database of financial information. Each new piece of data, whether it's a company earnings report, a merger announcement, or even a rumor is immediately integrated into stock prices. Thus, no investor has an informational advantage over others, as all pertinent information is already reflected in current stock prices as soon as it is available. That is why in an efficient market, it is impossible for an investor to consistently achieve returns higher than those of the market, as asset prices quickly and accurately adjust to any new information. Only future information, which is unpredictable, is susceptible to impacting the stock price. Superior risk adjusted returns compared to the market are thus considered to be caused by chance alone.

The Efficient Market Hypothesis (EMH) can be broken down into three distinct forms. The weak form asserts that financial asset prices incorporate all information contained in past prices. This means that technical analysis, which relies on studying historical price trends, does not allow for superior returns to those of the market. In other words, past price changes cannot predict future price changes. The semi-strong form suggests that asset prices incorporate not only all information contained in past prices but also all publicly available information, such as financial reports, economic announcements, and industry news. Therefore, fundamental analysis, which aims to assess the intrinsic value of an asset by studying its financial and economic data, also cannot allow beating the market as this information is already accounted for in current prices. The strong form asserts that asset prices reflect all information, both public and private. Even insiders, who have access to confidential information about a company, cannot achieve abnormal returns. This form of efficiency asserts that all types of information are instantaneously integrated into asset prices, making it impossible to exploit privileged information for profit.

Market efficiency has profound implications for investors, fund managers, and regulators. For investors, it means that trading strategies based on technical or fundamental analysis are unlikely to generate abnormal returns. For fund managers, it suggests that active management, which seeks to select winning stocks, may not offer better performance than passive management, which tracks a market index. Finally, for regulators, it implies that transparency and equal access to information are essential to maintaining fairness and market efficiency.

This thesis focuses on the weak form efficiency of the Belgian stock market by testing whether stock prices follow a random walk and if there are exploitable inefficiencies for investors. The objective is to determine whether past returns can predict future returns and to examine if the Belgian stock market, although located in a developed country, exhibits characteristics similar to those of less followed and less liquid markets due to its lower coverage compared to other developed countries.

First, the literature on the subject will be explored to gain an understanding of how the study of this concept has evolved over time. The literature review will also provide an overview of the current state of knowledge in the field. Secondly, we will examine the weak-form efficiency literature more specifically by diving deeper into the concept and its implications, highlighting popular statistical tools used to test it. We will then zoom in on empirical findings gathered in various markets and asset classes to gain a better understanding of its practical examination. Thirdly, the methodology related to the various statistical tests performed in this thesis will be detailed, followed by the presentation of their results. Finally, trading strategies based on these findings will be created and tested to determine if the potential inefficiencies identified with the statistical tools can be exploited for profit, providing a practical framework to confirm their presence.

2. Literature review

2.1. Origins and historical development

The history of market efficiency begins with the work of Louis Bachelier in 1900. In his doctoral thesis entitled "Théorie de la spéculation," Bachelier proposed for the first time that financial asset prices follow a random walk process. He used sophisticated mathematical tools for his time, including the concepts of Brownian motion and stochastic integration, to model price variations of financial securities. This revolutionary idea laid the foundation for modern market efficiency theory. However, Bachelier's work was largely ignored for several decades, mainly because his ideas were ahead of their time and the mathematical tools needed to fully understand them were not yet widely used in finance.

2.2. Contributions of Cowles and researchers of the 1930s

In the 1930s, Alfred Cowles and his colleagues conducted empirical studies on the performance of investors and financial analysts. Cowles (1933) analyzed the performance of financial forecasting services and investment funds, concluding that professionals do not consistently outperform the market, suggesting a form of market efficiency. Furthermore, the work of Holbrook Working (1934) and Cowles and Jones (1937) also confirmed Bachelier's ideas. Working (1934) demonstrated that the time series of commodity prices follow a random walk process, while Cowles and Jones (1937) found that financial analysts' forecasts were often incorrect, reinforcing the idea that asset prices already incorporate all available information.

2.3. The era of Samuelson and Mandelbrot

In the 1960s, the work of Paul Samuelson and Benoît Mandelbrot played a crucial role in the development of market efficiency theory. Samuelson (1965) formalized the idea that efficient markets should immediately incorporate all available information, making asset prices unpredictable. He mathematically showed that if a market is efficient, asset prices must follow a random walk. Samuelson also emphasized the importance of price anticipation by investors, explaining that current asset prices should reflect investors' future return expectations.

Mandelbrot (1966), for his part, studied the statistical properties of asset prices and introduced the concept of long-tail distributions, challenging the normality assumption of returns. He demonstrated that the variations in financial asset prices do not always follow a normal distribution but rather a leptokurtic distribution with fat tails. These findings led to a better understanding of volatility and risk in financial markets.

Other significant contributions during this period include the work of Fama himself, as well as that of Granger and Morgenstern (1963), who studied the predictability of stock prices and found additional evidence of market efficiency.

2.4. Formalization by Eugene Fama

The modern formalization of the Efficient Market Hypothesis (EMH) is attributed to Eugene Fama, who systematized the theory in his 1970 paper "Efficient Capital Markets: A Review of Theory and Empirical Work." Fama defined three forms of efficiency: weak, semi-strong, and strong. Weak form efficiency postulates that asset prices reflect all information contained in past prices. Semi-strong form efficiency adds that asset prices also incorporate all publicly available information. Finally, strong form efficiency asserts that asset prices reflect all information, both public and private. Fama also introduced several key concepts in understanding financial markets. He formalized the idea that asset prices follow a random walk, where price changes are independent and identically distributed. This means that past returns cannot predict future returns. He also acknowledged the existence of market anomalies, such as the small-cap effect and the calendar effect, but argued that these anomalies are not systematic enough to challenge the hypothesis of market efficiency. Fama developed and applied various empirical tests to evaluate market efficiency, including correlation tests and return predictability tests.

2.5. Post-Fama and the emergence of behavioral finance

Following Fama's work, research on market efficiency continued to evolve, with the emergence of antagonistic theories such as behavioral finance. This field opposes the EMH by arguing that investors are not always rational and that financial markets are influenced by psychological biases and irrational behaviors. Researchers like Daniel Kahneman and Amos Tversky showed that investors are subject to cognitive biases, such as overconfidence, loss aversion, and the disposition effect (Kahneman and Tversky, 1979). These biases can lead to market anomalies and asset prices that do not fully reflect all available information. Robert Shiller, another pioneer of behavioral finance, demonstrated that markets can be subject to speculative bubbles and financial crises due to irrational investor behavior (Shiller, 1981). Behavioral finance has generated considerable controversy in the financial literature, with some researchers arguing that markets are often inefficient due to irrational behavior, while others continue to defend the efficient market hypothesis. This controversy has enriched the academic debate and led to a better understanding of financial markets and their mechanisms.

2.6. Adaptive Markets Hypothesis

The Adaptive Markets Hypothesis (AMH), proposed by Andrew Lo, offers a new paradigm that reconciles the EMH with behavioral finance. According to Lo (2004, 2005), markets are not always efficient in the traditional sense but evolve over time as participants adapt to changing environments. The AMH applies principles of evolution such as competition, adaptation, and natural selection to financial markets, suggesting that market efficiency is not static but dynamic and context-dependent.

The AMH asserts that market participants use heuristics, or rules of thumb, that have evolved through trial and error. These heuristics may work well under certain conditions but can lead to inefficiencies when the environment changes. Unlike the EMH, which assumes that all participants are rational and markets are always efficient, the AMH acknowledges that behavior and efficiency vary over time (Lo, 2004, 2005).

The AMH implies that profit opportunities and market inefficiencies will arise and disappear as market conditions change. It suggests that financial markets are more like biological ecosystems, where different species (types of investors) interact and adapt to their environment. This perspective helps to explain why markets may appear inefficient at times and why some anomalies persist (Lo, 2004).

By viewing market efficiency through an evolutionary lens, the AMH provides a framework where the EMH and behavioral finance can coexist. It recognizes that while markets may be efficient in the long run, short-term inefficiencies and behavioral biases can exist due to the adaptive nature of market participants (Lo, 2005).

3. Weak Form Efficiency

Weak form efficiency implies that current asset prices incorporate all information from past prices. A market is considered weak form efficient if prices follow a random walk, where price changes are independent and identically distributed. This axiom of weak-form market efficiency implies that day-to-day price variations are independent of past variations, and it is impossible to predict future price movements based on past movements, they are therefore random and unpredictable.

