

# LUISS



Degree Program in *Economia e Management*

Course of *Financial Market Analysis*

**EVALUATING THE IMPACT OF ARTIFICIAL  
INTELLIGENCE INVESTMENTS ON STOCK PRICES:  
A Study of Microsoft and Nvidia within the framework of  
Efficient Market Theory**

Prof. Giovanni Rillo

SUPERVISOR

Giorgia Glorio, 271701

CANDIDATE

Academic Year 2023/2024

# TABLE OF CONTENTS

<b>LIST OF TABLES AND FIGURES</b> .....	3
<b>INTRODUCTION</b> .....	4
<b>CHAPTER 1 Understanding Artificial Intelligence: the value of AI investments</b> .....	5
1.1 Definition and history of Artificial intelligence.....	5
1.1.1 The state of AI.....	7
1.1.2 The impact of AI.....	9
1.2 AI Technologies.....	10
1.3 AI adoption at a corporate level.....	13
1.4 AI investment effects on firm value.....	16
1.5 AI potential to boost stocks .....	18
<b>CHAPTER 2 Theoretical Framework</b> .....	20
2.1 Efficient Markets Hypothesis .....	20
2.1.1 Definition.....	20
2.1.2 The Fair Game and Random Walk Theory.....	21
2.1.3 Consistency between market circumstances and efficiency.....	23
2.1.4 Different types of Efficiency Forms .....	24
2.2 Criticism .....	31
<b>CHAPTER 3 Field research</b> .....	33
3.1 Introduction to Nvidia and Microsoft.....	33
3.2 The Study Methodology .....	35
3.3 Statistical Method .....	38
3.4 Results.....	44
3.4.1 Nvidia.....	44
3.4.2 Microsoft .....	49
3.5 Discussion and limitations .....	59
<b>CONCLUSIONS</b> .....	62
<b>APPENDIX</b> .....	63
<b>REFERENCES</b> .....	64

# LIST OF TABLES AND FIGURES

## *Tables*

<b>Table 1</b> .....	45
<b>Table 2</b> .....	45
<b>Table 3</b> .....	50
<b>Table 4</b> .....	50
<b>Table 5</b> .....	54
<b>Table 6</b> .....	55

## *Figures*

<b>Figure 1 - Standardized Normal Distribution</b> .....	43
<b>Figure 2 - Abnormal Return (event window), Nasdaq Composite</b> .....	46
<b>Figure 3 - Cumulative Abnormal Return (event window), Nasdaq Composite</b> .....	47
<b>Figure 4 - Abnormal Return (event window), S&amp;P 500</b> .....	48
<b>Figure 5 - Cumulative Abnormal Return (event window), Nvidia, S&amp;P 500</b> .....	49
<b>Figure 6 - Abnormal Return (event window), Microsoft, Event 1, Nasdaq Composite</b> .....	51
<b>Figure 7 - Cumulative Abnormal Return (event window), Microsoft, Event 1, Nasdaq Composite</b> .....	52
<b>Figure 8 - Abnormal Return (event window), Microsoft, Event 1, S&amp;P 500</b> .....	53
<b>Figure 9 - Cumulative Abnormal Return (event window), Microsoft, Event 1, S&amp;P 500</b> .....	53
<b>Figure 10 - Abnormal Return (event window), Microsoft, Event 2, Nasdaq Composite</b> .....	56
<b>Figure 11 - Cumulative Abnormal Return (event window), Microsoft, Event 2, Nasdaq Composite</b> .....	57
<b>Figure 12 - Abnormal Return (event window), Microsoft, Event 2, S&amp;P 500</b> .....	58
<b>Figure 13 - Cumulative Abnormal Return (event window), Microsoft, Event 2, S&amp;P 500</b> .....	58
<b>Figure 14 - Cumulative Abnormal Return (event window), Nasdaq Composite</b> .....	60
<b>Figure 15 - Cumulative Abnormal Return (event window), S&amp;P 500</b> .....	60

All tables and figures presented in this thesis are original works.

# INTRODUCTION

In the business context, the importance of Artificial Intelligence is increasingly discussed, and there are few companies left that have not incorporated this technology into their business processes. Such rapid growth of this technology is likely because companies today are investing large sums of money in AI technologies to improve operations, develop innovative products, and increase customer satisfaction. Several studies have already attempted to measure the effects of AI on business performance, both individually and globally, but studies investigating how the market reacts to AI investments or how to quantify, in some way, the value of AI are limited. Therefore, the purpose of this thesis is to try to fill this gap in research and to propose a study that may be relevant to future research.

Underlying this study is the efficient markets theory (EMH), which serves as the theoretical framework for understanding how information influences stock prices and it is discussed in the second chapter. According to this theory, stock prices already incorporate and reflect all available information, preventing any possibility of obtaining abnormal returns by exploiting information in the public domain. However, the market's reaction to AI investment announcements raises questions about the ability of markets to quickly assimilate and reflect this information in stock prices.

This thesis aims to explore AI investment announcements' impact on stock prices by examining the market reaction and comparing it with theoretical predictions to test its efficiency. In this regard, after a theoretical discussion in the first two chapters, the conduct and results of three event studies concerning AI investment announcements by Microsoft and Nvidia are reported, which help us answer our questions and try to bridge the gap in research.

# CHAPTER 1

## Understanding Artificial Intelligence: the value of AI investments

---

The introductory chapter of this thesis discusses the birth of Artificial Intelligence, its technologies, and its progress to date, highlighting the impact this technology has on various industries worldwide. Next, the chapter examines the effects of investments in Artificial Intelligence and other technologies on corporate value, drawing on empirical and theoretical evidence. Finally, the potential relationship between AI and financial markets is analyzed, answering the question: *can the use of Artificial Intelligence increase stock value?*

---

### 1.1 Definition and history of Artificial intelligence

Since the early days of Artificial Intelligence, its concept has revolved around the capacity to act with an appearance of intelligence. Alan Turing, a British polymath, through several tests, declared that machines are intelligent when humans cannot distinguish their actions from those of human beings. Turing, in the 1950s, explored methods for creating intelligent machines and evaluate their intelligence, asking: *if humans are capable of processing information rationally to solve problems and make decisions, why can't machines do the same?*<sup>1</sup>

To date, there isn't just one definition of Artificial Intelligence but rather multiple ones. The European Commission, in 2018, told us that *"AI refers to systems that display intelligent behavior by analyzing their environment and taking action - with some degree of autonomy - to achieve specific goals"*.<sup>2</sup> This definition seems incomplete and does not

---

<sup>1</sup> <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>

<sup>2</sup> The European Commission's High-Level Expert Group on Artificial Intelligence. *A Definition of Ai: Main Capabilities and Scientific Disciplines*, Brussels, 18 December 2018.

define the methodologies employed to attain intelligence. AI is a hypernym, which means that it encompasses a set of processes, methods, tools, and techniques aimed at enabling computers to mimic human intelligence. These capabilities include *"reasoning, problem-solving, memory recall, planning, learning, natural language processing, perception, manipulation, social intelligence, and creativity"*.<sup>3</sup>

The term *Artificial Intelligence* was coined in 1956 during the Dartmouth Summer Research Project on Artificial Intelligence (DSRPAI) event, which brought together a group of scientists. In the workshop proposal, released a year earlier, researchers stated that *"an attempt will be made to find how to make machines use language, form abstractions, and concepts, solve kinds of problems now reserved for humans, and improve themselves"*.<sup>4</sup>

This summer marked the turning point, and from this time Artificial Intelligence officially became a field of study, thriving more and more. Computers were able to process and store larger amounts of information faster and more economically. Learning algorithms were implemented, and people were more educated in the operation of them. Optimism and expectations were so high that government agencies such as the Defense Advanced Research Projects Agency (DARPA) began funding research in this field. In 1970 Marvin Minsky, in an interview with *Life Magazine*, said: *"from three to eight years we will have a machine with the general intelligence of an average human being"*.<sup>5</sup>

The great initial optimism was short-lived, as researchers soon realized that computers were not yet at the desired level and that things were more complex than anticipated. McCarthy's doctoral student, Hans Moravec, claimed that *"computers were still millions of times too weak to exhibit intelligence"*.<sup>6</sup> Skepticism took the place of optimism, and we entered a phase called the "AI winter", which lasted from the mid-1970s to the early 1980s.

---

<sup>3</sup> Luce, L. (2019). *Artificial intelligence for fashion: How AI is revolutionizing the fashion industry*. Berkley, CA: Apress.

<sup>4</sup> McCarthy, J., Minsky, M., Rochester, N., Shannon, (1955). *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. AI Magazine Volume 27 Number 4 (2006).

<sup>5</sup> <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>

<sup>6</sup> <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>

Interest in AI grew again in the 1980s, thanks to increased funding and an expansion of algorithmic resources. "Deep learning" techniques enabling computers to learn from experience became widespread, and Edward Feigenbaum proposed expert systems capable of emulating the decision-making process of a human expert in the field.

By the 1990s and 2000s, Artificial Intelligence had made numerous strides, and many of the fundamental goals had been achieved. Most notable is the defeat of world chess champion Gary Kasparov against IBM's Deep Blue, a computer program capable of playing chess. This was, unquestionably, a major win for Artificial Intelligence, and studies continued.

In the new century, AI - supported by all the efforts of computer science - no longer represents a mere trend but emerges as one of the key players in shaping the future of human society.

### **1.1.1 The state of AI**

The modernity we experience today represents the result of a series of innovations that have occurred over time. For over a century, the concept of artificial intelligence has inspired countless science fiction authors and futurists. Today advancements in computer science and big data have turned it into reality, with machines now being deployed extensively across several fields. From the inception of the wheel, which transformed agriculture, to the intricate robotic assembly lines we see today, machines have played a crucial role in shaping the world as we know it. Yet, despite recognition of their incredible and inexhaustible usefulness, there is often a sense of fear towards machines; particularly, the fear that they may, one day, develop human-like intelligence and act autonomously.

Nowadays, intelligent machines are an essential component of our daily lives and culture landscape. They have reached extremely high levels of complexity, with astonishing computational speed, capable of performing as many calculations within a second as a person could accomplish in more than 31 billion years. But it extends beyond mere calculation: these devices have acquired and continue to acquire abilities and insights that were previously the exclusive domain of humans.

The application of Artificial Intelligence can extend across multiple levels, stimulating individual, economic, and corporate growth. As the working population percentage decreases in many countries worldwide, AI provides essential automation to support and accelerate productivity growth at both micro and macro levels. At the micro level, businesses are implementing various AI technologies to gain advantages such as lower labor costs, increased production, enhanced quality, and reduced downtime. At the macro level, automation is expected to drive significant productivity growth.

The potential for automation also varies significantly across different industrial sectors. For example, a McKinsey report indicates that the manufacturing sector has an automation potential of 64%, compared to 34% in the educational services sector. Several factors influence the speed and scope of automation, including openness to change, the cost of developing new technologies, labor market dynamics, regulatory frameworks, and economic benefits.

The highest levels of disruption are found in sectors such as health care, financial services, automotive, and education, where advances in AI are causing impressive changes in business models for the stakeholders involved.

However, despite these advances, particularly those regarding the application of cognitive thinking, machines remain constrained in their ability to improvise. They mainly follow programmed algorithms, limiting their actions to predetermined responses for each conceived scenario, thereby encountering a fundamental constraint of data-driven statistical infraction. They fail when faced with new situations because they lack the "common sense" that is the trademark of human cognition.

Indeed, we are witnessing a gradual but steady improvement in the capabilities of these new machines, even capable of outperforming humans in various activities, including those that require superior cognitive skills. A prime example is *Libratus*, an AI developed by Carnegie Mellon University and the first computer to defeat four professional poker players. Another example is *AlphaGo*, an artificial intelligence system developed by Google's DeepMind, which was able to beat Lee Sedol, considered the world's best Go player, an ancient Chinese board game rich in permutations and combinations. Also noteworthy is Dr. Fill, an AI computer that, in 2021, defeated nearly 1300 people in the annual American crossword puzzle championship.



All these developments have led to much speculation and fear about the future development of AI and the impact it will have on human life and evolution. Elon Musk, the co-founder of Tesla, has called AI "*our greatest existential threat*". Philosopher Nick Bostrom, in his work "Superintelligence", talks about the possibility of computers surpassing humans in cognitive ability. In addition, scientists from Google and Oxford University published a paper in September 2022 stating that superintelligent AI could become unruly and wipe out humanity in the long run. However, in light of recent advances and benefits from AI, as well as the necessity for human intervention at different levels, the foreseeable future is likely to witness a collaborative approach between AI and humans.

In this regard, cognitive scientist Alison Gopnik provides us with an important insight into the human-machine relationship:

*"One of the fascinating things about the search for AI is that it's been so hard to predict which parts would be easy or hard. At first, we thought that the quintessential preoccupations of the officially smart few, like playing chess or proving theorems—the corridas of nerd machismo—would prove to be hardest for computers. In fact, they turn out to be easy. Things every dummy can do, like recognizing objects or picking them up, are much harder. And it turns out to be much easier to simulate the reasoning of a highly trained adult expert than to mimic the ordinary learning of every baby."*<sup>7</sup>

### **1.1.2 The impact of AI**

According to a study conducted by Accenture and Frontier Economics, AI has a significant impact on a country's gross value added (GVA). Annual GVA growth rates in 2035 are estimated for both a base case scenario, based on current economic growth assumptions, and a scenario in which AI is an integral part of economic processes. The impact varies by country, with increases ranging from 0.8% in Italy or Spain to 2.0% in Finland or the United States. AI is also estimated to have an impact on labor productivity

---

<sup>7</sup> <https://www.edge.org/response-detail/26084> and Statista, *Artificial Intelligence: in-depth market analysis*

in different countries. In this case, the impact varies from an 11% rise in Spain to a 37% increase in Sweden.