The random walk concept has significant implications for investors and analysts. If prices follow a random walk, it means that past returns cannot be used to predict future returns. Consequently, technical analysis, which relies on studying past price trends to predict future movements, would be ineffective in a weak form efficient market. Additionally, this implies that arbitrage opportunities, where investors could exploit price anomalies for risk-free profits, would be rare or non-existent.

Methods commonly used to test for weak form efficiency include statistical tests aimed at detecting deviations from a random walk such as the runs tests, the Variance Ratio test and the Hurst exponent, among others. Another approach to test for weak form efficiency is to implement trading strategies and observe whether these can consistently generate abnormal returns, as this would indicate predictability in past price information. This way of testing efficiency echoes Samuelson (1965). According to Paul Samuelson's definition of market efficiency, a market is considered efficient if it is impossible to consistently generate abnormal returns using available information. Thus, even if there are anomalies or patterns in the market, as long as no one can consistently exploit them for profit, the market is considered efficient. Consequently, the ultimate test of efficiency, according to Samuelson's view, is to determine whether any strategies can consistently outperform the market. This perspective aligns with the idea that markets are "informationally efficient," meaning all known information is already incorporated into prices, making it difficult for traders to gain an edge.

3.1. Runs tests

As early as 1940 was the year in which two statisticians named Wald and Wolfowitz published the first test used to measure the market efficiency hypothesis. The so-called "runs test". The essence of a runs test is to examine sequences (runs) of similar signs in a time series, such as consecutive days of positive or negative returns. This method does not assume a specific distribution for returns, making it versatile for analyzing various financial datasets.

The runs test helps in identifying whether the occurrences of returns above or below a median value are random. If the number of observed runs significantly deviates from what is expected under the assumption of randomness, it suggests the presence of serial correlation in the data, thus indicating potential market inefficiency. Runs tests have been widely applied in finance to challenge the weak form of the Efficient Market Hypothesis (EMH), particularly in detecting patterns and predictability in stock prices that could be exploited for profit. Its application has revealed inefficiencies in various markets, providing insights into the behavior of financial assets and guiding further research on market anomalies.

3.2. Hurst exponent

The Hurst exponent, named after hydrologist Harold Edwin Hurst who introduced it in 1951, measures the long-term memory of time series data. It provides insights into the tendency of a time series to regress to the mean or persist in a certain direction. A Hurst exponent value of 0.5 indicates a random walk (no memory), values greater than 0.5 suggest a persistent behavior (trending), and values less than 0.5 indicate anti-persistent behavior (mean-reverting).

In the context of market efficiency, the Hurst exponent is a valuable tool for detecting long-range dependencies in financial time series. Its dynamic variation, where the Hurst exponent is calculated over moving windows, allows researchers to observe how market efficiency evolves over time. Studies utilizing the Hurst exponent have revealed periods of both efficiency and inefficiency in financial markets, challenging the static view of the EMH. By capturing the changing nature of market behavior, the Hurst exponent has become an essential metric in the study of market dynamics, contributing to a deeper understanding of the factors influencing asset prices and market stability.

3.3. Variance Ratio Test

The Variance Ratio Test, introduced by Lo and MacKinlay in 1988, is a pivotal tool for examining the random walk hypothesis in financial markets. This test compares the variance of multi-period returns to the variance of single-period returns to detect serial dependencies in the data. Under the random walk hypothesis, the variance of returns should increase linearly with the time interval, meaning the Variance Ratio should be one.

The significance of the Variance Ratio Test lies in its ability to uncover both mean-reversion and momentum effects, which are critical in evaluating market efficiency. Lo and MacKinlay's application of this test to U.S. stock market data provided compelling evidence against the random walk hypothesis, suggesting that stock returns exhibit serial correlation over certain time periods. This finding was instrumental in questioning the weak form of EMH and influenced numerous studies applying the Variance Ratio test across different markets and asset classes, from equities to cryptocurrencies.

3.4. Trading strategies

Momentum and contrarian strategies are two widely studied approaches in the realm of trading strategies. Momentum strategies are based on the premise that assets that have performed well in the past will continue to perform well in the future, and those that have performed poorly will continue to underperform. Traders using momentum strategies typically buy assets that have shown strong performance over a specific period and sell those that have shown weak performance, expecting the trends to persist. This approach leverages the inertia of price movements, capitalizing on the continuation of existing trends (Jegadeesh & Titman, 1993). On the other hand, contrarian strategies operate on the opposite principle, suggesting that markets tend to overreact to news, causing prices to deviate from their intrinsic values. Contrarian traders buy assets that have recently performed poorly, anticipating a price correction, and sell assets that have recently performed well, expecting them to revert to the mean. These strategies are grounded in the belief that market participants often exhibit herd behavior, leading to excessive optimism or pessimism that can be exploited for profit (De Bondt & Thaler, 1985). By understanding and applying momentum and contrarian strategies, traders aim to capitalize on the predictable patterns of asset price movements.

4. Empirical studies on Weak Form Efficiency

The literature on weak form efficiency is vast and varied, covering a wide range of financial markets around the world. We investigate the findings where different regions were studied but also the results of different types of research, whether in terms of style or asset class analyzed.

4.1. Studies on U.S. markets

The United States has been the center of numerous empirical studies on market efficiency, largely due to the availability of high-quality data and market depth. Fama (1970) conducted pioneering research on the U.S. stock market. In his study, he showed that stock returns follow a random walk, thus supporting the weak form efficiency hypothesis. He used correlation tests and autoregressive models to examine the independence of stock returns.

Lo and MacKinlay (1988) used correlation tests and the Variance Ratio Test to examine the independence of U.S. stock returns. They found evidence that there was positive autocorrelation of weekly and monthly returns, thus refuting weak form efficiency. In a 1997 study, they concluded “Our results show that predictable components are indeed present in the stock market, and that sophisticated forecasting models based on measures of economic conditions do have predictive power." (Lo, and MacKinlay, 1997, p31). They assert that the stock market not following a random walk is a well-established fact.

Malkiel (2003) reviewed several studies and provided evidence that U.S. stock markets are generally efficient, with prices reflecting all available information. Regarding predictable patterns and prices irregularities, he thinks that these can exist but not in the long term and that there are too few of them to disprove the EMH. He adds that a small degree of inefficiency is necessary to provide the incentive for market participants to collect and analyze the information and exploit those short-lived opportunities. This equilibrium allows for the market to stay mostly efficient. In his own words "whatever patterns or irrationalities in the pricing of individual stocks that have been discovered in a search of historical experience are unlikely to persist and will not provide investors with a method to obtain extraordinary returns. If any $100 bills are lying around the stock exchanges of the world, they will not be there for long." (Malkiel, 2003, p80).

4.2. Studies on European markets

European markets have also been the subject of numerous empirical studies, although results are often more varied due to differences between countries.

The study of DeLong and Becht (1992) on the German market revealed evidence of excess volatility in the post-World War II but an absence of excess volatility pre-World War I.

Cuthbertson and Hyde (2002) studied the German and French markets using the Campbell-Shiller VAR methodology and found evidence of excess volatility, but only when the model assumed a constant risk premium.

Maria Rosa Borges (2010) explored the weak-form efficiency of stock market indexes of UK, France, Germany, Spain, Greece and Portugal, from January 1993 to December 2007. She found mixed evidence, with some periods showing signs of inefficiency. The efficient hypothesis was rejected on daily data for Portugal and Greece. However, the empirical tests show random behavior after 2003. France and UK reject the null hypothesis, due to the presence of mean reversion in weekly data. Germany and Spain show efficiency, the latter being the most efficient.

Worthington and Higgs (2004) analyzed 20 European stock markets, of which 16 were developed and 4 emerging. The results, indicate out of the emerging markets only Hungary follows a random walk. In the developed markets only Germany, Ireland, Portugal, Sweden and the United Kingdom exhibit random walk behavior.

4.3. Studies on emerging markets

Emerging markets often exhibit different characteristics in terms of liquidity and information transparency, which can influence their level of efficiency.

Chaudhuri and Wu (2003) studied 17 emerging markets (Argentina, Brazil, Chile, Colombia, Greece, India, Jordan, Korea, Malaysia, Mexico, Nigeria, Pakistan, Philippines, Taiwan, Thailand, Venezuela, and Zimbabwe) and found that ten of them did not show signs of weak form efficiency, possibly due to lower competition and limited information availability.

Narayan and Smyth (2005) investigated the stock markets of 22 developed OECD countries, including some Asian markets, revealed evidence of weak form efficiency, although some inefficiencies persist in certain countries. Overall, results provided support for the random walk hypothesis.

Karemera and al (1999) tested weak-form efficiency for 15 emerging capital markets, including Argentina, Brazil, Chile, Hong Kong, Indonesia, Israel, Jordan, Korea, Malaysia, Mexico, Philippines, Singapore, Taiwan, Thailand, and Turkey. Most of them were found to respect the null hypothesis which suggests that "investors are unlikely to make systematic nonzero profit by using past information in many of the examined markets, thus, investors should predicate their investment strategies on the assumption of random walks." (Karemera, 1999, p171).

4.4. Research on the Belgian stock market

Specific research on the Belgian stock market is relatively limited, which further justifies the importance of this thesis. The lack of extensive published literature on this market highlights the need to study its efficiency.

In the study of Lee (1992) the random walk hypothesis was tested with the Variance Ratio for 10 industrialized countries, including Belgium for the period 1967-1988. Using weekly returns, Belgium was one of the only 2 countries along with Australia to reject the random walk hypothesis. However, the hypothesis was validated for all countries when the holding period return was lengthened from one week to four weeks.

Chan (1997) tested weak-form efficiency and cointegration of 18 nations studying the period between 1961 and 1992, the weak-form hypothesis was not rejected for the Belgian market.

In the previously mentioned Worthington and Higgs (2004) study which looked at data from 1987 to 2003 and employed multiple test such as serial correlation, unit root and multiple Variance Ratio tests, the Belgian market did not respect the random walk hypothesis. This was the case with all the tests used.

4.5. Studies on the cryptocurrency market

More recent studies of weak-form market efficiency have also been performed on the rapidly growing cryptocurrency market. Even though the asset class is different, the principles remain the same and the methodologies are therefore analogous.

Kang, Lee, and Park (2022) conducted a comprehensive study on the information efficiency of the cryptocurrency market. They employed three random walk tests to verify the weak-form EMH, namely the runs test, the Variance Ratio and an autocorrelation test named the Durbin-Watson test. Furthermore, they used the event study method to test the semi-strong-form EMH. Their findings indicated that only 6.04% of the 893 cryptocurrency units satisfied the weak-form EMH, and 2.695% met the semi-strong-form EMH. Additionally, they observed that larger exchanges established before November 2017 were more likely to exhibit both forms of efficiency (Kang, Lee, & Park, 2022).

Palamalai, Kumar, and Maity (2021) conducted a study on the weak-form efficiency of the top ten cryptocurrencies by applying non-parametric and parametric random walk testing methods robust to unknown structural breaks and asymmetric effects. The tests included the runs test, unit root tests like the ADF test, multiple Variance Ratio tests, among others. Their results did not support the random walk hypothesis, validating weak-form inefficiency for daily cryptocurrency returns. This inefficiency was attributed to the presence of asymmetric volatility clusters (Palamalai, Kumar, & Maity, 2021).

Lopez-Martin, Benito, and Arguedas (2021) carried out a comprehensive study on the efficiency of various cryptocurrencies, including Bitcoin, Litecoin, Ethereum, Ripple, Stellar, and Monero. They utilized five tests applied in both static and dynamic contexts, including the Variance Ratio test and the Hurst exponent. Their results indicated that the degree of market efficiency tends to increase over time. However, the efficiency of cryptocurrencies like Ripple, Stellar, and Monero fluctuated, showing alternating periods of efficiency and inefficiency, consistent with the Adaptive Market Hypothesis (Lopez-Martin, Benito, & Arguedas, 2021).

These studies collectively indicate that the cryptocurrency market, while exhibiting some characteristics of traditional financial markets, often deviates from the EMH, with efficiency varying significantly across different cryptocurrencies and over time.

4.6. Studies using trading strategies

As mentioned above another part of this literature looks at the potential profitability of trading strategies to assess a market's efficiency.

Technical analysis, which involves forecasting future price movements based on past price data, has been extensively tested in various markets to determine whether these strategies can consistently generate abnormal returns. Early studies often indicated limited profitability in stock markets, with pioneering works such as Fama and Blume (1966) finding that many technical trading rules, did not outperform a simple buy-and-hold strategy after accounting for transaction costs. However, later studies have added some nuance. For instance, Park and Irwin (2007) conducted a comprehensive review of technical trading strategies across different markets. Their review of 95 more modern studies revealed that 56 of these studies reported positive profitability of technical trading strategies, particularly in non-equity markets such as foreign exchange and futures markets. The review highlighted that while these strategies were profitable in stock markets through the late 1980s, their effectiveness has diminished thereafter. On the contrary, in foreign exchange markets, technical trading rules demonstrated consistent profitability over several decades, though their effectiveness has reportedly declined since the early 1990s. However, the authors also caution that many of these studies face significant methodological challenges, such as data snooping and issues with transaction cost estimation, which hurts the robustness of the findings.

In their study, Rad, Low, and Faff (2016) investigate the profitability of pairs trading strategies using three different methods: distance, cointegration, and copula techniques, applied to the U.S. equity market over a long period from 1962 to 2014. They find that while all three methods show significant excess returns before transaction costs, the performance diminishes notably after costs are considered, with the copula method particularly underperforming. These findings highlight the importance of considering transaction costs and practical implementation challenges in such strategies.

Gerritsen et al. (2020) examine the effectiveness of several technical trading rules in the Bitcoin market, using daily price data from 2010 to 2019. Their study finds that certain trading rules, particularly the trading range breakout rule, consistently outperform the buy-and-hold strategy, especially during periods of strong market trends. The study highlights the potential for technical analysis to generate excess returns in cryptocurrency markets, though it also underscores the importance of adapting strategies to the unique characteristics of this asset class.

In their study, Ni et al. (2022) examine the effectiveness of momentum and contrarian trading strategies within the Korean and Chinese stock markets. Their findings reveal that these markets are not fully efficient, as investors could exploit specific technical indicators to generate abnormal returns. The study demonstrates that contrarian strategies are particularly effective in the Korean market whereas in the Chinese market, momentum strategies yielded significant returns. These findings suggest that while technical trading strategies can exploit market inefficiencies, the choice of strategy depends heavily on the specific market conditions prevalent in different regions.

4.7. Implications of empirical results

The implications of empirical results on market efficiency are vast. For investors, market efficiency means that trading strategies based on technical or fundamental analysis are unlikely to generate abnormal returns. However, the presence of inefficiencies may offer opportunities for those who can exploit information not yet integrated into prices. The presence of inefficiencies in the Belgian market may offer arbitrage opportunities for investors able to identify and exploit these anomalies. Strategies based on fundamental analysis or the use of forecasting models may be more effective in a less followed market like Belgium.

For regulators, ensuring transparency and access to information is crucial to maintaining or improving market efficiency. Detected anomalies and inefficiencies can also inform policies aimed at improving market oversight and regulation. Measures to improve liquidity and access to information can help reduce inefficiencies in the Belgian market.

4.8. Purpose of this research

Out of the 3 studies that tested the weak-form efficiency of the Belgian market, 2 of them showed inefficiency, including the most recent one. The latter covered a period that stopped in 2003, that leaves us with a huge gap since, that this present thesis is looking to fill in.

5. Methodology

5.1. Data

For this study, the data was extracted from Yahoo Finance, covering the period from January 1, 2010, to July 19, 2024. The dataset includes the daily closing prices of the BEL 20 index, as well as the closing prices of all individual stocks currently constituting the BEL 20 index. Additionally, six stocks that were part of the BEL 20 in the past were included to ensure a comprehensive analysis. For 5 of the 26 individual stocks, the data didn't go as far back as January 2010 so the earliest start date available was chosen. These stocks are Argenx (01/08/2014), Azelis (01/10/2021), Syensqo (01/01/2024), Bpost (01/07/2013) and Ontex (01/07/2014). As a consequence, the Hurst exponent calculated with 1024 observations was not able to be calculated for Azelis and Syensqo due to lack of observations.