Furthermore, we can estimate that artificial intelligence will affect economic growth rates at a weighted average of 1.7% by 2035, involving 16 industries. The analysis predicts that the information and communications, manufacturing, and financial services sectors will be the main beneficiaries in terms of annual gross value added (GVA) growth rates, with increases of 4.8%, 4.4%, and 4.3%, respectively.

In addition, the economic output of industrial sectors is projected to increase. For the manufacturing sector, AI can increase GVA by nearly \$4 trillion in 2035, followed by wholesale and retail trade (\$2.2 trillion), professional services (\$1.5 trillion), financial services (\$1.2 trillion), information and communication (\$1 trillion), transportation and storage (\$0.8 trillion) and construction (\$0.5 trillion).

## **1.2 AI Technologies**

The technologies behind Artificial Intelligence handle large amounts of data with intelligent algorithms, allowing software to detect patterns and learn from them. Following a sufficient period of training, the AI system uses what it has learned to make decisions, make predictions, and perform tasks without human support.

AI algorithms make use of several underlying components for their operation, such as Machine Learning, Artificial Neural Networks (ANNs), and Natural Language Processing (NLP).

### **Machine learning**

Machine learning, a branch of artificial intelligence, employs data and algorithms to emulate the learning process of humans, progressively enhancing its accuracy over time. Machine learning has been defined as *"the field of study that gives computers the ability to learn without being explicitly programmed"*<sup>8</sup> and focuses on designing new learning

---

<sup>8</sup> A. L. Samuel. (1959). *Some Studies in Machine Learning Using the Game of Checkers*, IBM Journal of Research and Development, vol. 3, no. 3, pp. 210-229.

algorithms and improving existing ones. In this type of technology, programs acquire knowledge from existing data and use it to analyze new data or to make future predictions by forecasting the value of non-existing data, thus enabling machines to learn without being manually programmed.

Machine learning begins with a collection of data (numbers, images, bank transactions, repair records, time series of data...), called training data, which constitutes the information on which the machine learning model is trained. This model is selected by the programmers, who, once provided with the data, allow the computer model to train itself to make predictions or find patterns. The programmers can, of course, later modify the model or parameters and test its accuracy through evaluation data. The outcome is a model that can be applied to different data sets in the future.

Machine learning algorithms can perform several functions: *descriptive function*, that is, the system uses data to explain the occurrence of events; *predictive function*; that is, the system uses data to predict future occurrences; or *prescriptive function*, using data to suggest an action to be taken.

In addition, machine learning has three subcategories:

- *Supervised Learning*: includes all those techniques aimed at responding appropriately to given stimuli. It is a machine learning method where algorithms are trained using labeled datasets and corresponding outputs to learn how to make predictions and identify specific patterns. Typical algorithms in this type of machine learning include neural networks, linear regression, decision trees and support vector machines.
- *Unsupervised Learning*: in this case, the models are fed with unlabeled datasets and are left to solve the problem and find patterns autonomously, without explicit instruction. In summary, the system is not provided with the correct answer upfront, but it is tasked with learning it autonomously, based on its experience in solving related problems. Algorithms that fall into this category include Gaussian mixture models, hierarchical clustering and Hidden Markov Models.

- *Reinforcement learning*: According to this approach, the algorithm learns by trial and error to associate situations with actions to optimize a numerical reward signal (Sutton & Barto, 2014, 2015). It then creates a memory of each experience so that it can use what it has learned in future situations.

An important field of Machine Learning is Deep Learning, which involves neural networks with multiple layers. The purpose of these neural networks is to simulate how the human brain operates, enabling the system to store a large amount of data. While achieving parity with the human brain is impossible, a multi-layer neural network can optimize and improve accuracy, increasingly simulating human reasoning. Deep Learning serves as the basis for many artificial intelligence functions and solutions capable of performing physical and analytical tasks independently. This technology underpins products and services that we use daily, such as systems detecting credit card fraud or digital assistants.

Unlike Machine Learning algorithms, Deep Learning ones do not require preprocessing of data (which is typically found in machine learning). They can acquire and process unstructured data, such as texts and images, and involve greater automation in extracting components, thereby almost eliminating dependence on humans.

### **Artificial Neural Networks (ANNs)**

Artificial Neural Networks are computer programs designed to mimic how the human brain processes information. Each neural network consists of hundreds of small units called "artificial neurons" or processing elements, interconnected with weights, creating a neural structure organized in layers that vaguely resembles the arrangement of neurons in the human brain. This type of technology acquires knowledge through experience rather than programming, learning from the detection of patterns and links in the data they analyze. Each processing element receives adjustable weighted inputs, processes them, and produces an output. To ensure maximum efficiency and accuracy in predictions, this technology needs to be trained and tested, and once training is complete, the networks can be used to make predictions on new data.

Artificial neural networks play an important role in handling complex problems where information sources are incomplete or unclear or in recognizing patterns by analyzing information. For example, this technology can help accurately predict market prices by being able to handle complicated and uncertain data; or it can help us assess the risk of an asset, allowing us to predict and identify the magnitude of the risk, enabling people to make more informed decisions.

### **Natural Language Processing**

Natural Language Processing (NLP), is a subset of artificial intelligence that empowers computers to create, understand and manage human language. This technology serves as a meeting point between users and computers, allowing the latter to interact with users using natural language, whether in the form of text or voice. We take advantage of this technology whenever we interface with voice assistants such as Siri, Cortana or Alexa, and their ability to respond in natural language to our requests is directly linked to Natural Language Processing. Other examples of this technology include grammar or spell-checking, automatic translation of a text, spam filtering of emails, and automatic response suggestions to a message or email that are consistent with their content.

Natural Language Processing has several practical applications that greatly simplify business operations. Indeed, it is of great help in optimizing search engines, automating routine tasks, enabling employees to work on more challenging and interesting tasks, giving market insights, and analyzing and organizing large collections of documents and data.

In the financial world, this technology is leveraged by traders to extract information from corporate documents and press releases for the purpose of extracting information relevant to their portfolios and trading decisions.

### **1.3 AI adoption at a corporate level**

A McKinsey study on the use of Artificial Intelligence, conducted from 2017 to 2022, shows an expansion in the use of this technology, stating that adoption has widely

increased in this time frame. In 2017, in fact, only 20% of respondents were involving AI in at least one operational area of their company, while in 2022 the figure reached 50% of respondents, peaking at 58% in 2019.

In addition, the average number of artificial intelligence capabilities used by organizations (such as, for example, robotic process automation, natural language understanding and processing, computer vision, and speech recognition) increased from 1.9 in 2018 to 3.8 in 2022, doubling its value. Robotic process automation and computer vision are the most frequently adopted skills over the 5 years studied, and natural-language text understanding climbed the rankings, ranking right after computer vision.

Regarding AI use cases, they have remained relatively stable over the considered timeframe. Service operations optimization has held the top spot since 2019, followed, in order, by creation of new AI-based products, customer service analytics, customer segmentation, customer acquisition and lead generation, contact-center automation, product feature optimization, risk modeling and analytics, and predictive service and intervention.

Furthermore, investments in artificial intelligence have shown increasing adoption. In 2018, 40% of the respondents stated that over 5% of their “*digital budgets*” was allocated to artificial intelligence. By 2022, however, the percentage had risen to over 50%, and looking ahead, 63% of the respondents expect an increase in investments by their organization in the next 3 years.

There has also been a shift in terms of the specific areas where companies see the value of artificial intelligence. While in 2018 the most valued functions were production and risk, in 2022, the greatest effects are in marketing and sales, product and service enhancement, strategic planning, and corporate finance. In addition, respondents point to several cost benefits of using artificial intelligence in the supply chain. A significant portion of respondents report that the implementation of AI systems in their companies has contributed to generating at least 5% of their organization's EBIT in 2021.

In the case of artificial intelligence, as in any other industry, it is possible to outline the figures of leading companies. These companies stand out for their ability to gain economic benefits from artificial intelligence, with 20% or more of EBIT coming from

the use of this technology and the great advantages they have, primarily, in terms of returns. These high-performing companies are becoming increasingly committed to developing AI on a large scale, or, as it is called by some, "industrializing AI". They are also distinguished by their efforts to link AI strategy to business results. These companies usually include the use of a modular data architecture to rapidly accommodate new AI applications and automation of data-related processes to improve efficiency in AI development.

Investment is another aspect that contributes to widening the gap: high-performing AI companies are more likely to spend large sums of money on AI-related efforts. In concrete terms, these companies spend more on average than others, which assumes a higher percentage of revenue. On average, among the high-performance companies surveyed, 20% of their technology budgets are allocated to artificial intelligence.

Continuing with our analysis, we can say that 2023 represents the breakthrough year for AI, as this technology has left a niche position and landed in a more mainstream dimension in a very short space of time. We all interact with artificial intelligence daily without realizing it: voice assistants such as Siri and Alexa are based on this technology, as are chatbots often made available by customer service on the web.

McKinsey's recent annual global survey on artificial intelligence (August 2023) introduces and confirms the explosive growth of a new tool: Generative Artificial Intelligence. One-third of respondents confirm that they use AI in at least one operational aspect of their business, and business leaders and boards personally use AI tools for work and discussions, showing that AI deserves attention not only from technology employees but from all business roles. Advances made with generative AI come into play in this landscape, which deserves credit for driving companies to invest more overall in AI, as reported by 40% of respondents. Moreover, it could be seen that organizations that had already embraced artificial intelligence were the first to explore the potential of generative artificial intelligence; in fact, 60% of organizations already open to artificial intelligence reported using generative artificial intelligence.

Expectations about the impact of generative artificial intelligence are high, especially for those employed in the tech and financial sectors, who are more inclined to anticipate

significant upsetting change from this new technology. Overall, all sectors will experience some degree of disruption; what will be different is the level of impact.

#### **1.4 AI investment effects on firm value**

Artificial intelligence is a disruptive technology. Companies in all industries are increasingly incorporating AI into business operations, and managers continue to invest in AI adoption to improve their operations. As reported by McKinsey Global Institute, AI has the potential to contribute an extra \$13 trillion annually to global economic output by 2030.

Despite the enthusiasm, concern remains, and some companies are hesitant to embrace and use these technologies. Companies are concerned that machines may lack the capacity to align with the goals of management, potentially resulting in poor decisions or critical errors during operations. AI adoption is not as straightforward as people think and it will continue to be a significant challenge for organizations, as they could face the consequences of reputational damage, revenue loss, and decreased public trust if the AI machine stumbles. While, on the contrary, proper use of it and a targeted strategy can bring immense value to the company.

The assessment of AI implementation costs and benefits plays a crucial role in determining its impact on a company's market value. The market will evaluate the discounted present value of both the long-term benefits derived from AI investment and the transient costs associated with it. If the expected long-term benefits outweigh the expected costs of the investment, the market value will increase and vice versa.

Artificial Intelligence typically requires large investments and entails high risk, and its benefits and payback time are difficult to measure. There are many costs associated with AI, and beyond the implementation expenses, such as system costs, project fees, consulting, and system integration and upgrade fees, organizations are likely to encounter other additional costs as they embrace this technology.

However, expectations are mostly positive, and it is believed that AI will bring long-term benefits while costs are temporary. Consequently, the market may assess AI adoption as positive news for companies, as shown by research and surveys conducted by various



organizations over the past decade. Many international companies, such as Amazon, Apple, and Facebook, have involved AI in their business to successfully automate their operations. The industries to benefit from AI are varied, Domino's Pizza, for example, has reduced delivery times thanks to AI, being able to predict delivery times more accurately. Barclays has used AI for fraud detection and enhancing the customer experience through chatbots. A study also predicted that the use of AI can reduce up to 20% of financial institutions' operating expenses costs. This kind of news usually strikes a positive chord with investors and allows them to have high expectations for AI. As a result, investors look with curiosity and interest at companies that claim to implement AI and may consider investing in them.

The implementation of AI has a significant impact on costs, risks, and revenues of companies, factors that are taken into consideration before making investment decisions. Based on greater or lesser aversion to risk, investors may view AI adoption as good news, while others may steer clear of companies implementing this technology.

A recent study, conducted by Lui et al. (2020), helps us quantify the impact of AI investments and the value of companies based on company announcements and market reactions; their purpose, in fact, is to assess the value of AI investments. To understand how AI investments influence the value and reputation of companies in the stock market, they adopted an event study methodology, focusing on the day before the announcement, the announcement day, and the day after the announcement, thus covering a three-day time window. The study concludes by stating that significant changes occur on the day of the announcement and that the average market reaction to an AI implementation announcement is -1.77%.

Although empirical evidence regarding market reactions to investments in AI specifically is, at currently limited, several studies, over the years, have analyzed stock market reaction to announcements of investments in other disruptive technologies, such as IT, for example. Based on a study conducted by Dos Santos, Peffers and Mauer (1993) we can say that when companies make investments in new technologies or IT technologies, these investments yield some direct benefits that contribute to future cash flows. For instance, an investment in an innovative technology project can improve a firm's ability to leverage this new technology in future projects, thereby positively affecting the firm's future

investment prospects. Unfortunately, it is not easy to quantify the direct benefits of investments in these technologies, so they are often underestimated or ignored.

Theoretically, if the net discounted cash flows arising from an investment (NPV) are positive, as the direct and indirect benefits of these investments are expected to generate a return greater than the required rate of return, then the value of the company should increase. The stock prices of these companies should reflect this change, and if the company's securities are traded in an efficient market, the change in value should occur even more rapidly, allowing it to be observed and measured.

The conclusions drawn from this study reveal that the market reacts differently to innovative IT investment announcements than it does to follow-up or non-innovative ones, showing that innovative IT investments increase company value, while non-innovative investments do not.

Research and studies in this field, although still relatively unexplored, do not seem to have stopped. A recent study conducted by Tania Babina, Anastassia Fedyk, Alex He, and James Hodson, analyzes the relationship between investment in artificial intelligence and economic growth, using detailed data on human capital and job postings across different industries. The research findings suggest that investment in artificial intelligence is associated with significant growth in sales and employment at the industry level. This indicates that industries with higher AI adoption tend to experience faster growth in their businesses and workforce.

### **1.5 AI potential to boost stocks**

The growing interest in artificial intelligence has led to a significant rise in technology stocks over the past year, with a set of U.S. corporations driving market performance upwards. According to a report by Goldman Sachs, those dubbed "early winners", primarily including US semiconductor manufacturers and cloud service providers, have seen a 60% increase by August 2023. Despite the fast and sudden rise in the share price of these companies, a strategist at Goldman Sachs Research, Oppenheimer, denies that they are in a bubble, saying that we are going through the *"relatively early stages of a*

*new technology cycle that is likely to lead to further outperformance*".<sup>9</sup> Moreover, in contrast to what happened in 2022, technology stock prices have still risen despite rising market interest rates, pointing out that investors expect higher growth rates in the future for these companies.

The recent stock market growth, fueled by interest in AI, has raised questions about future implications. While AI adoption may increase corporate profits in the long run, the extent and timing of the impact remain uncertain. However, according to another report published by Goldman Sachs on the U.S. landscape, economists estimate that the widespread adoption of artificial intelligence could increase productivity growth by 1.5% annually over a 10-year time frame, leading to a compound annual growth rate of the S&P 500 EPS of 5.4% over the next 20 years, compared to the current compound annual growth rate of 4.9%, marking a 9% increase in the fair value of the S&P 500 compared to its current value.

Focusing on the price-to-earnings (P/E) ratio of the U.S. technology sector, it is, currently, at its highest compared to the median and the 10-year range, placing technology sector valuations in a position significantly higher than historical standards. However, there is still an air of awareness and responsibility in the market, as unlike during the Internet bubble in the late 1990s, major technology companies today are less extreme in valuations and have a much more moderate P/E than seen in past bubble periods, demonstrating strong balance sheets and stable returns on investments. Moreover, US technology sector valuations have developed an unusual premium compared to technology companies in Europe. This highlights the importance of AI in driving market gains, as most of the leading companies in this sector are based in the United States. However, market growth is concentrated, with only a few companies contributing significantly to the overall returns of the S&P 500 Index.

---

<sup>9</sup> <https://www.goldmansachs.com/intelligence/pages/why-ai-stocks-arent-in-a-bubble.html>

# CHAPTER 2

## **Theoretical Framework**

---

The second chapter aims to provide the theoretical framework necessary to understand the topic discussed and to create a theoretical framework that will be used throughout the study. The chapter begins with a definition of the Efficient Markets Hypothesis, followed by a review of the existing literature on this theory and a discussion of the different forms of efficiency, with related empirical tests. Finally, the chapter concludes with an analysis of the criticisms that have been raised over the years, providing a comprehensive and critical view of this theory.

---

### **2.1 Efficient Markets Hypothesis**

A widely used theory in financial analysis of stock price movements is the efficient market hypothesis. According to this theory, markets react instantaneously to information, thus incorporating any new information about a company's value into stock prices. This concept is essential for understanding how the market responds to announcements and for interpreting the results of event studies.

#### **2.1.1 Definition**

In general terms, the capital market plays a crucial role in the economy by allocating ownership of capital between investors and firms. Essentially, this means that the capital market facilitates the transfer of ownership of economic resources, such as stocks and bonds, between those who own them and those who wish to invest in them. In general, the optimal situation would be one in which the market is able to provide precise guidance on how to allocate resources. This would imply assuming that stock prices and other financial securities fully reflect all available information, enabling companies to make informed decisions about what to produce and invest in, and investors to make informed

choices about which securities to buy. A market with these characteristics is called "efficient".

Although the term was coined before 1970, Eugene Fama is often recognized as the founder of the Efficient Markets Hypothesis. This is probably because he and his collaborators conducted the first study of events, and the findings derived from their research offered substantial backing to the assertion that the stock market operates efficiently. The breakthrough work on the Efficient Markets Hypothesis, published in 1970, is the first of three review articles by Eugene F. Fama, entitled "*Efficient Capital Markets: a review of theory and empirical work*". In this paper, Fama defines efficient as "a market in which prices always fully reflect available information".<sup>10</sup>

### 2.1.2 The Fair Game and Random Walk Theory

From a theoretical point of view, the definition provided by Fama is fully understandable, but, in itself, it lacks empirically testable implications. Therefore, to evaluate this theory of market efficiency, it is necessary to define the price formation process more precisely. Essentially, stating that prices "fully reflect" available information, what implications does it entail empirically?

Most theoretical models and empirical studies in this area are based on the assumption that market equilibrium can be represented on the basis of expected returns, defining the so-called *Expected Return* or "*Fair Game*" Models. This entails assuming that, conditional on a set of relevant information, the equilibrium expected return of a security is influenced by the perception of its "risk." This "risk" can be interpreted in many ways, and different market theories primarily differ in their definition of risk. The notational representation of these expected returns theories is as follows:

$$E(\tilde{p}_{j,t+1} | \Phi_t) = [1 + E(\tilde{r}_{j,t+1} | \Phi_t)]p_{jt}$$

---

<sup>10</sup> Fama, E. F. (1970). *Efficient Capital Markets: A Review of Theory and Empirical Work*. The Journal of Finance, 25(2), page 383.

Where  $E$  represents the expected value operator;  $p_{j,t}$  and  $p_{j,t+1}$  are the prices of stock  $j$  at time  $t$  and time  $t + 1$ , respectively;  $\tilde{r}_{j,t+1}$  represents the percentage return in the period  $t - t + 1$ , in expanded form  $(p_{j,t} - p_{j,t+1})/p_{j,t}$ ;  $\Phi_t$  is a symbol to indicate a set of information assumed to be fully reflected in the price at the time  $t$ ; and the tilde, which crowns  $\tilde{p}_{j,t+1}$  and  $\tilde{r}_{j,t+1}$ , is used to indicate that they are random variables at time  $t$ . The left-hand side of the equation  $E(\tilde{p}_{j,t+1} | \Phi_t)$ , represents the expected value of the future stock price at the time  $t + 1$ , conditioned on the information available up to time  $t$ . This expectation takes into account currently available information and forecasts of the future performance of the security. The right-hand side of the equation  $[1 + E(\tilde{r}_{j,t+1} | \Phi_t)]p_{jt}$  is a product of the current price of security  $j$  and the expectations of its future return, also conditioned on the information available up to time  $t$ . The value of these conditional return expectations  $E(\tilde{r}_{j,t+1} | \Phi_t)p_{jt}$  is determined by this expected return theory, attempting to represent mathematically that, whatever the reference model, information (represented by  $\Phi_t$ ) plays a key role in determining equilibrium expected returns.

Stating that equilibrium expected returns are formed based on available information, however, has significant empirical implications. These assumptions rule out the possibility of the existence of trading systems that, relying only on the information available at the time, allow for "easy" profits or expected returns higher than those predicted by the equilibrium model. This is because any relevant information would already be embedded in prices; therefore, there would be no opportunity to exploit information not already reflected in expected returns. In this context, market efficiency can be expressed as the inability to obtain excessive profits or expected returns above the norm by exploiting trading strategies based only on the available information at the time.

Furthermore, in the early days of studies of efficient markets, it was believed that if the price of a security fully reflects all available information, then future changes in prices should be independent of each other and follow the same probability distribution. These two assumptions are at the basis of the so-called "Random Walk" model, which can be mathematically represented as follows:

$$f(r_{j,t+1} | \Phi_t) = f(r_{j,t+1})$$

Formally, this means that the conditional probability of a change in future price (more commonly understood as a change in returns over a period), given a set of information available at the time  $t$ , is equal to the marginal probability of the same. It is further assumed that the nature of price variations does not change over time and that, therefore, the density function describing this probability distribution must be the same at all times.

The Random Walk model is, of course, more comprehensive and specific than the expected return model. If we were to limit the latter by assuming that the expected return of a security is constant over time, we would obtain an equality between the mean of the distribution of returns at time  $t$  and the average of projected returns to the future  $t + 1$ . However, the Random Walk model goes further than this, claiming that the entire distribution of returns is independent of time. In simpler terms, while in the expected return model only the mean remains constant over time, in the Random Walk model, both the variance and the shape of the distribution of returns remain unchanged.

Consequently, we can assert that the Random Walk Model can be considered an expansion of the "Fair Game" model, as the latter does not describe the stochastic process of yield formation, but merely states that market equilibrium conditions can be expressed in terms of expected returns. Indeed, if investors' reactions to new information and the process of generating that information are unpredictable, then it is not possible to predict future returns based on past ones. And that is precisely the essence of the Random Walk model: each step in the process is independent and unpredictable.

### **2.1.3 Consistency between market circumstances and efficiency**

To better understand market efficiency, it is important to outline the ideal conditions for its achievement. In an ideal setting, the market appears frictionless and free of transaction costs, allowing investors to sell and buy securities without incurring additional expenses. This would remove one of the main obstacles to market efficiency, allowing for perfect alignment between security prices and available information. Moreover, in an ideal market, all the most significant information regarding securities would be open to all, allowing free access to anyone. This would entirely remove barriers to information access, and all investors would have a uniform knowledge of securities and their

characteristics. Thus, a situation would be created in which investors are in complete agreement on how to interpret the available information.

However, this scenario remains unattainable in real-world markets, where these ideal conditions are not always met. Indeed, practical considerations such as transaction costs – including brokerage fees and taxes – constrain investors' ability to freely view and buy securities as discussed above. In addition, information is hardly available to everyone; in fact, it is often restricted to institutional or privileged investors, resulting in obvious disparities in its interpretation.

Fortunately, the ideal conditions described above are sufficient for market efficiency, albeit not necessary. Indeed, as long as all market participants consider all available information in the market, even the presence of large transaction costs that may potentially inhibit the transactions flow does not necessarily imply that prices do not fully reflect available information at the time of trading. Similarly, any divergence of opinion among investors on the implications of available information does not by itself prove that the market is not efficient.

However, it must be considered that these obstacles exist and impact real markets to some extent. Assessing the extent of this impact on the price formation process is, of course, the main objective of empirical research in this field. Numerous studies on this subject have followed, greatly expanding the study and importance of the field, and also giving rise to divergent schools of thought on this market efficiency hypothesis.

#### **2.1.4 Different types of Efficiency Forms**

Empirical research on Efficient Market Theory investigates how security prices reflect available information. This has evolved over time and has gone through several stages of investigation, marking the emergence of different types of market efficiency. Initially, studies focused on so-called *weak-form tests*, which consider only past prices and investigate their possible usefulness in predicting future price movements. Most of the findings in this field came from the literature on Random Walk Theory. Next, *semi-strong form tests* were conducted, which were interested in how quickly prices adjust to publicly available information, such as annual reports, new stock issues, or dividend



announcements. Finally, *strong-form tests*, which focused on whether investors or specific groups monopolized information relevant to price formation.

### **Weak Form tests of the Efficient Markets Model**

The weak form of the Efficient Markets Hypothesis, commonly associated with Random Walk theory, asserts that fluctuations in stock prices from the previous day are unrelated to one another. This means that there is no trend in prices and that any increase or decrease occurring today will not affect the stock price tomorrow. In 1965, Fama himself stated that *"the past history of series of share price changes cannot be used to predict the future in any meaningful way"*.

Empirical tests of Random Walk Theory can be performed by evaluating, from a statistical point of view, whether there is a serial correlation between share prices. Serial correlation (a term now considered synonymous with autocorrelation and using the correlation coefficients of a consecutive series of stock price changes) occurs when changes in stock prices in a time series exhibit a trend caused by the influence of previous price changes on current price changes.

Tests on the weak form of efficiency have shown that the acceptance of Random Walk Theory, while sufficient, is not entirely necessary. This means that in order to claim the existence of weak-form efficiency conditions, it is sufficient to have a time series of prices whose degree of serial correlation, if any, is not exploitable economically to achieve above-average returns. In simpler terms, even if there may exist a (small) serial correlation, it will be conveniently exploitable only if the expected excess profits generated are sufficient to offset the transaction costs related to the brokerage of securities. Nearly all the tests have found the existence of a small serial correlation (positive or negative), without, however, leading to a rejection of the acceptance of efficiency in a weak form, precisely because of the costs associated with trading securities that make the trend not conveniently exploitable, eroding, or, even, exceeding the profits that could be obtained by exploiting these correlations.