The chosen period, spanning over 14 years, provides a substantial sample size for robust analysis. This period was selected to focus on the modern financial era, characterized by technological advancements as well as changes in market conditions and regulatory landscapes which significantly evolved after the 2008 financial crisis, illustrated by Basel III and the Dodd-Frank Act. By starting from 2010, the data reflects a more stabilized and current financial environment, allowing for more relevant and accurate conclusions about today's market efficiency. Daily returns were calculated from this dataset by dividing the closing price of a given day by the closing price of the previous trading day and then subtracting one from the result.

5.2. Runs test

The runs test, proposed by Wald and Wolfowitz (1940), is a non-parametric statistical method used to test the randomness of a sequence of data points. In the context of financial markets, it is employed to determine whether the sequence of stock returns follows a random pattern, which is a key assumption of the weak form of market efficiency.

To perform the runs test, the first step involves transforming the daily return data into a binary sequence. Each return is classified as either a positive return or a negative return. This binary transformation simplifies the analysis by focusing on the direction of the returns rather than their magnitude.

Once the binary sequence is established, the next step is to identify and count the "runs" within the sequence. A run is defined as a consecutive series of the same type of returns (positive or negative). For example, a sequence of three positive returns followed by two negative returns and then four positive returns would contain three runs: one run of positive returns, one run of negative returns, and another run of positive returns, of lengths 3, 2 and 4.

The total number of runs observed in the data is then compared to the expected number of runs under the null hypothesis of randomness. The expected number of runs can be calculated using the following formula:

$$E(R) = \frac{2n_1n_2}{n_1+n_2} + 1$$

where n_1 is the number of positive returns, n_2 is the number of negative returns, and n is the total number of returns ($n = n_1 + n_2$)

To assess the statistical significance of the difference between the observed and expected number of runs, the standard deviation of the number of runs is calculated as:

$$\sigma_R = \sqrt{\frac{2n_1n_2(2n_1n_2 - n_1 - n_2)}{(n_1+n_2)^2(n_1+n_2-1)}}$$

The test statistic for the runs test is then computed as:

$$Z = \frac{R - E(R)}{\sigma_R}$$

where R is the observed number of runs. This Z-score follows a standard normal distribution under the null hypothesis of randomness. If the absolute value of the Z-score is greater than the critical value from the standard normal distribution, the null hypothesis is rejected, indicating that the returns are not random and that there may be some form of dependency in the data. If the null hypothesis is rejected because the number of observed runs is too high, it implies that the average length of runs is shorter than under random conditions and that there is a back and forth between positive and negative returns. In other words, it indicates that the series of returns is mean reverting. On the contrary, If the null hypothesis is rejected because the number of observed runs is too low, it implies longer runs and a persistent behavior of stock returns. In other words, the series of returns is trending.

5.3. Variance Ratio test

The property on which the Variance Ratio test relies is that in a series in which observations are independent and identically distributed, variance scales linearly. In other words, if a time series of financial daily returns is indeed independent and identically distributed, as it is asserted by the random walk hypothesis, the variance of k-day returns of the series should be equal to the variance of 1 day returns of that series, times k.

Therefore the core idea of the Variance Ratio Test, introduced by Lo and MacKinlay in 1988, is to compare the variance of returns over multiple periods (k-period returns) with the variance of single-period returns and see if it deviates from what is expected under the random walk assumption.

In this study, the test was performed for k-periods ranging from 2 to 64 days. The Variance Ratio (VR) for each k is computed using the following steps:

First, calculate the variance of daily returns ($\sigma^2(1)$).

Next, calculate the variance of k-period returns ($\sigma^2(k)$), where k-period returns are the sum of k consecutive daily returns.

Then, compute the Variance Ratio:

$$VR(k) = \frac{\sigma^2(k)}{k \cdot \sigma^2(1)}$$

The Variance Ratio is adjusted by subtracting 1 to center it around zero for easier interpretation.

The standard error for the Variance Ratio is then calculated as:

$$SE = \sqrt{\frac{2(2k-1)(k-1)}{3kN}}$$

where N is the total number of observations.

The test statistic Z is computed as:

$$Z = \frac{VR(k)-1}{SE}$$

Finally, the p-value is determined to assess the statistical significance of the test statistic which follows a normal distribution. This whole process is repeated for the 63 different k periods. Because we are conducting multiple tests, there's an increased chance of obtaining at least one small p-value purely by random chance.

To confirm that the lowest p-value we obtained out of the 63 tests was not due to chance, we apply the Chow-Denning test, a procedure proposed by Chow and Denning in 1993. The resulting p-value from the Chow-Denning test is derived from the distribution of the maximum absolute value Z-statistic we obtained, which adjusts for the fact that multiple tests are being conducted. It reflects the likelihood that the observed maximum absolute value Z-statistic could occur under the null hypothesis of no serial correlation, considering the multiple tests we conducted. The Chow-Denning statistic is calculated as:

$$\text{Chow-Denning} = 1 - (1 - \min \text{p-value})^{\max_t}$$

where min p-value is the smallest p-value obtained from the individual Variance Ratio tests, and $\max_t$ is the maximum k-period tested. If the Chow-Denning p-value is greater than 10%, we consider the market as efficient. If it is less than 10% and the maximum absolute value(abs) Z-statistic is positive, the market is trending. If the maximum abs Z-statistic is negative, the market is mean-reverting. This is because a positive max abs Z-statistic indicates that the Variance Ratio of a certain k-period is significantly higher than expected, which can be attributed to the trending behavior of the time series during those periods, which exacerbates returns. Conversely, when the Variance Ratio is significantly lower than expected, it suggests mean-reversion, where returns tend to revert to the mean over the given k-period, resulting in a smoothing effect on the returns.

For visualization, the results of the Variance Ratio test were plotted with the Z-statistic values on the y-axis and the different k periods on the x-axis. This allows for a clear visual representation of the time series' behavior over various time horizons. The plotted data helps to easily identify the specific k periods where the Variance Ratio deviates the most from what is expected under the random walk hypothesis.

5.4. Dynamic Hurst exponent

The Hurst exponent, developed by hydrologist Harold Edwin Hurst in 1951, is a statistical measure used to evaluate the long-term memory of time series data. This means that if the time series movements seem to be impacted by long-term past movements, they retain a "memory" of them. In the context of financial markets, it helps to determine whether stock prices exhibit a random walk, persistent behavior, or mean-reverting behavior. The Hurst exponent H can take values between 0 and 1. If H=0.5, it indicates a random walk, meaning that the series is completely random and past returns do not predict future returns. Values of H>0.5 indicate persistent behavior, suggesting that positive returns are likely to be followed by positive returns, and negative returns by negative returns, which can be seen in trending markets. Conversely, values of H<0.5 indicate mean-reverting behavior, where high returns are likely to be followed by low returns and vice versa, leading to a smoothing effect over time.

The dynamic Hurst exponent extends this concept by calculating the Hurst exponent over rolling windows of time, allowing for the analysis of how market efficiency evolves dynamically. Indeed, instead of calculating one Hurst exponent for the whole time series the dynamic variations allows to have Hurst exponents calculated for each rolling windows (a rolling window captures a continuous sequence of data points, and as the window moves, it drops the oldest data point and includes the next new point, maintaining the same window size) of a chosen length and we can thus see the its evolution through time over the broader period we are studying. This approach is particularly useful for detecting changes in market behavior and identifying periods of inefficiency. By observing the Hurst exponent over different periods, we can gain insights into the temporal stability of market conditions and the potential for exploiting these inefficiencies.

The procedure involves several steps. First, the daily returns are computed and segmented into rolling windows of a specified length, usually powers of 2. In this analysis 1024 trading days was chosen which is about 4 years, that is why the first Hurst exponent available for the time series is from 2014 onwards. Each window is analyzed separately to calculate the Hurst exponent. Within each window, the returns are further divided into subsamples of varying lengths, where the subsample length increases in powers of two. This hierarchical division allows for a detailed examination of the rescaled range across different time scales. For each subsample, the rescaled range is computed by first calculating the cumulative deviation from the mean, then finding the range (difference between the maximum and minimum cumulative deviations), and finally dividing by the standard deviation. The average rescaled range for each subsample length is then transformed using the logarithm base 2, which makes the relationship more linear, making it suitable for regression analysis.