In addition to Random Walk Theory, there are other popular assumptions in financial econometrics, such as Rational Expectations and the Martingale model. Rational

Expectations (RE) theory, introduced by John F. Muth in 1961, states that individuals formulate future forecasts rationally, using all available information at the time of the forecast. Furthermore, Muth assumes that future price changes are random, as no one knows what will happen at time  $t + 1$ . However, the novelty of his theory is that future expectations are "optimal," indicating that the expected forecast error (i.e., the difference between the forecast and the actual outcome) will balance out to zero in the long run. From a mathematical point of view, this theory can be explicated as follows:

$$P_{t+1} = E_t[P_{t+1}] + \varepsilon_{t+1}$$

Where  $P_{t+1}$  is the expected price at time  $t + 1$ ;  $E_t[P_{t+1}]$  represents the expected value of the price at time  $t + 1$ , given the set of information available up to the present time  $t$ ; and  $\varepsilon_{t+1}$  is the forecast error at time  $t + 1$ , representing the divergence between the expected and actual price at time  $t + 1$ . The key concept here is that the expected forecast error,  $E_t[\varepsilon_{t+1}]$ , cancels out, indicating that individuals' expectations of future prices are unbiased and, on average, correctly predict the future price. This implies that investors cannot earn abnormal profits consistently by yielding predictable patterns in asset prices, since these patterns would already be reflected in investors' expectations. In expanded form:

$$E_t[\varepsilon_{t+1}] = E_t[P_{t+1} - E_t[P_{t+1}]] = 0$$

This equation demonstrates that the expected error is not correlated with the information available at time  $t$ , showing that past information or patterns in the data do not allow market participants to systematically predict future changes in asset prices and suggesting that the error process should exhibit serial correlation, meaning that past forecast errors would influence future errors.

The martingale model, on the other hand, is based on probability theory and the previously discussed "Fair Game." The underlying assumption of the model defines that the price of a variable at time  $t + 1$  is equal to its price at time  $t$ , using all past information. From an analytical point of view, if the conditions  $E[P_{t+1} | I_t] = P_t$  and  $E[(P_{t+1} - P_t) | I_t] = 0$  are met, the martingale is a stochastic process of  $P_t$ . This means that the expected value of the future price conditional on the information available up to time  $t$  is

equal to the current price. This model further implies that if the two conditions above are met, then the expected variance is zero, indicating that the probability of positive or negative variance occurring is the same. Essentially, we can conclude that there is no systemic relationship between price changes in different time periods.

### **Semi-strong Form tests of the Efficient Markets Model**

This second form states that stock prices already incorporate not only historical information but also all public information. This information concerns past returns, balance sheets, earnings, or dividend announcements, and, more generally, the financial situation of companies.

One of the first to test this type of efficiency was Alfred Cowles, an American economist, in 1933. He examined the predictive capabilities of 45 financial analysis agencies with respect to the stock market, comparing the forecasts made by these financial experts with the market outcome. At the end of his study, he discovered that the experts' recommendations did not produce better returns than the market itself, concluding that a more informed investor or professional does not perform better than a less informed one.

Another important notable test is the one conducted in 1969 by Fama, Fisher, Jensen and Roll. Their study focuses on the effects of stock splits on stock prices, starting from the assumption that, since a stock split is typically a formal change in a company's capital structure, the impact on the stock price itself should be limited or almost nonexistent. This may be the first documented case of using the event study as a methodology.

Initially, they identified the month in which a stock split occurred as time zero for that specific stock. Next, they estimated the expected returns for each month of the stocks in their sample using the single index model, expressed mathematically as follows:

$$R_{i,t} = a + b_i R_{m,t} + e_{i,t}$$

Where the expected residual value  $e_{i,t}$  is zero. This group of scholars paid close attention to the analysis of residuals, which are the difference between actual and expected returns. Then they examined the residuals ( $e_{i,t}$ ) for each month and for each stock in their sample and proceeded to calculate the average for each month considered. The most relevant

result was found when analyzing whether there were differences in the results between companies that experienced an increase in dividends at some point after the split and those that did not. Companies with increasing dividends showed a positive residual performance after splits, showing that the market promptly reacted to the announcement of a stock split, adjusting share prices only according to expected changes in future dividend payments. In conclusion, they point out that there is no way to exploit a stock split to achieve superior returns from it, unless one possesses privileged (and therefore non-public) information, but we will discuss this further later. However, this study has been a key guide for subsequent studies and tests on semi-strong form efficiency.

Michael Arthur Firth also conducted a test to assess semi-strong form efficiency. He focused on the effect of communications regarding the acquisition of over 10% of a company, a potential sign of an impending merger. His analysis started by calculating Cumulative Abnormal Returns (CARs), a measure of the deviation of actual stock returns from expected ones, going back 30 days prior to the announcement. The results obtained showed that most of these CARs occur between the last trade before the announcement and the first trade after the announcement (to a greater extent). Firth concluded that the market can be considered efficient in terms of semi-strong form because returns change almost immediately after the (near) merger announcement. This finding suggests that there are no opportunities to achieve excess returns solely based on this information.

In conclusion, the Challenger explosion is an event worth mentioning in this context. The Challenger shuttle disaster occurred on January 28, 1986, precisely at 11:47 am. The news quickly spread, and the spotlight was on several companies involved, such as Rockwell International and Lockheed, who preferred not to make statements. The accident was caused by poor quality rocket booster rings supplied by Morton Thiokol, which failed at low temperatures. On the same day, Morton Thiokol saw its stock performance plummet, registering -11.86%, and experienced an unprecedented trading volume, with 1.74 million shares traded that same day (compared to the average daily trading of 100,000 shares over the previous three months). The other companies involved, mentioned earlier, also exceeded the average trading volume, but to a lesser extent than Morton. This event shows how quickly public information is reflected in the stock prices and how swiftly the market weighed Morton Thiokol compared to the other companies.

## **Strong Form tests of the Efficient Markets Model**

Tests of the strong form of the Efficient Market Hypothesis are aimed at demonstrating that prices fully reflect all available information, such that higher profits than the market cannot be achieved by exploiting an information monopoly. In this context, it is important to point out that, while empirical evidence provided for the two previous forms effectively demonstrates the feasibility of the hypotheses, in the specific case of the strong form there is contradictory evidence that undermines the effectiveness of the model, not allowing for a full fit of the model to reality. In this regard, a study conducted by Victor Niederhofer and Osborne, two pioneers of modern finance, confirms that some specialized operators can concretely enjoy a monopoly of information. Specifically, they discuss privileged access to information on limit orders by NYSE specialists, allowing them to make larger profits.

Despite the validity of the model being, in part, compromised, studies on this topic have not stopped, investigating the extent of deviations from the Efficient Markets Hypothesis model, and raising questions about investment strategies and access to privileged information in the stock market. In this regard, great credit is given to Jensen for his study to investigate the performance of mutual funds. In his assessment, Jensen outlines two main objectives: to understand whether the managers of these funds have privileged access to information such that they are able to achieve above-average returns, and to assess whether some funds make better use of such information than others. In evaluating the performance of funds, Jensen uses the equilibrium model of expected returns proposed by Sharpe and Lintner, which provides a guideline for establishing a benchmark that is consistent with the very objectives of the study.

In the empirical development of his study, Jensen evaluated the performance of 115 mutual funds by analyzing the relationship between risk, estimated based on historical data on annual interest rates with continuous interest reinvestment, and return, measured through analysis of annual nominal returns with continuous compound interest over a period of ten years. To compare these performances with those of the market, he used the S&P 500 index as a proxy in constructing the market portfolio.

To determine whether the mutual funds under consideration are able to achieve higher than market returns by exploiting privileged information, Jensen proceeds systematically. Initially, he tries to determine whether the funds are able to achieve returns that compensate investors for the expenses strictly related to fund management and maintenance, emphasizing that these expenses can be easily avoided by investing in a combination of risk-free investments and market portfolio (while maintaining the same level of risk as the funds). The answer is negative, as nearly 80% of the funds failed to earn a return sufficient to compensate for the risk taken compared to an investment in a representative market portfolio. In fact, on average, there was a -14.6% difference between the returns achieved by investing in the funds and the expected returns along the market line.

What if we ignored these management fees? In the next step of his study, Jensen confirms that the answer is still negative, even if we eliminate these expenses. In fact, more than 60% of the funds are unable to record returns at least equal to those expected along the market line, showing a difference of -8.9%.

Ultimately, Jensen tried to understand whether, ignoring all previously discussed expenses, mutual funds could somehow identify securities capable of outperforming market returns. Due to delays and discrepancies in publishing data on the actual amount of these fees, it is not possible to provide a specific answer for individual funds. He suggests that, although most of these funds are unable to take advantage of information not yet fully incorporated into prices, there may still be specific cases that achieve above-average returns, albeit without being able to provide empirical evidence of this. Generally speaking, his study shows that the number of funds actually capable of outperforming market returns constitutes a tiny percentage compared to the overall number examined, concluding that mutual fund managers do not have such abilities to predict returns and exploit them advantageously. In light of these results, Jensen states that, despite the skills of the managers of these funds, *“the fact that they are apparently unable to forecast returns accurately enough to recover their research and transactions costs is a striking*

*piece of evidence in favor of the strong form of the martingale hypothesis - at least as far as the extensive subset of information available to these analysts is concerned”<sup>11</sup>*

## **2.2 Criticism**

The theory of efficient markets, despite its rigidity, acknowledges the possibility that the market price may not correspond to the intrinsic value of a stock. This is because investors' reactions to available information can sometimes be excessive or insufficient. However, these deviations should be random and should not be linked to any specific variable, thus assuming that stocks are equally likely to be overpriced or underpriced.

Initially, the theory of efficient markets received a positive response from the entire economics community of the time, mainly due to Fama's valuable contribution, as discussed earlier. However, with the advent of the twenty-first century, the first skepticisms began to emerge, pointing out the presence of several limitations in the model.

Critiques primarily target the underlying assumptions. Indeed, claiming that all market participants have the same rational expectations and have access to the same information becomes empirically problematic. In reality, investors, in situations of uncertainty and risk, do not possess the pure rationality assumed by the EMH. Simon refers to this deviation as Bounded Rationality. Other scholars, such as Kahneman and Riepe, support the same theory, stating that in financial decision-making, which by definition is uncertain and lacks fixed rules, investors tend to rely on instinct, deviating from a process of rational decision-making. Economist John Conlisk, through an empirical study on biases, also demonstrates that people make many reasoning errors that greatly impact economic decisions. He specifies that the types and severity of these errors are influenced by economic variables such as incentives and level of experience, emphasizing that this is not a deviation from economic reasoning but, rather, a necessary extension.

Kahneman and Tversky, on the other hand, tell us about the imperfections found in financial markets, stating that they are attributable to inherent human cognitive biases and

---

<sup>11</sup> Fama, E. F. (1970). *Efficient Capital Markets: A Review of Theory and Empirical Work*. The Journal of Finance, 25(2), page 413.

recurrent errors in reasoning and information processing. This thought has been theorized, acknowledging these two scholars for the birth of Behavioral Finance Theory. In addition, some studies have highlighted market anomalies that contradict the market efficiency. For example, Jegadeesh and Titman state that if stocks have shown positive (or negative) returns in the previous 3-12 months, it is very likely that they will continue with the same trend for the next 3-12 months. As Fama himself admitted, this deviation, known as Price Momentum, constitutes one of the most relevant anomalies of the efficient market model.

Other relevant phenomena that do not align with efficiency theory include volatility not linked to events or news and calendar effects. Proceeding in order, Cutler states that prices can fluctuate sharply even in the absence of major announcements or news, as the collapse of the Dow Jones Industrial Average in 1987 demonstrated. Regarding calendar effects, several studies have identified patterns in stock returns correlated with the days of the week, the change of the month, and the beginning of the year, the latter known as the "*January effect*".

In conclusion, the disagreement between different schools of thought on the subject is mainly based on how market rationality and anomalies are interpreted. While proponents of market efficiency attribute these anomalies to imperfections in pricing models such as the Capital Asset Pricing Model (CAPM), behaviorists believe that their presence is due to market inefficiency. In general terms, despite moments in which investors' judgment falters and stock prices are not consistently regular, the market tends to make good use of available information, confirming an underlying efficiency. While there is evidence of extraordinary market events such as bubbles or anomalies in stock pricing, these do not remain persistent over time and, therefore, do not provide a reliable method for achieving higher returns. In fact, if there were opportunities to make greater profits, this would be immediately exploited by investors, inhibiting their ability to persist over time. This makes us realize that despite the presence of anomalies and irrationality, the market will always tend to correct itself in the long run.



# CHAPTER 3

## Field research

---

The third chapter focuses on conducting three event studies to analyze the effects of announcements of investments in artificial intelligence on the stock prices of two different companies: Nvidia and Microsoft. Through rigorous analysis, the methodology used to collect and analyze the data, the statistical model employed, and how the calculations were performed will be thoroughly illustrated, aiming to assess the ability of financial markets to reflect relevant information promptly and accurately. Finally, the results of the studies, and their limitations will be presented and interpreted, providing a rigorous interpretation of stock market behavior in response to these announcements.