A linear regression is performed between the logarithm of the rescaled range and the logarithm of the subsample lengths, with the slope of the regression line providing an estimate of the Hurst exponent. We also get the standard error of that estimate with the regression analysis. The t-statistic is then computed to test the null hypothesis that the Hurst exponent is equal to 0.5, indicating a random walk. The p-value is then determined to assess the statistical significance of the deviation from 0.5. The dynamic Hurst exponent can thus be plotted over time to visualize changes in price movements behavior, with periods where the Hurst exponent deviates significantly from 0.5 indicating potential inefficiencies.

6. Empirical results

6.1. Runs test

Results for the 27 assets are presented in Appendix I (see Appendix I: Runs test results). A significance level of 10% was used, and a two-tailed normal distribution was employed to calculate the p-values. The BEL 20 index failed to reject the null hypothesis of randomness with a p-value of 91.45% which highly suggests weak-form efficiency. Out of the 26 individual stocks, only 7 rejected the null hypothesis. Namely Aedifica, Cofinimmo, Lotus, Sofina, Syensqo, WDP and Colruyt. It is interesting to point out that for all of these stocks the Z score was positive, which implies that their number of runs was significantly higher than their expected number of runs, which indicates a mean reversion behavior of their daily returns. P-values were particularly low for Aedifica (0,11%), Cofinimmo (1,7%) and Lotus (0,09%).

The results of this runs test suggest that the Belgian stock market is overall mostly efficient. However some inefficiencies were found, especially for the 3 stocks with very low p-values.

6.2. Variance Ratio test

Results for the 27 assets are illustrated with graphs in Appendix II (see Appendix II: Variance Ratio test graphs). For visualization, the results of the Variance Ratio test were plotted with the Z-statistic values on the y-axis and the different k periods on the x-axis. This allows for a clear visual representation of the time series' behavior over various time horizons. The plotted data helps to easily identify the specific k periods where the Variance Ratio deviates the most from what is expected under the random walk hypothesis.

Since we use the Chow-Denning procedure, we focus only on the maximum absolute value Z-statistic (minimum p-value) for the hypothesis testing, for which the chosen significance level is 10%. All the p-values were determined using a two-tailed normal distribution.

The BEL 20 again failed to reject the null hypothesis, the Variance Ratio test thus reinforces the findings of the runs test suggesting weak-form efficiency for the index. Regarding individual stocks, 5 out of 26 rejected the null hypothesis. Namely KBC at the 2-day period, Lotus at 3, WDP at 39, Colruyt at 4, and Proximus at the 8-day period. The evidence from the Variance Ratio test confirms the majority of the conclusions from the runs test except for Aedifica, Cofinimmo, Sofina, Syensqo, KBC and Proximus. The Chow-Denning p-value for KBC was 7%, and 5,19% for Proximus. As for Lotus, WDP and Colruyt who rejected the null hypothesis yet again, they respectively had Chow-Denning p-values of 0.00000126%, 4,25% and 0.0035%. All the stocks that rejected the null hypothesis had a negative maximum absolute value Z-statistic, which suggests mean-reversion in the data, except for KBC who had a positive maximum absolute value Z-statistic which suggests a trending behavior of the data.

6.3. Dynamic Hurst exponent

Results for the 27 assets are illustrated with graphs in Appendix III (see Appendix III: Dynamic Hurst exponent graphs). Each asset has 2 graphs to illustrate the findings. On the first one, we can see the evolution of the dynamic Hurst exponent over the period starting from its first available computation, which is approximately 4 years after the first data point of the time series (2014 for most assets), until the end of the period studied which is in July 2024. The second graph complements the first one by showing the value of the corresponding t-statistic over the same period along with the positive and negative critical t-statistic thresholds for the two-tailed confidence intervals of 0.5%, in order to have a better visualization of its significance.

Results could not be obtained for Azelis and Syensqo due to the number of observations being too low. Since this is a dynamic result, a conclusion can not be given for the whole time series, instead we examine the overall behavior of the Hurst exponent over time. Looking at the graphs for the 25 assets we can note that only Ontex does not reject the null hypothesis of absence of long memory in the data. It is interesting to observe that for the other 24, the Hurst exponent surpasses 0.5 in a statistically significant manner at least at some point during the period. It indicates long memory in the data and more specifically trending behavior.

We can point out that the majority of those assets saw their Hurst exponent increase substantially around the start of 2020, suggesting that data from around the start of that year had trending patterns. This phenomenon might be explained by inefficiencies occurring during the covid shock. The highest Hurst exponents tend to be seen during the period 2020-2024, even though the BEL 20 index, Ageas, Cofinimmo, GBL, KBC, Lotus, Melexis, Proximus, and Bekaert showed some temporary high values during the 2014-2016 period, especially around 2015. High levels were also seen during the period 2016-2020 for Ackermans, Argenx, Lotus, Melexis, Solvay, UCB, Umicore, Proximus, Bpost and Bekaert.

The dynamic nature of this test allowed us to analyse this time series through a different lens. Alternating periods of inefficiency and efficiency were the most frequently observed finding, consistent with the Adaptive Market Hypothesis mentioned in the literature review.

7. Trading strategies based on findings.

As mentioned above, another way to test for efficiency and even in Samuelson's view point the only real way, is to investigate if certain trading strategies could generate abnormal returns. In an attempt to enrich our analysis, we will try to see if the inefficiencies discovered with our three statistical tests could help develop profitable trading strategies, which would confirm them.

Each strategy will be tested under conditions with and without short selling to evaluate their effectiveness across different scenarios to see if private investors, who generally do not have the ability to short stocks, could also benefit from the potential inefficiencies.

7.1. Runs test strategies

Regarding the runs test, inefficiencies suggesting mean-reverting behavior were found for 7 stocks. We will thus develop and test 2 trading strategies that aim to profit from the mean-reversion. The first strategy is a threshold-based mean-reversion strategy. This strategy relies on the expectation that sequences of positive/negative returns tend not to last long and to be followed by a movement in the opposite direction, suggested by the findings of the runs test. For each stock, trades were executed when the previous x days' cumulative return exceeded or fell below a specified threshold, which was set at 0.5, 1, 1.5 and 2 standard deviations of the stock's x-day returns. If the return exceeded the positive threshold, the strategy consisted in short selling the stock the following day to profit from the expected reversion. On the other hand, if the return fell below the negative threshold, the strategy involved buying the stock the next day. Introducing a threshold helps to filter out minor fluctuations that are less likely to represent true mean-reversion opportunities. Regarding the length of the x-day periods, 1-, 2-, 3-, 4- and 5-day periods were analyzed. Longer periods would be less appropriate because they would greatly reduce the number of trades and miss a lot of potential reversal opportunities. Furthermore, long trends would not be expected to be seen often in data that exhibits mean-reversion.

The second strategy focuses on consecutive day trends to identify potential mean-reversion opportunities. This strategy hypothesizes that sequences of consecutive gains or losses might signal upcoming reversals in price movements. Specifically, the strategy tests periods of 1, 2, 3, 4 and 5 consecutive days of positive or negative returns. After identifying such sequences, the strategy involves short selling the stock following a sequence of consecutive gains, anticipating a reversal, or buying the stock after consecutive losses, expecting a rebound.

7.2. Variance Ratio test strategies

As for the Variance Ratio test, 4 out of the 5 stocks that rejected the null hypothesis had Variance Ratios far below the expected value which suggests mean-reversion in the particular k-period concerned. Therefore, we developed a trading strategy that aims to capitalize on the mean-reverting behavior over the critical k-period for those 4 stocks as well as a version of it that exploits trending behavior over the k-period concerned for the fifth stock. Similar to the previous strategies, each implementation was tested both with and without short selling.

The strategy involves splitting the identified k-period into two parts: an observation period (the first part) and a trading period (the second part). For instance, if the k-period is 8 days, possible splits might include a 3-5 split, where the first 3 days are used to observe the cumulative returns that occurred over those first 3 days, and the next 5 days are used to execute trades. Depending on the observed cumulative return during the first part, a trading decision is made for the second part. The multiple splits tested for each asset maintain a balance between the first and second periods. Indeed, if the first part is too short relative to the entire k-period, it might not provide enough information to reliably indicate a trend that is likely to persist/revert over the rest of the period. Therefore, it would increase the risk of trading based on false signals. On the other hand, if the first part of the split is too long relative to the k-period, a significant portion of the trend/reversion might already have played out, which would diminish the predictive power of that first period. This is especially true for the trending strategy and less so for the mean-reversion strategy. However, it is also not ideal for the latter to have a very short second part because it limits the strategy to brief exposures and prevents it from profiting from shorter yet still well-established trends that are likely to revert.

In the trending strategy, the cumulative returns calculated during the first part of the split are then compared against predefined thresholds, based on the standard deviation of the cumulative returns of the same length as the first part of the split over the entire time series. In our example of the 3-5 split the standard deviation used for the thresholds is the standard deviation of cumulative 3-day returns for the asset over the whole time series. Adding a threshold for the cumulative return before taking a position helps filter out minor fluctuations that might not indicate a strong trend, thereby reducing the likelihood of false signals. The thresholds used in this analysis are 0.5, 1, 1.5, and 2 standard deviations. If the cumulative return exceeds the positive threshold, it is assumed that the trend will continue over the k-period more often than not, leading to this higher Variance Ratio we observed in the results. Therefore in that case, we buy the asset to get the returns of the second part of the split. On the contrary, if the cumulative return is below the negative threshold, it is assumed that the trend will continue downwards, so we short the asset for the second part of the split.

For the mean-reversion strategy, cumulative returns are similarly calculated during the first part of the split, and the same set of thresholds (0.5, 1, 1.5, and 2 standard deviations) are used to determine the significance of the observed cumulative returns. If the cumulative return exceeds the positive threshold, it is assumed that the stock is likely to revert in the second part of the split, leading to the lower Variance Ratio observed in the results. In this case, we would short the asset for the second part of the split to benefit from its fall. When the cumulative return is below the negative threshold, it is also assumed that the stock will revert, so we buy the asset for the second part of the split to profit from its rise.

7.3. Dynamic Hurst exponent strategies

And lastly the dynamic Hurst exponent. As mentioned above, the results showed that all assets which had sufficient data to be analyzed showed a statistically significantly higher than 0.5 hurst exponent at some point during the period except for one. It suggests that these assets had trending behavior at some point that could potentially be exploited. The idea is then to identify the periods where the data is most inefficient and implement a strategy only around those. Looking at the graphs of the Hurst exponent helps us see at which point was it the highest and how it evolved over time, however just looking at the highest point and then implementing the strategy on the whole window with which that hurst exponent was calculated is not the most relevant way of doing it and I'm going to explain why. Let's first remember that we calculated the Hurst exponents based on rolling windows of 1024 days which is about 4 years of the time series' data. Imagine that for a given stock the Hurst exponent in 2016 had a value of 0.5. It would suggest that the data with which it was calculated, being 2012 to 2016 approximately, shows no long memory. Now let's say the line plotted on the graph representing the Hurst exponent stays flat at 0.5 until 2018. Keeping in mind that since it's calculated based on a 4-year rolling window, as the window moves forward, the chronologically next new data point is added at the end while the oldest data point is simultaneously removed from the beginning, keeping the length of the window constant. Knowing that, we can infer that the data from 2016 to 2018 has the same behavior as the data from 2012 to 2014. That is because the Hurst exponent of 2016 is the same as the one from 2018 and the slope of the line of the Hurst exponent was 0 in between these 2 points so the exponent hasn't moved. This is only possible if the next new data that is continually being added as the 4-year window "rolls" between 2016 and 2018 is very similar in terms of behavior regarding long memory to the oldest data that is being replaced by it.

Indeed, the output not changing over time indicates that the new input and the old one it is replacing are very similar, in this case in terms of behavior, because that is the characteristic of the data that impacts the result. Otherwise, it would change the output. Old data being replaced in the input by more trending/less mean-reverting data would lead to an increase in the Hurst exponent. On the contrary old data being replaced by more mean reverting/less trending data would lead to a decrease in the Hurst exponent. Therefore, only old data being replaced by new data of similar behavior would result in the Hurst exponent not moving. Now let's imagine that the same graph shows a significant and linear rise of the Hurst exponent going from 0.5 in 2018 to 0.65 in 2020 and then staying at 0.65 during the following 2 years. The continued and significant rise from 2018 to 2020 suggests that it is the new data being added during this period that is causing the rise of the Hurst exponent meaning that the data containing trending behavior is probably precisely this period from 2018 to 2020. The value of 0.65 was thus potentially obtained with 2 years of efficient data (2016-2018) and 2 years of inefficient, trending data (2018-2020). The fact that the line stays flat at 0.65 during the following 2 years (2020-2022) indicates that the data from this period is very similar to that of the 2016-2018 period. As mentioned just above, that is because as the windows rolls during the 2-year period in which the Hurst exponent stays flat (2020-2022) it incrementally excludes the data from 2016-2018, which we consider efficient, and incorporates the data from 2020-2022. The fact that the Hurst exponent stays flat while this rolling occurs proves that the data from 2020-2022 has the same behavior as the data from 2016-2018. Therefore, we can consider in this example that the 2020-2022 period is also efficient.

This analysis allows us to find the specific period containing inefficiencies and thus to determine exactly where we should focus the implementation of our trading strategies, which in this instance would be the 2018-2020 period. Simply searching for the highest points of the graph and implementing the trading strategies in its 4-year window would have resulted in picking a point between 2020 and 2022 and so testing the strategies in a 4-year period between 2016 and 2022. Whatever the 4-year period chosen, it would have most likely contained 2 years of trending data and 2 years of random data, which would have been a suboptimal choice for a period. Therefore, carefully analyzing the evolution and the slope of the plotted Hurst exponent, trying to find the specific period causing certain deviations in the Hurst exponent is a more relevant way of identifying periods of interest.

The idea that a Hurst exponent not moving during a certain period implies that the new data incorporated into the rolling window over this period has the same behavior as the old data that has been removed from the rolling window over this same period, is a perfectly sound syllogism. Therefore, our deduction in the example above that the 2012-2014, 2016-2018 and 2020-2022 periods have the same behavior regarding long memory is correct. However, there is a big caveat to make. There is a good reason why I use the words “suggest” “probably” “consider” “potentially” when I say that these periods have an efficient behavior with an absence of long memory. This is because I postulated in my example that the 0.5 Hurst exponent in 2016 was due to the entire 4-year period of 2012-2016 having random behavior and absence of long memory. Although this is the typical interpretation, it is not a certainty. Indeed, it is theoretically possible for the Hurst exponent to be close to 0.5 even with data within the rolling window having periods of trending behavior followed by periods of mean-reverting behavior. For example, we could imagine the 2012-2014 period having trending behavior, followed by the 2014-2016 period having mean-reverting behavior. If these opposing behaviors within the window balance each other out, they could result in an overall Hurst exponent close to 0.5, which would misleadingly suggest randomness.

Consequently, there are multiple possibilities consistent with the shape of the graph in our example. Instead of having the 3 similar periods (2012-2014, 2016-2018 and 2020-2022) all being efficient, the 2014-2016 also being efficient and the 2018-2020 having trending behavior, we could have exactly the same graph with those 3 similar periods all having trending behavior, the 2014-2016 period having mean reverting behavior, and the 2018-2020 period having random behavior. The 2012-2014 trending period could balance out with the 2014-2016 mean-reverting period to result in the 0.5 value we observe in 2016. Afterwards, the 2012-2014 period being replaced by the similarly trending 2016-2018 as the windows rolls would lead to no significant change in the Hurst exponent, staying near 0.5 until 2018. Then, the mean-reverting 2014-2016 period being replaced by the random 2018-2020 period would result in the increase we see in 2020 up to 0.65. Finally, the trending 2016-2018 period being replaced by the similarly trending 2020-2022 period in the input would result in the Hurst exponent not changing up to 2022.

We’ve seen that different scenarios can lead to the same graph, which complicates the determination of the right period to implement our trading strategies. Indeed, in the example just above in the first scenario the inefficient period is the 2018-2020 period. However, in the second scenario based on the exact same graph the inefficient period is every period but 2018-2020! which is quite ironic.

To fix this problem, we will add granularity to our analysis by performing the same dynamic Hurst exponent calculations but with a much shorter window of 256 days, which is very close to the length of a trading year. This approach complements our initial analysis and helps confirm our suspicions about specific periods that may show long memory. This shorter window is more sensitive to shorter-term trends and changes, allowing us to cross-verify the findings of the 1024-day window. If a period shows long memory in both windows, it represents a stronger signal and would therefore be adequate to test our trading strategy. The graphs displaying the 256-day window Hurst exponents can be observed in Appendix IV (see Appendix IV: 256-day window Hurst exponents).