---

### 3.1 Introduction to Nvidia and Microsoft

Before delving into the explanation of the actual studies, it is crucial to provide a brief overview of the two main companies involved: Nvidia and Microsoft.

#### **Nvidia**

Nvidia Corporation is one of the largest U.S. technology companies specializing in the production of semiconductors and Graphics Processing Units (GPUs) designed to optimize and accelerate the processing of complex graphics used in a wide range of sectors, including video games, video editing, three-dimensional (3D) graphics, artificial intelligence and machine learning. It was founded in 1993 by Jensen Huang, Chris Malachowsky and Curtis Priem, driven by the shared ambition to innovate the gaming and multimedia markets by introducing 3D graphics. Nvidia, based in Santa Clara, California, holds a dominant position in the GPU industry, commanding approximately 80% of the global GPU semiconductor chip market share through 2023.

Since the invention of the GPU in 1999, Nvidia has continuously strived to enhance artificial intelligence computing, which plays a prominent role in the company's value proposition. As AI evolves and high-performance computing increases, NVIDIA GPUs are enabling the world's leading industries to harness accelerated computing and unleash AI to its fullest potential. In addition, several products and systems in Nvidia's portfolio have contributed significantly to neural networks training, providing high-performance computing densities and large-scale cloud services.

Nvidia's contribution to the progress and implementation of intelligence is immense and will continue to be so in the future. As Nvidia's CTO, Michael Kagan, states, the company's goal is to *"build the AI factories of the future"* leveraging the *"wealth of Nvidia technologies."*<sup>12</sup>

## **Microsoft**

Microsoft Corporation is a leading American multinational technology company. Originally focused on operating systems such as MS-DOS and Windows, Microsoft has since expanded its presence into the areas of personal computing, entertainment, corporate cloud solutions, and AI technology. The company boasts research facilities in various locations worldwide and operates through three primary segments, such as productivity and business processes, intelligent cloud, and personal computing, providing a broad array of products and services across multiple technology sectors.

It was founded in 1975 on the initiative of Bill Gates and Paul Allen to create software for one of the first personal computers on the market: the Altair 8800. Through an agreement with IBM to provide an operating system for the company's first personal computer and the launch of Windows, a graphical operating system for PCs, Microsoft became the world's largest personal computer software company within just ten years after its founding.

Microsoft is heavily involved in innovating and implementing artificial intelligence, conducting ongoing research to expand the company's efforts in the field of AIs, mainly

---

<sup>12</sup> <https://www.eetimes.eu/nvidias-michael-kagan-building-on-ai-iphone-moment-to-architect-data-processings-future/>

through the establishment of a large group of researchers, known as "*Microsoft Research AI*". The group's goal is to accelerate the adoption of AI in as many aspects as possible, including services, agents, infrastructure and applications. In addition, Microsoft has shown considerable commitment to addressing ethical and social issues related to AI. As stated by CEO, Satya Nadella, the company aims to make artificial intelligence accessible to everyone by "*democratizing access to intelligence to help solve our most pressing challenges.*"<sup>13</sup>

### **3.2 The study Methodology**

This study follows a deductive approach, using existing models and theories in the research field to formulate hypotheses that will be tested with an empirical data set. The efficient markets theory is employed to quantify and evaluate the impact of artificial intelligence investment announcements on the stock market by analyzing the change in stock prices of two major market players, Microsoft and Nvidia. The objective of the study is to analyze the abnormal returns caused by such announcements, thereby quantifying the impact of AI investments on the stock price. The change in stock prices on the announcement day is expected to reflect investors' perceived increase in value in market returns resulting from AI investments, as will be illustrated later.

The theory underlying the study, namely that of efficient markets, states that if the market in which the company's shares are traded is efficient, the increase in share price will reflect the increase in value of the analyzed company. This can be due to the fact that, since investors act rationally and efficiently, they will be able to recognize the change in company value as a result of its investment in AI and react accordingly.

The methodology employed for the analysis is the event study, a highly regarded technique in finance for assessing stock market reactions to specific events, such as investment announcements, mergers, or acquisitions. This method is based on the idea that investors immediately react and reflect such events through changes in the stock prices of the companies involved. In this specific case, the events whose impact is being

---

<sup>13</sup> <https://news.microsoft.com/2016/09/29/microsoft-expands-artificial-intelligence-ai-efforts-with-creation-of-new-microsoft-ai-and-research-group/>

analyzed consist of investment announcements in artificial intelligence by Microsoft and Nvidia.

In particular, for Nvidia, reference is made to an announcement dated January 8, 2024, in which the company reveals significant investments in Tensor Core GPUs, LLMs, and tools for PCs, while also announcing the forthcoming release of new cutting-edge AI software and tools. As for Microsoft, however, two events are considered. The first, dated January 11, 2024, announces a significant investment in artificial intelligence, unveiling new generative AI solutions and data to improve the customer experience, with the introduction of copilot features through Microsoft Cloud for Retail. The second event, on May 2, 2024, sees Microsoft announcing a \$2.2 billion investment to support Malaysia in its journey towards embracing cloud and artificial intelligence.

To study the impact of these events, quantitative methods are employed to capture and analyze the relationship between investments in artificial intelligence and stock prices. The market reaction is assessed by collecting daily stock prices over a specific time frame, which were downloaded from Yahoo Finance. Subsequently, the observations were divided into two time sub-windows, called *estimation window* and *event window*, which will be explained in detail later.

In the specific case:

- For **Nvidia**: the estimation period extends from May 2, 2023, to December 20, 2023, and includes 162 observations. The event window, on the other hand, includes the stock prices of the ten days before the announcement, the day of the announcement, and the following ten days, for a total of 21 observations between December 21, 2023, and January 23, 2024.
- For **Microsoft**: the estimation window includes 200 observations for both the events analyzed. For the first one, the observations span from July 3, 2023 to December 26, 2023; while for the second one, they span from July 3, 2023 to April 17, 2024. The event window, as for Nvidia, includes the share price of the ten days prior to the announcement, the day of the announcement, and the following ten days, for a total of 21 observations. They are located between December 27, 2023,

and January 26, 2024, for the first event and between April 18, 2024 and May 16, 2024, for the second event.

Firstly, the normal returns of the examined securities were calculated based on a statistical benchmark model. In statistical models, the calculation of returns typically assumes that asset returns follow a multivariate normal distribution. This means that the returns are distributed such that any linear combination of them follows a normal distribution. Furthermore, returns are assumed to be independent and identically distributed over time, implying that the returns in a given period are not affected by the returns in previous ones and that their distribution remains constant over time.

The statistical model chosen for conducting this analysis is the *market model*, which relates the return of a specific security to the return of the overall market. According to this model, assets returns are assumed to follow a joint linear distribution, allowing a linear specification to be used. In addition, the market model is believed to be more comprehensive and accurate compared to other statistical models because, by considering the relationship between the stock and the market, it eliminates the part of the return that is not explained by the market, reducing the variance of abnormal returns and making it easier to detect the effects of events.

Two market indexes were then selected to replicate market performance: the Nasdaq Composite, a stock index of the U.S. Nasdaq stock exchange composed mainly of technology companies, and the S&P 500, a capitalization-weighted index designed to reflect the performance of the 500 largest publicly traded companies in the United States. This results in a dual analysis of returns, carried out by first considering one index and then the other, in order to compare the results obtained and assess any differences in market behavior.

In the third step of this study, abnormal returns were calculated as the difference between actual returns and expected returns (the latter calculated using the market model) in the estimation window. This allows for the isolation of the impact of specific events on stock returns, independent of general market fluctuations. Next, the null hypothesis was defined, stating that there is no effect of the announcements on stock returns, implying that abnormal returns are zero. This serves as our benchmark for significance testing. To test the validity of the hypothesis, aggregate abnormal returns (CARs) were calculated

over time, as they provide greater significance of effects, reduce variability, and have greater statistical power than simple abnormal returns.

This explanation of the methodology provides a general overview of the analysis conducted. However, quantifying and evaluating the impact of these events on stock prices requires further statistical analysis, which is addressed in the following section.

### 3.3 Statistical Method

The study begins with the calculation of logarithmic returns, which were preferred over the calculation of simple returns, as advised by Fama in one of his studies (1965). The calculation equation requires the returns to be calculated as the difference between the natural logarithm of the price at time  $t$  and the logarithm of the price on the previous day  $t - 1$ :

$$r_t = \ln(P_t) - \ln(P_{t-1})$$

Fama explains that logarithmic returns offer advantages over simple returns in terms of greater stability of their variance and greater accuracy in assessing percentage changes in prices over time.


After introducing the method of calculating returns, it is essential to statistically define the reference statistical model chosen for the analysis, namely the *market model*. By establishing a linear correlation between stock returns and market returns, it provides an estimation of expected returns, allowing for the quantification of deviations from expected returns attributable to the specific events under study. For any one security  $i$ , the market model can be represented as follows:

$$R_{it} = \alpha_i + \beta_i R_{mt} + \varepsilon_{it}$$
$$E(\varepsilon_{it}) = 0 \quad \text{var}(\varepsilon_{it}) = \sigma_{\varepsilon_i}^2$$

Where  $R_{it}$  represents the returns of stock  $i$  for period  $t$ ;  $R_{mt}$  denotes the returns of the selected market index chosen as a proxy for market performance for period  $t$ ;  $\alpha_i$  and  $\beta_i$  are the parameters of the model; and  $\varepsilon_{it}$  represents the residual error of the model or

disturbance term, which, in expectation, has zero mean and variance equal to  $\sigma_{\varepsilon_i}^2$ . As mentioned earlier, the market indices used in this study are the Nasdaq Composite and the S&P 500.

This model provides a basis for the analysis of returns and their deviations from expected values. To facilitate the measurement and analysis of these deviations, the data have been divided into different time windows, as previously anticipated.  $t = 0$  represents the *event date*, the period from  $t = T_0 + 1$  to  $t = T_1$  identifies the *estimation window*, and the period from  $t = T_1 + 1$  to  $t = T_2$  represents the *event window*. We can summarize these time windows by defining  $L_1 = T_1 - T_0$  as the length of the estimation window and  $L_2 = T_2 - T_1$  as the length of the event window. Optionally, we can define an additional post-event window from  $t = T_2 + 1$  to  $t = T_3$ , with a length  $L_3 = T_3 - T_2$ . This time division can be represented as follows:



Since the estimation of market model parameters, aimed at calculating expected returns, takes place along the estimation window, it is important that the two time windows, the event and estimation ones, do not overlap, so as to avoid distortions.

To measure the expected returns of a stock, it is first necessary to estimate the parameters of the market model. The ordinary least squares (OLS) estimation method allows for the identification of an unbiased estimator for each parameter of the model. They are:

$$\hat{\beta}_i = \frac{\sum_{t=T_0}^{T_1} (R_{it} - \hat{\mu}_i)(R_{mt} - \hat{\mu}_m)}{\sum_{t=T_0}^{T_1} (R_{mt} - \hat{\mu}_m)^2}$$

$$\hat{\alpha}_i = \hat{\mu}_i - \hat{\beta}_i \hat{\mu}_m$$

$$\hat{\sigma}_{\varepsilon_i}^2 = \frac{1}{L_1 - 2} \sum_{t=T_0+1}^{T_1} (R_{it} - \hat{\alpha}_i - \hat{\beta}_i R_{mt})^2$$

where

$$\hat{\mu}_i = \frac{1}{L_1} \sum_{t=T_0+1}^{T_1} R_{it}$$

$$\hat{\mu}_m = \frac{1}{L_1} \sum_{t=T_0+1}^{T_1} R_{mt}$$

In which,  $\hat{\mu}_i$  and  $\hat{\mu}_m$  are the averages of the returns for stock  $i$  and market returns within the estimation window  $L_1$ .

Once the parameters of the market model have been estimated, the abnormal returns can be calculated and analyzed. They are the result of the following formula:

$$\widehat{AR}_{it} = R_{it} - \hat{\alpha}_i - \hat{\beta}_i R_{mt}$$

Where, defining  $t = T_1 + 1, \dots, T_2$ ,  $AR_{it}$  identifies the abnormal returns of security  $i$  in the event window. This difference between the actual return of a security and the return predicted by the market model is considered the disturbance term of the model, as it captures unexpected deviations from the expected return. The calculation of abnormal returns occurs out-of-sample, as the market model parameters are estimated in an estimation window prior to the event.

The calculation of abnormal returns is useful for conducting a hypothesis test to aimed at evaluating whether an event has had a significant impact on the stock price of the two companies under consideration. Thus, defining the null hypothesis  $H_0$ , we assume that, taking into account the general market conditions during the event period, the abnormal returns follow a multivariate normal distribution, with conditional mean equal to zero and conditional variance  $\sigma^2(AR_{it})$ , which can be represented as follows:



$$\widehat{AR}_{it} \sim N(0, \sigma^2(AR_{it}))$$

Where

$$\hat{\sigma}^2(AR_{it}) = \sigma_{\varepsilon_i}^2 + \frac{1}{L_1} \left[ 1 + \frac{(R_{mt} - \hat{\mu}_m)^2}{\hat{\sigma}_m^2} \right]$$

Essentially, under this assumption, it is posited that the event does not affect the returns in terms of both mean and variance. As can be seen, the variance formula has two components. The first  $\sigma_{\varepsilon_i}^2$  represents the error variance or disturbance term, that is, the intrinsic variability of returns that the market model fails to capture. Conversely, the second component represents the additional variance due to uncertainty in the estimation of the  $\alpha_i$  and  $\beta_i$  parameters. As the number of observations in  $L_1$  increases, this second component tends to zero and can be ignored in the calculation.