Going even further in terms of granularity would have downsides because there is a trade-off when choosing a window length to calculate Hurst exponents (Cajueiro and Tabak, 2004).

Indeed, longer windows have less sensitivity to shorter terms and more recent changes of trends as well as less granularity but are more reliable because of the sheer amount of data that leads to more robust conclusions. In contrast, shorter windows allow to see trends in more details and have more flexibility but are more susceptible to being impacted by short term noise that could lead to misleading conclusions. Having the cross-analysis with both 1024-day and 256-day windows allows us to have the best of both worlds and going even further would expose the analysis to too much noise for no real upside since the 256 days window already adds plenty of granularity.

For each asset, the period(s) that demonstrate(s) the strongest presence of long memory are identified through this cross-verification with both dynamic Hurst exponent calculations and are our periods of choice for the implementation of our trading strategies. Since all of those critical periods have Hurst exponents significantly greater than 0.5, trending strategies will be used.

Both strategies are similar to the two runs test strategies, except that they are adapted to trending data. For the first one, we look at the past x days' cumulative return to identify beginnings of trends. If that return exceeds a positive threshold which was calculated as z times the standard deviation of the entire period's x-day cumulative returns, a long position is initiated to get the next day's return. If the past x days' cumulative return falls below the negative version of the threshold, a short position is initiated to profit from the next day's potential negative return. Various values of x and z were tested to determine the most effective parameters. Namely 2, 3, 4, 5, 10, 20 and 50 days. This mix of short and long term periods is explained by the fact that the Hurst exponent is intrinsically a general assessment of the behavior of the data. Indeed, the value of the Hurst exponent does not pinpoint a specific time window in which the time series is more inefficient, it rather gives us the general tendency by taking into account sub samples of different lengths in its calculation. Therefore, a varied mix of values of x was chosen to ensure trends of different lengths were spotted and exploited. As it pertains to values of z, 0.25, 0.5 and 1 were chosen. These values are lower than for the previous strategies of mean-reversion because of the very nature of the behavior. Indeed, more extreme movements can be good signals of reversals in mean-reverting data. However, in trending data the goal is to enter the trend as early as possible and not get in when the movement is already at an extreme point, more than likely nearing its end. Consequently, values of z are intentionally lower.

The second strategy relies on the occurrence of consecutive days of positive or negative returns as a signal to enter a trade. The idea is to take a long position the days after x consecutive days of positive returns have been observed, and to short after x consecutive days of negative returns. Number of days from 1 to 5 were tested. Going beyond 5 would start to significantly diminish the number of opportunities and has diminishing marginal returns in terms of trend detection.

8. Performance of trading strategies.

In order to assess whether a strategy has generated returns indicating that it indeed has been able to capitalize on inefficiencies, the results from each combination of the strategies' parameters are displayed in tables and are compared to a buy and hold strategy over the same period. A buy and hold strategy consists in buying the asset at the start of the period and holding it until the end of it. For every table, combinations that outperformed the buy and hold strategy are highlighted in light green, with the highest performance being colored in a darker green.

8.1. Runs test strategies

Results for both strategies are shown in Appendices V and VI (see Appendix V: First runs test strategy performance) (see Appendix VI: Second runs test strategy performance). The first strategy performed better than the buy and hold strategy for 4 out of the 7 stocks concerned. Out of 20 different combinations of parameters tested for each stock (number of days and standard deviations thresholds), 5, 7, 20 and 12 outperformed the benchmark (buy and hold) for Aedifica, Cofinimmo, Syensqo and Colruyt respectively. The most notable excess returns were observed with Aedifica (425%) and Colruyt (383,58%) although those with Cofinimmo (115,81%) and Syensqo (26,2% over less than a year) were significant as well.

Without short selling, only 2 stocks were able to beat their benchmark with the strategy. Namely Colruyt (177,71%) and Syensqo (12,36%).

The second strategy also saw outperformances. Surprisingly, the without short selling variation was able to generate excess returns with 4 stocks but only 3 with the short selling variation. All of these were the same that were able to be exploited with the first strategy. The short selling variation had the highest outperformance for Cofinimmo (225,42%) and Syensqo (36,7%), whereas the other variation generated the highest excess returns for Aedifica (30,56%) and Colruyt (115,33%). Shorter number of consecutive days were responsible for most of the strategy's success.

We can safely say that these results from the runs test's strategies confirm the presence of inefficiencies by proving these can be exploited for profit for certain stocks at least. Analyzing the past days cumulative return and having thresholds seems to yield better results than using consecutive days to detect reversal opportunities, highlighted by the greater success of the first strategy.

8.2. Variance Ratio test strategies

Results for each stock can be seen in Appendix VII (see Appendix VII: Variance Ratio test strategy performance). We can see that for 3 out of 5 stocks, outperformance was achieved. All of the highest excess returns were obtained with the presence of short selling and were substantial. Colruyt is yet again part of the list with an outperformance reaching 425,04%. As for Proximus and KBC, 381,93% and 239,82% was achieved respectively in excess of their benchmark. However, the gains obtained without short selling were not something to dismiss. 200,63% for KBC, 302,36% for Colruyt and 241,58% for Proximus.

This strategy strongly confirms some of the inefficiencies detected by the Variance Ratio test as demonstrated by the high excess return achieved for the 3 stocks. The consistent outperformance of the variation with short selling compared to the already strong performance of the version without it, suggests that the logic behind the strategy is the main reason behind its success rather than being a mere statistical artifact.

8.3. Dynamic Hurst exponent strategies

Results of both strategies are displayed in Appendix VIII and IX (see Appendix VIII: First dynamic Hurst exponent strategy performance) (see Appendix IX: Second dynamic Hurst exponent strategy performance). There are 35 lines in these tables because some stocks had multiple periods tested. Each period is mentioned next to the stock's ticker in the first columns. The buy and hold return for each line is calculated based on these periods precisely. Out of those 35 periods, returns in excess of the buy and hold benchmark were observed for 26 periods with the short selling variation of the first trading strategy, and for 27 periods for the variation without short selling. The most substantial outperformances were seen with Bpost (238,98%, 2018-2022), Barco (222,55%, 2019-2024), Ageas (165,35%, 2020-2022), KBC (163,72%, 2011-2013), Galapagos (152,46%, 2019-2021), Argenx (134,02%, 2017-2021), Aedifica (114,96%, 2020-2024), Sofina (114,69%, 2020-2024), Bekaert (111,03%, 2020-2024), AB Inbev (102,79%, 2017-2020), UCB (92,33%, 2020-2024) and the BEL20 index (69,95%, 2019-2024). Some notable excess returns were also observed over shorter time periods with Colruyt in the year 2019 (51,94%), UCB in 2016 (34,82%), Melexis in 2023 (20,24%) as well as Umicore in the 2014-2016 period (57,43%). All of the very strong performances came from the short selling variation with the exception of Argenx. Moreover, 24 out of the 35 periods obtained their highest return with the short selling variation, strengthening the validity of the strategy.

As it pertains to the second strategy, excess returns were obtained in 16 out of the 35 periods with the short selling variation but the other variation saw 19 periods showing outperformance. Furthermore, the highest return for each of the 35 periods came mainly from the without short selling variation with 21 periods. This suggests that this strategy was not working as well, also highlighted by the lower excess returns generated by it in comparison with the other strategy. Nonetheless, the benchmark was substantially surpassed by KBC (536,51%, 2011-2013), Aedifica (160,03%, 2020-2024), Bpost (78,58%, 2018-2022), UCB (74,52% in 2016), Galapagos (54,54%, 2019-2021), Umicore (39,02%, 2018-2020) and Proximus (19,27%, 2023). It is worth noting that, with the exception of Bpost, all of these significant outperformances were obtained with the short selling variation even though the other variation worked more frequently overall, indicating that while the strategy worked less often, when it did it was particularly effective.

These strategies using the dynamic Hurst exponent also showed success in exploiting the inefficiencies suggested by the test results, for a significant portion of the assets analyzed. The cross-verification with both windows for the Hurst exponent yielded some good insights, proved by the success of the strategies. The first one performed especially well, confirming the conclusions of the runs test trading strategies suggesting that trading based on the past x-day cumulative return was a better opportunity detecting method than looking at the past x consecutive positive/negative day.