To test the validity of the hypothesis, it is necessary to construct test statistics to test whether the anomalous returns significantly differ from zero. An initial approach would be to standardize the anomalous returns by creating the test statistic  $\widehat{SAR}_{it}$ , which relates the anomalous return  $AR_{it}$  and its standard error  $\sqrt{\hat{\sigma}^2(AR_{it})}$ . Asymptotically, this statistic follows a standard normal distribution:

$$\widehat{SAR}_{it} = \frac{\widehat{AR}_{it}}{\sqrt{\hat{\sigma}^2(AR_{it})}} \sim N(0,1)$$

If the absolute value of  $\widehat{SAR}_{it}$  is large, it means that the abnormal return is significantly different from zero, suggesting that there could be a significant impact of the event on the return of security  $i$  at day  $t$ , thus rejecting the null hypothesis.

Since there is likely to be considerable variation in returns on days within the event window, the observations of abnormal returns need to be aggregated so that more complete overall conclusions could be drawn about the events under this study. Indeed, Cumulative Abnormal Returns provide a more robust assessment of the overall impact of the event compared to simple Abnormal Returns. Aggregation can occur along two

dimensions: over time and across stocks. In the specific case, aggregation occurs over time, as the aim is to analyze each stock individually.

The Cumulative Abnormal Return is calculated over the multi-period event window by summing the abnormal returns over the period between  $t_1 = T_1 + 1$  and  $t_2 = T_2$ . In mathematical terms, the formula is:

$$\widehat{CAR}_i(t_1, t_2) = \sum_{t=t_1}^{t_2} \widehat{AR}_{it}$$

In this specific case, an event time interval of 21 days was used for all three events analyzed.

The variance of  $CAR_i$ , as the number of observations contained in  $L_1$  increases, can be defined as:

$$\hat{\sigma}_i^2(t_1, t_2) = (t_2 - t_1 + 1)\hat{\sigma}^2(AR_{it})$$

Under the null hypothesis,  $H_0$ , as seen for abnormal returns, the cumulative abnormal returns follow a multivariate normal distribution, with conditional mean equal to zero and conditional variance  $\sigma_i^2(t_1, t_2)$ :

$$CAR_i(t_1, t_2) \sim N(0, \sigma_i^2(t_1, t_2))$$

Therefore, to test this hypothesis, it is possible to construct a test statistic for the cumulative anomalous returns, which will again be distributed according to a standard normal distribution:

$$SCAR_i(t_1, t_2) = \frac{\widehat{CAR}_i(t_1, t_2)}{\sqrt{\hat{\sigma}_i^2(t_1, t_2)}} \sim N(0,1)$$

To test the acceptance or rejection of the null hypothesis,  $H_0$ , it is necessary to compare the test statistic  $SCAR_i(t_1, t_2)$  with the critical value of the standard normal distribution. To calculate it, it is necessary to define a significance level  $\alpha$ , which in this study is set at 0.05, and, since a two-sided test will be performed, it is generally denoted  $Z_{\alpha/2}$ . Having

a suitably high number of observations, the corresponding critical value is approximately  $\pm 1.96$ .

Defining the hypotheses:

- **Null hypothesis,  $H_0$ :** The event has no significant impact on returns, implying that cumulative abnormal returns are zero or that there is no difference from what would be expected under normal conditions. In statistical terms, this would mean assuming that  $E(CAR_i(t_1, t_2)) = 0$
- **Alternative hypothesis,  $H_1$ :** The event has a significant impact on returns, implying that the cumulative abnormal returns are different from zero or that there is a significant difference from what would be expected under normal conditions. In statistical terms, this would mean assuming that  $E(CAR_i(t_1, t_2)) \neq 0$

it can be stated that:

- If  $-Z_{\alpha/2} < |\widehat{SCAR}_i(t_1, t_2)| < Z_{\alpha/2}$  : the value of the test statistic falls within the Acceptance Region, indicating that there is not enough evidence to reject the null hypothesis  $H_0$ , which is, therefore, accepted.
- If  $|\widehat{SCAR}_i(t_1, t_2)| > |Z_{\alpha/2}|$  : the value of the test statistic falls within the Critical Region, rejecting the null hypothesis and accepting the alternative hypothesis  $H_1$ .

Visually, this can be represented as in Figure 1.




Figure 1 - Standardized Normal Distribution

A study like this allows for testing of market efficiency. Indeed, by using statistical tests to evaluate the null hypothesis, we can infer whether market inefficiencies are present. Acceptance of the null hypothesis might indicate greater market efficiency, since, being the abnormal returns not significantly different from zero, prices accurately reflect all available information. If, on the contrary, the null hypothesis is rejected, we may be facing market inefficiency and, therefore, abnormal profit opportunities, since the abnormal returns are significantly different from zero.

### **3.4 Results**

In this section, the results obtained from the three event studies conducted are presented, two for Microsoft and one for Nvidia, aimed at assessing the impact of announcements of investments in artificial intelligence and related products on stock prices. Tables will be provided to illustrate the results obtained from hypothesis testing, along with interpretations and explanations of these results. For simplicity, we will analyze the data obtained for each of the two companies separately.

#### **3.4.1 Nvidia**

As previously introduced, on January 8, 2024, Nvidia announced significant investments in core Tensor GPUs, LLMs, and PC tools, with the goal of democratizing generative AI for millions of users. In the event study conducted, the impact of this announcement on Nvidia's stock price performance was examined, considering 162 observations in the estimation window and 21 in the event window. For a complete evaluation, the event study was repeated twice, using two different market indexes, the Nasdaq Composite and the S&P 500, to derive the market model (Appendix, Table A and B). The results are summarized in the following tables: Table 1 and Table 2.

Table 1

*Impact of the announcement on Nvidia's abnormal returns (AR) and cumulative abnormal returns (CAR), using the Nasdaq Composite as the market benchmark*

Panel A: Abnormal returns and test statistics in the event window -10, 0, +10

T-test	[-10]	[-5]	[-1]	[0]	[1]	[5]	[10]
AR	-0.704%	0.955%	1.994%	1.914%	1.415%	3.293%	-0.548%
SAR	-0.322	0.455	0.911	0.875	0.647	1.505	-0.251

Panel B: Cumulative abnormal returns and test statistics for given event windows

T-test	[-10, 10]	[-5, 5]	[-2, 2]
CAR	11.73%	13.98%	7.96%
SCAR	1.143	1.845	1.481

Table 2

*Impact of the announcement on Nvidia's abnormal returns (AR) and cumulative abnormal returns (CAR), using the S&P 500 as the market benchmark*

Panel A: Abnormal returns and test statistics in the event window -10, 0, +10

T-test	[-10]	[-5]	[-1]	[0]	[1]	[5]	[10]
AR	-0.543%	0.454%	1.720%	3.069%	1.843%	3.654%	-0.416%
SAR	-0.222	0.185	0.703	1.255	0.753	1.493	-0.170

Panel B: Cumulative abnormal returns and test statistics for given event windows

T-test	[-10, 10]	[-5, 5]	[-2, 2]
CAR	11.14%	12.04%	8.99%
SCAR	0.971	1.420	1.501

Panel A of both tables displays the abnormal returns (AR) and the respective test statistics (SAR) obtained before and after the announcement in the event window. The results refer to different days in the event window: with [0] representing the announcement day, [-10], [-5], and [-1] denote, respectively, ten days before, five days before, and the day before

the announcement; while [1], [5], and [10] represent the day after, five days after, and ten days after the announcement, respectively.

Panel B, on the other hand, reports the results for the cumulative abnormal returns and related test statistics considering three time intervals within the event window: [-10, 10], [-5, 5] and [-2, 2] denote intervals of 21 days, 11 days and 5 days, respectively.

Referring to Table 1, the analysis of Nvidia's abnormal returns, considering the Nasdaq as the market index, reveals that the announcement had a generally positive impact on stock performance, especially on the days close to the announcement. However, the corresponding test statistics (SARs) fail to reach statistically significant levels, indicating a lack of robustness in the observed positive effect. In fact, the SAR values are significantly lower than the critical value of 1.96, at a 5% significance level. The most pronounced effect appears on the fifth day post-announcement, both in terms of abnormal returns and test statistics, suggesting an amplification of the announcement's effects in the days following the vent day. However, the value of the test statistic (1.505) does not exceed the critical value (1.96), suggesting, once again, that there is not a level of statistical significance sufficient to reject the null hypothesis. The effect of the announcement on abnormal returns is illustrated in Figure 2.




Figure 2 - Abnormal Return (event window), Nasdaq Composite

To test the significance of these results more accurately, it is necessary to analyze the results for the cumulative abnormal returns (CAR) and their test statistics over the three time windows considered. The analysis, in this case as well, suggests that although the cumulative abnormal returns are positive for all time windows considered, they are not statistically significant. Specifically, the widest interval [-10, 10] shows a high cumulative abnormal return, yet the corresponding SCAR test statistic is 0.971, showing a lack of statistical significance. The result holds true for the other two shorter intervals, [-5, 5] and [-2, 2], which, despite showing a positive CAR and higher values of the test statistics, fail to reach the 5% significance level. The effect of the announcement on cumulative abnormal returns is illustrated in Figure 3.

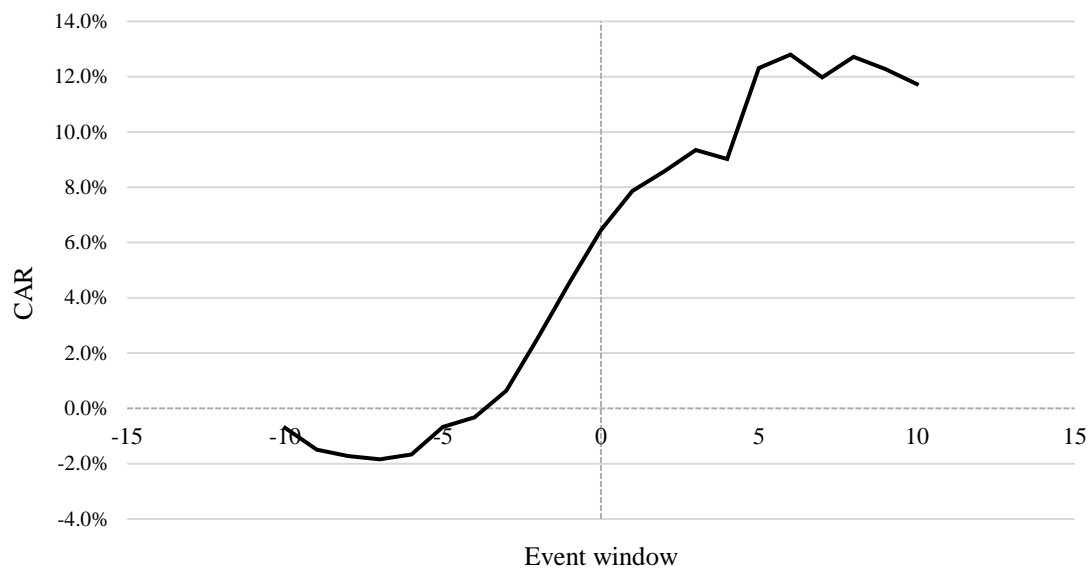


Figure 3 - Cumulative Abnormal Return (event window), Nasdaq Composite

In light of this analysis, we can state that the results obtained in the event window, including the day of the announcement itself, do not have a level of statistical significance that allows us to reject the null hypothesis  $H_0$ . Consequently, we can state that the event has no significant impact on returns and that, in expectation, the simple and cumulative abnormal returns are not significantly different from zero. This absence of significant changes in abnormal returns before, during, and after the event suggests that the market

has already incorporated the information about Nvidia's announcement into stock prices, demonstrating a certain degree of efficiency in the market price formation process.

To test whether significant differences exist in the results when considering a different market index, the event was repeated using the S&P 500 as the market benchmark for estimating expected returns. The results are outlined in Table 2.

Overall, the results obtained using the S&P 500 are in line with those previously discussed, indicating no compelling evidence to reject the null hypothesis  $H_0$  across any of the examined intervals or individual days, either for simple or cumulative expected returns. However, a comparison between the two indexes showed some disparities in the percentages of simple and cumulative abnormal returns. For instance, on the event day, the abnormal return was 3.069% with the S&P 500 compared to 1.914% with the Nasdaq Composite. Additionally, slight differences were observed in the respective test statistics. Regarding SARs, the Nasdaq Co. tends to exhibit a higher value on the fifth day post-event, while for SCARs, the Nasdaq Co. shows a higher value in the period  $[-5, 5]$ , although no significant differences are noted in the other two intervals. The impact of the announcement on simple and cumulative abnormal returns, considering the S&P 500, is illustrated in Figures 4 and 5.




Figure 4 - Abnormal Return (event window), S&P 500






Figure 5 - Cumulative Abnormal Return (event window), Nvidia, S&P 500

### 3.4.2 Microsoft

The analysis performed on Microsoft focuses on studying the effects of two events: the announcement of a significant investment in new generative artificial intelligence solutions, for convenience, will be referred to as *Event 1*, and the announcement of a \$2.2 billion investment to support Malaysia in its cloud and artificial intelligence opening process, referred to as *Event 2*. As already seen for Nvidia, the study of each event has a double run, considering first the Nasdaq Composite and then the S&P 500 as the market index (Appendix, Table A and B).

#### Event 1

On January 11, 2024, Microsoft announced a significant investment to release new generative artificial intelligence and data management solutions aimed at retailers. Among the proposals are new copilot models on *Azure OpenAI Service* and *Microsoft Dynamics 365 Customer Insight* and the launch of *Retail Media Creative Studio* in the *Microsoft Retail Media Platform* to improve the customer experience. In the event study conducted, the impact of this announcement on Microsoft's stock price performance was

studied, considering 200 observations in the estimation window and 21 in the event window. The results are summarized in Tables 3 and 4, one for each market index.

*Table 3*

*Impact of the announcement on Microsoft's abnormal returns (AR) and cumulative abnormal returns (CAR), using the Nasdaq Composite as the market benchmark*

Panel A: Abnormal returns and test statistics in the event window -10, 0, +10

T-test	[-10]	[-5]	[-1]	[0]	[1]	[5]	[10]
AR	-0.378%	-0.204%	1.024%	0.426%	0.921%	-0.549%	0.074%
SAR	-0.348	-0.188	0.943	0.393	0.849	-0.506	0.068

Panel B: Cumulative abnormal returns and test statistics for given event windows

T-test	[-10, 10]	[-5, 5]	[-2, 2]
CAR	3.84%	1.822%	3.12%
SCAR	0.754	0.485	1.171

*Table 4*

*Impact of the announcement on Microsoft's abnormal returns (AR) and cumulative abnormal returns (CAR), using the S&P 500 as the market benchmark*

Panel A: Abnormal returns and test statistics in the event window -10, 0, +10

T-test	[-10]	[-5]	[-1]	[0]	[1]	[5]	[10]
AR	-0.405%	-0.442%	1.136%	0.464%	0.819%	-0.208%	-0.255%
SAR	-0.327	-0.357	0.917	0.375	0.660	-0.168	-0.206

Panel B: Cumulative abnormal returns and test statistics for given event windows

T-test	[-10, 10]	[-5, 5]	[-2, 2]
CAR	2.96%	3.21%	3.55%
SCAR	0.509	0.748	1.170

Panel A of both tables displays the abnormal returns (AR) and the respective test statistics (SAR) obtained before and after the announcement in the event window. The results refer to different days in the event window: with [0] representing the announcement day, [-10], [-5], and [-1] denote, respectively, ten days before, five days before, and the day before the announcement; while [1], [5], and [10] represent the day after, five days after, and ten days after the announcement, respectively.

Panel B, on the other hand, reports the results for the cumulative abnormal returns and related test statistics considering three time intervals within the event window: [-10, 10], [-5, 5] and [-2, 2] denote intervals of 21 days, 11 days and 5 days, respectively.

Panel A of Table 3 shows that in the interval [-10, -5] prior to the announcement Microsoft's stock performance was below expectations, gradually recovering as the announcement day approached. In fact, the day before the announcement recorded an abnormal return of 1.024%, a value significantly higher than that recorded on the event day, which was 0.464%, indicating a possible leak of information in the market before the announcement. The results of the test statistics show that the abnormal returns are not statistically significant at a 5% significance level on any of the days under study. The effect of the event on abnormal returns is illustrated in Figure 6.




Figure 6 - Abnormal Return (event window), Microsoft, Event 1, Nasdaq Composite

These findings are further supported by Panel B of Table 3, where the data on Cumulative Anomalous Returns are displayed. The test statistics obtained across the three intervals all fall within the acceptance region, indicating non-significance at a 5% significance level. The effect of the announcement on the Cumulative Anomalous Returns is illustrated in Figure 7.




Figure 7 - Cumulative Abnormal Return (event window), Microsoft, Event 1, Nasdaq Composite

To test for significant differences in the results when considering a different market index, the event was repeated using the S&P 500 as the market benchmark for estimating expected returns. The results obtained are summarized in Table 4.

Overall, with the S&P 500 as the market proxy, positive returns are more evident on days near the event,  $[-1, 0, 1]$ , compared to those obtained using the Nasdaq Composite index, suggesting a more positive impact of the event on stock prices. However, the change of the index does not alter the significance of the results. Indeed, for the S&P 500, none of the test statistics (SAR) exhibit statistically significant values, as all are strictly below 1.96 and above -1.96. Regarding the cumulative abnormal returns, using the S&P 500, lower data are recorded compared to those obtained using the Nasdaq Composite in the time window  $[-10, 10]$ , while in the other two intervals,  $[-5, 5]$  and  $[-2, 2]$ , the roles are reversed. The test statistics (SCAR) show no significant differences between the two

indices, confirming the absence of statistical significance in the results even with the S&P 500 index. Considering this index, the effect of the announcement on simple and cumulative abnormal returns is illustrated in Figures 8 and 9.




Figure 8 - Abnormal Return (event window), Microsoft, Event 1, S&P 500




Figure 9 - Cumulative Abnormal Return (event window), Microsoft, Event 1, S&P 500

In conclusion, it can be stated that the event analyzed doesn't have a significant impact on Microsoft's stock prices, regardless of the market index used as a market proxy. This, from a market efficiency point of view, may be due to the fact that the market had already incorporated the relevant information prior to the announcement itself.

## Event 2

On May 2, 2024, Microsoft announced its commitment to investing \$2.2 billion over the next four years to drive Malaysia towards digital transformation. This investment involves the development of new cloud and artificial intelligence infrastructure, enhancing the collaboration with the Government of Malaysia to establish a *National Center of Excellence for Artificial Intelligence*, and strengthening the country's cybersecurity. An event study was conducted to evaluate the impact of this announcement on Microsoft's share price performance, considering 200 observations in the estimation window and 21 in the event window. The results are summarized in Tables 5 and 6, one for each market index.

Table 5

*Impact of the announcement on Microsoft's abnormal returns (AR) and cumulative abnormal returns (CAR), using the Nasdaq Composite as the market benchmark*

Panel A: Abnormal returns and test statistics in the event window -10, 0, +10

T-test	[-10]	[-5]	[-1]	[0]	[1]	[5]	[10]
AR	-1.415%	-1.928%	1.699%	-0.680%	0.347%	0.149%	-0.293%
SAR	-1.568	-2.137	1.88	-0.753	0.385	0.166	-0.324

Panel B: Cumulative abnormal returns and test statistics for given event windows

T-test	[-10, 10]	[-5, 5]	[-2, 2]
CAR	-4.19%	-3.26%	0.50%
SCAR	-0.989	-1.042	0.225

Table 6

*Impact of the announcement on Microsoft's abnormal returns (AR) and cumulative abnormal returns (CAR), using the S&P 500 as the market benchmark*

Panel A: Abnormal returns and test statistics in the event window -10, 0, +10							
T-test	[-10]	[-5]	[-1]	[0]	[1]	[5]	[10]
AR	-1.630%	-1.982%	1.797%	-0.341%	0.731%	-0.179%	-0.285%
SAR	-1.673	-2.035	1.845	-0.350	0.750	-0.184	-0.292

Panel B: Cumulative abnormal returns and test statistics for given event windows			
T-test	[-10, 10]	[-5, 5]	[-2, 2]
CAR	-4.34%	-2.70%	1.19%
SCAR	-0.950	-0.801	0.497

Panel A of Table 5 shows that in the days before the announcement, significantly negative abnormal returns are observed: - 1.415% and - 1.928% on the tenth and fifth days before the announcement, respectively. This suggests a performance diverging from market expectations, potentially anticipating a negative reaction to the content of the announcement. However, on the day immediately preceding the event, abnormal returns return to being positively divergent. Nevertheless, on the announcement day, the abnormal return turns negative again (- 0.680%), showing an unfavorable reaction to the announcement. In the following days, small positive abnormal returns are recorded, indicating a possible market correction or a slightly delayed positive reaction. As for the values of the test statistics, they do not exceed the critical value  $\pm 1.96$  on any of the days analyzed, except for day -5, where the value exceeds the critical threshold ( $-2.137 < -1.96$ ) and is, therefore, statistically significant. This means that, on that specific day, the abnormal returns were non-zero. For a visual representation, the impact of the announcement on abnormal returns is illustrated in Figure 10.

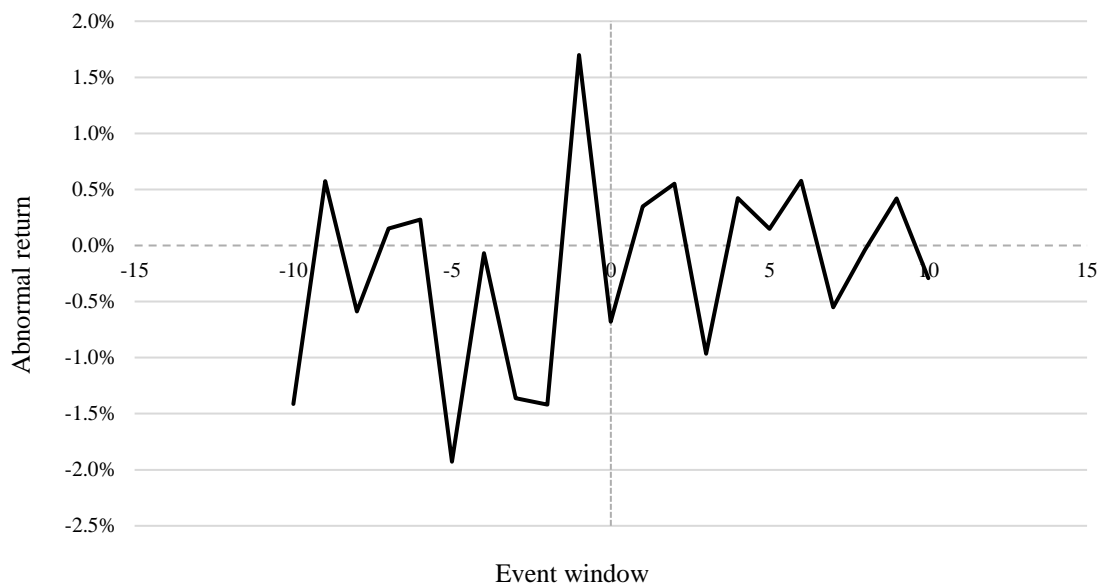


Figure 10 - Abnormal Return (event window), Microsoft, Event 2, Nasdaq Composite

The analysis of cumulative abnormal returns (panel B), indeed, confirms Microsoft’s performance significantly below market expectations in the [-10, 10] interval, showing a value of - 4.19%. The negative trend even in the narrower window [-5, 5], with an outlier performance of - 3.26%. However, the market seems to correct itself in the [-2, 2] interval, showing a slightly positive CAR of 0.50%, indicating that the days around the announcement had a marginal but positive impact on stock returns. The values of the test statistics, on the other hand, show two negative values in the broader intervals and one positive value in the [-2, 2] interval. Nevertheless, none of these values reach the significance level in any of the three intervals, thus enabling the acceptance of the null hypothesis and the rejection of the alternative one. The effect of the announcement on cumulative abnormal returns is illustrated in Figure 11.



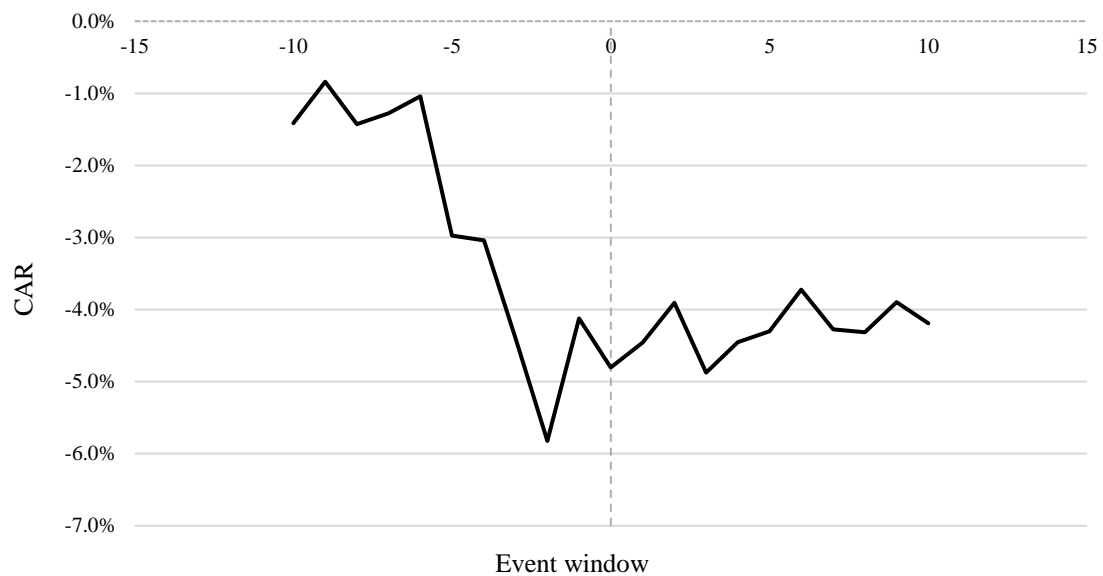


Figure 11 - Cumulative Abnormal Return (event window), Microsoft, Event 2, Nasdaq Composite

Turning to Table 6, it is evident that, using the S&P 500, the abnormal returns are even more negative than those in Table 5, especially on days - 10 and - 5, further reinforcing the hypothesis of an early negative reaction to the content of the announcement. On the event day, however, the abnormal return with the S&P 500 is nearly half of that resulting from using the Nasdaq Composite, though it remains negative. In the post-announcement days, Table 6 shows a relatively strong recovery on the day immediately following the announcement, while in Table 5 the recovery is less obvious and more gradual. As for the test statistics (SAR), there are no statistically significant values, except for the one recorded on day - 5. Indeed, the corresponding test statistic value is -2.035 and, therefore, falls within the rejection region of the distribution, making the abnormal return statistically significant. Being, however, an isolated case, it is likely that the price movement is due to random market fluctuations or other variables rather than the announcement itself.