9. Conclusion

The aim of this thesis was to investigate the weak-form efficiency of the Belgian stock market by analyzing whether stock prices follow a random walk and if any inefficiencies could be exploited for profit. The study covered the period from January 2010 to July 2024, capturing the modern financial era. The first objective was to assess whether past returns of our dataset consisting of the BEL 20 index and its constituents, could predict future returns by deviating from a random process.

For that purpose, we conducted three distinct statistical test : the runs test, Variance Ratio test and the dynamic Hurst exponent, to detect possible inefficiencies. The results of these tests provided mixed evidence. While the overall market (as represented by the BEL 20 index) mostly showed signs of efficiency, several individual stocks displayed behaviors inconsistent with the random walk hypothesis, suggesting the presence of exploitable inefficiencies. The runs test found mean-reverting behavior in the data of Aedifica, Cofinimmo, Lotus, Sofina, Syensqo, WDP and Colruyt, while the other 20 assets did not reject the null hypothesis of randomness. The Variance Ratio test confirmed the findings regarding Lotus, WDP and Colruyt. It also detected mean-reversion in Proximus' returns, and trending behavior for KBC. The dynamic Hurst exponent suggested periods of trending behavior for all assets analyzed, with the exception of Ontex.

To further validate these findings, we developed and tested multiple trading strategies based on the detected inefficiencies. These strategies, designed to capitalize on either mean-reverting or trending behaviors, were applied to the relevant assets in the corresponding periods. Strategies were implemented both with and without the possibility of short selling, to assess whether private investors could also exploit the inefficiencies found. The results were compelling, with most strategies outperforming the simple buy-and-hold benchmark for a significant portion of the assets concerned. Those incorporating short selling usually saw the biggest gains however, outperformances were often seen without it as well. This suggests that certain inefficiencies in the Belgian stock market can indeed be exploited for profit, both for institutional and private investors. Consequently, those findings challenge the weak-form efficiency of this market.

However, this study is not without its limitations. While the trading strategies developed were usually successful, their applicability in the real world may be limited. Indeed, transactions costs were not taken into account in this study. The presence of trading costs could potentially diminish the excess returns obtained with these strategies in a substantial manner. Moreover, the results observed were dependent on the assumption that there were no liquidity constraints when trading these assets, which is not a guarantee in a smaller market like the Belgian stock market. Lastly, even though the results from the statistical tests and the trading strategies show that inefficiencies were present and able to be exploited, we exploited them with hindsight. Indeed, the data we used to detect those inefficiencies was not available during the period in which our strategies were implemented. Therefore, while these findings represent strong evidence that inefficiencies might be exploitable in the Belgian stock market, finding success implementing strategies in real time would be needed to make it a certainty.

Looking forward, future research could build on this study by incorporating limiting factors such as transaction costs and implementing some strategies in real time. Moreover, incorporating additional statistical tools and modern machine learning techniques could offer deeper insights regarding the patterns present in this market. Expanding the scope to include event studies to assess the semi-strong form of efficiency could also provide value to have a more detailed understanding of the informational efficiency of the Belgian stock market.

References

- Bachelier, L. (1900). Théorie de la spéculation. *Annales scientifiques de l'École Normale Supérieure*, 17, 21-86.
- Cajueiro, D. O., & Tabak, B. M. (2004). The Hurst exponent over time: testing the assertion that emerging markets are becoming more efficient. *Physica A: statistical mechanics and its applications*, *336*(3-4), 521-537.
- Chan, K. C., Gup, B. E., & Pan, M. S. (1997). International stock market efficiency and integration: A study of eighteen nations. Journal of business finance & accounting, 24(6), 803-813.
- Chaudhuri, K., & Wu, Y. (2003). Random walk versus breaking trend in stock prices: Evidence from emerging markets. *Journal of Banking & Finance*, 27(4), 575-592.
- Cowles, A. (1933). Can stock market forecasters forecast? *Econometrica*, 1(3), 309-324.
- Cowles, A., & Jones, H. E. (1937). Some a posteriori probabilities in stock market action. *Econometrica*, 5(3), 280-294.
- Cuthbertson, K., & Hyde, S. (2002). Excess volatility and predictability of stock returns in the United Kingdom and Germany. *Economics Letters*, 75(3), 340-345.
- De Bondt, W. F. M., & Thaler, R. (1985). Does the stock market overreact? *The Journal of Finance, 40*(3), 793-805. https://doi.org/10.1111/j.1540-6261.1985.tb05004.x
- DeLong, J. B., & Becht, M. (1992). Excess volatility in the German stock market, 1876-1990. *The Quarterly Journal of Economics*, 107(4), 1251-1289.
- Fama, E. F., & Blume, M. E. (1966). Filter rules and stock-market trading. *The Journal of Business*, *39*(1), 226-241.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383-417.
- Gerritsen, D. F., Bouri, E., Ramezanifar, E., & Roubaud, D. (2020). The profitability of technical trading rules in the Bitcoin market. *Finance Research Letters*, *34*, 101263.
- Granger, C. W. J., & Morgenstern, O. (1963). Spectral analysis of New York stock market prices. *Kyklos*, 16(1), 1-27.
- Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers, 116*(1), 770-799.
- Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance, 48*(1), 65-91. https://doi.org/10.1111/j.1540-6261.1993.tb04702.x

- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-291. https://doi.org/10.2307/1914185
- Kang, H. J., Lee, S. G., & Park, S. Y. (2022). Information efficiency in the cryptocurrency market: The efficient-market hypothesis. Journal of Computer Information Systems, 62(3), 622-631.
- Karemera, D., Ojah, K., & Cole, J. A. (1999). Random walks and market efficiency tests: Evidence from emerging equity markets. Review of Quantitative finance and Accounting, 13, 171-188.
- Lee, U. (1992). Do stock prices follow random walk?:: Some international evidence. International Review of Economics & Finance, 1(4), 315-327.
- Lo, A. W. (2004). The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective. The Journal of Portfolio Management, 30(5), 15-29.
- Lo, A. W. (2005). Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis. The Journal of Investment Consulting, 7(2), 21-44.
- Lo, A. W., & MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *The Review of Financial Studies*, 1(1), 41-66.
- Lo, A. W., & MacKinlay, A. C. (1997). Maximizing predictability in the stock and bond markets. *National Bureau of Economic Research*.
- López-Martín, C., Benito Muela, S., & Arguedas, R. (2021). Efficiency in cryptocurrency markets: New evidence. Eurasian Economic Review, 11(3), 403-431.
- Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1), 59-82.
- Mandelbrot, B. (1966). Forecasts of future prices, unbiased markets, and martingale models. *Journal of Business*, 39(1), 242-255.
- Maria Rosa Borges (2010). Efficient market hypothesis in European stock markets. *European Journal of Finance*, 16:7, 711-726.
- Narayan, P. K., & Smyth, R. (2005). Are OECD stock prices characterized by a random walk? Evidence from sequential trend break and panel data models. *Applied Financial Economics*, 15(8), 547-556.
- Ni, Y., Day, M. Y., Cheng, Y., & Huang, P. (2022). Can investors profit by utilizing technical trading strategies? Evidence from the Korean and Chinese stock markets. *Financial Innovation*, *8*(1), 54.
- Palamalai, S., Kumar, K. K., & Maity, B. (2021). Testing the random walk hypothesis for leading cryptocurrencies. Borsa Istanbul Review, 21(3), 256-268.
- Park, C. H., & Irwin, S. H. (2007). What do we know about the profitability of technical analysis?. *Journal of Economic surveys*, *21*(4), 786-826.

- Rad, H., Low, R. K. Y., & Faff, R. (2016). The profitability of pairs trading strategies: distance, cointegration and copula methods. *Quantitative Finance*, *16*(10), 1541-1558.
- Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6(2), 41-49.
- Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *The American Economic Review, 71*(3), 421-436.
- Wald, A., & Wolfowitz, J. (1940). On a test whether two samples are from the same population. *Annals of Mathematical Statistics*, 11(2), 147-162.
- Working, H. (1934). A random-difference series for use in the analysis of time series. *Journal of the American Statistical Association*, 29(185A), 11-24.
- Worthington, A., & Higgs, H. (2004). Random walks and market efficiency in European equity markets. The Global Journal of Finance and Economics, 1(1), 59-7