From the results in Panel B of Table 6 it can be seen that, regardless of the index used, CARs are still negative in the first two intervals. For the [-2, 2] interval, the CAR, resulting from using the S&P 500, is significantly more positive than the one obtained in Table 5 (1.19% versus 0.50%). Test statistics (SCAR) are not statically significant across

all time windows, which nullifies the statistical significance of the abnormal result recorded on day -5. A graphical representation of the impact of the announcement on simple and cumulative abnormal returns, considering the S&P 500 as a proxy for the market, is shown in Figures 12 and 13.

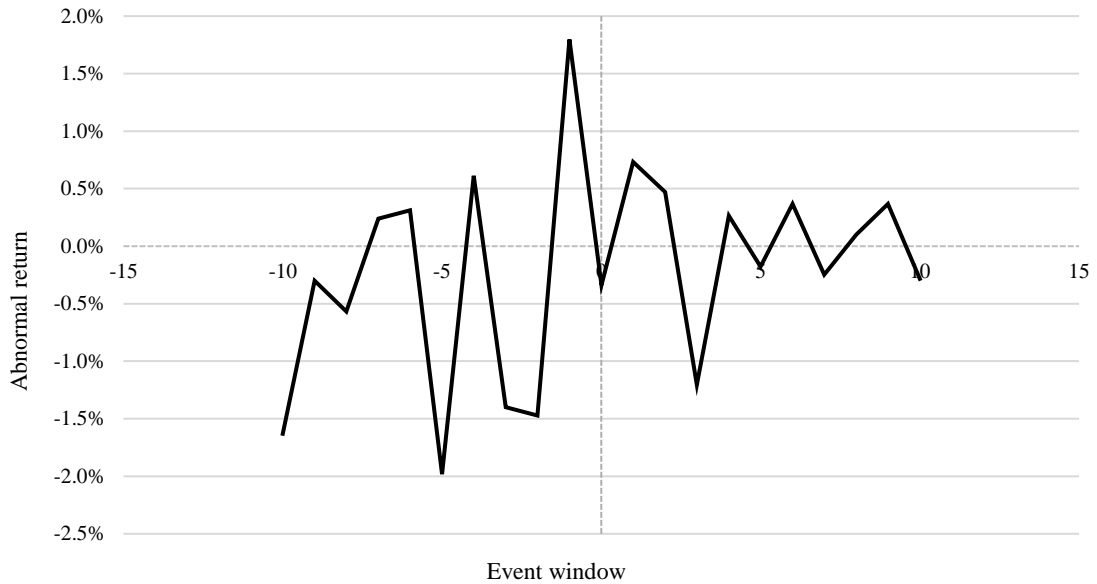


Figure 12 - Abnormal Return (event window), Microsoft, Event 2, S&P 500




Figure 13 - Cumulative Abnormal Return (event window), Microsoft, Event 2, S&P 500

In general, these results suggest that, although there have been changes in daily returns, these are not consistent enough to be considered statistically significant. In addition, the strong negativity of the cumulative abnormal returns indicates that the announcement may not have been positively received by market participants, or that other unrelated events may have negatively affected the market during the same time period, an uncertainty that, unfortunately, cannot be confirmed in this study.

### **3.5 Discussion and limitations**

The theory of efficient markets suggests that there should be no expected abnormal returns following an event, such as the AI investment announcements by Microsoft and Nvidia. As observed in the six tables above, simple and cumulative daily abnormal returns tend to vary their sign from day to day. This behavior is consistent with the predictions of efficient markets theory, which asserts that stock prices immediately reflect and incorporate all available new information, resulting in random and unpredictable changes in returns. Although the ARs recorded a daily change over the time interval considered, indicating a persistent deviation from expected returns, statistical tests conducted showed an absence of statistical significance of these results at a significance level of 5%. This supports the null hypothesis,  $H_0$ , validating market efficiency.

The three events studied had different impacts on the market. Nvidia's announcement was received very positively by the market, as the values of CARs show in Tables 1 and 2. On the other hand, Microsoft's announcements produced different results: the first event generated a fairly positive reaction, falling short, however, of the levels of Nvidia's announcement; while the second event was not positively received by the market, as shown by the CARs in Tables 5 and 6. Indeed, they were negative in most of the analyzed intervals and only showed a weak recovery in the interval closest to the event. These results are illustrated in Figures 14 and 15.

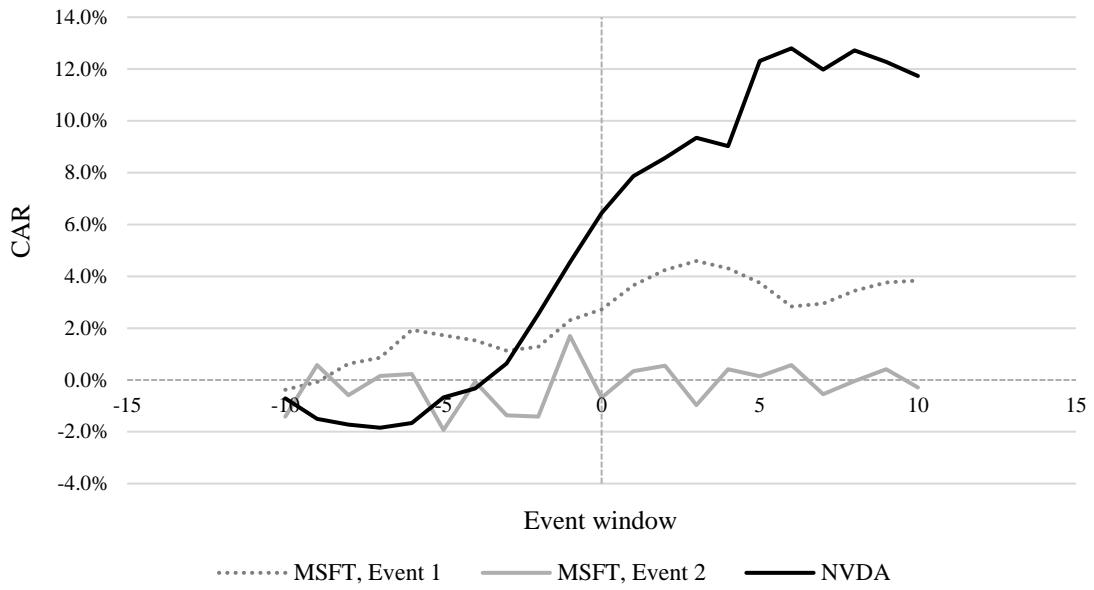


Figure 14 - Cumulative Abnormal Return (event window), Nasdaq Composite



Figure 15 - Cumulative Abnormal Return (event window), S&P 500

The model used in conducting this study, however, does come with limitations. In fact, in testing the validity of the hypotheses, test statistics, such as  $t_{CAR}$ , were employed; the use of which presupposes, however, knowledge of its statistical distribution. The approximation used in this study is the one proposed by MacKinlay, which states that, under the null hypothesis,  $t_{CAR}$  follows a standard normal distribution:

$$t_{CAR} \sim N(0, 1)$$

This approximation is based on the assumptions that (1) the data are independent and identically distributed following a distribution with zero mean and variance  $\sigma^2$ , which is not known a priori, and that (2) the sample size tends to infinity. In general, as long as the number of observations, in this case those contained in the estimation window, is greater than 100, the approximation holds, as the number of degrees of freedom rises, the  $t$  distribution tends to a standard normal distribution. The main limitation comes from the fact that these tests tend to be used for smaller time intervals, such as the event window, and this is a very strong assumption.

Moreover, the study does not allow for control over whether other variables are present; in fact, the change in stock price could be influenced not only by the company-specific announcement but also by other external factors not considered in the model. This may also be due to the fact that the results do not reflect the general market dynamics because the event studies were conducted separately for each announcement, presenting an aggregation of abnormal returns over time instead of across securities.

## CONCLUSIONS

The studies provide supported evidence that artificial intelligence investment announcements do not cause statistically significant abnormal returns for the analyzed companies' shares, confirming market efficiency and indicating that these announcements haven't had a significant impact on the share price.

The only statistically significant deviation (event 2) observed lasts for just one day, demonstrating that, although there is evidence of anomalies in stock pricing, these are not a reliable method from which to derive greater profits, since they are immediately exploited and do not persist over time.

In addition, the results show that the market tends to perceive the investment strategies of the two companies differently. Nvidia's strategy appears to be more solid in the eyes of investors, who reacted positively to the company's announcement, compared to Microsoft's, which produces mixed results. Of course, market reactions may depend on other variables not directly considered in the study, such as political events or macroeconomic factors, highlighting a limitation of the study.

Ultimately, this study has contributed to the understanding of stock market dynamics in response to Artificial Intelligence investment announcements, albeit limited to two companies, and has provided a basis for further research and investigation in the field.

## APPENDIX

**Table A**

*Summary output of Market Model regression, using Nasdaq Composite as market proxy*

<b>Regression Statistics</b>	<b>NVDA</b>	<b>MSFT, ev. 1</b>	<b>MSFT, ev. 2</b>
Multiple R	0.664	0.687	0.712
R Square	0.440	0.471	0.506
Adjusted R square	0.437	0.469	0.503
Standard Error	0.022	0.109	0.009
Observations	162	200	200
<b>Coefficients</b>			
Alpha	0.0009	0.0005	0.0004
Beta	1.9426	1.0141	0.9164

**Table B**

*Summary output of Market Model regression, using S&P 500 as market proxy*

<b>Regression Statistics</b>	<b>NVDA</b>	<b>MSFT, ev. 1</b>	<b>MSFT, ev. 2</b>
Multiple R	0.548	0.558	0.652
R Square	0.300	0.311	0.424
Adjusted R square	0.296	0.308	0.422
Standard Error	0.024	0.012	0.010
Observations	162	200	200
<b>Coefficients</b>			
Alpha	0.0015	0.0009	0.0003
Beta	2.1439	1.0831	1.1483

## REFERENCES

- A. L. Samuel. (1959). *Some Studies in Machine Learning Using the Game of Checkers*, IBM Journal of Research and Development, vol. 3, no. 3, pp. 210-229.
- Babina T., Fedyk A., He A., Hodson J. (2024), *Artificial intelligence, firm growth, and product innovation*, Journal of Financial Economics, Volume 151.
- Beresford R., Agatonovic-Kustrin S. (2000), *Basic concepts of artificial neural network (ANN) modeling and its application in pharmaceutical research*. Journal of Pharmaceutical and Biomedical Analysis, Volume 22, Issue 5, Pages 717-727.
- Cohen, M. K., M. Hutter, and M. A. Osborne (2022). *Advanced artificial agents intervene in the provision of reward*. AI Magazine 43: 282–93.
- Fama, E. F. (1965). *The Behavior of Stock-Market Prices*. The Journal of Business, 38(1), 34–105.
- Fama, E. F., Fisher, L., Jensen, M. C., & Roll, R. (1969). *The Adjustment of Stock Prices to New Information*. International Economic Review, 10(1), 1–21.
- Fama, E. F. (1970). *Efficient Capital Markets: A Review of Theory and Empirical Work*. The Journal of Finance, 25(2), 383–417.
- Firth, M. (1977). *The Valuation of Shares and the Efficient-Markets Theory*. *Studies in Finance and Accounting*, 117–139. Palgrave, London.
- Jensen, M. C. (1969). *Risk, The Pricing of Capital Assets, and The Evaluation of Investment Portfolios*. The Journal of Business, 42(2), 167–247.
- Kahneman, D. and Riepe, M. (1998). “*Aspects of investor psychology*”, Journal of Portfolio Management 24, pp.52-65.
- Luce, L. (2019). *Artificial intelligence for fashion: How AI is revolutionizing the fashion industry*. Berkeley, CA: Apress.
- Lui, A.K.H., Lee, M.C.M. & Ngai, E.W.T. (2022). *Impact of artificial intelligence investment on firm value*. Ann Oper Res 308, 373–388.



- MacKinlay, A. C. (1997). *Event Studies in Economics and Finance*. Journal of Economic Literature, 35(1), 13–39.
- Maloney, Michael T. and Mulherin, J. Harold (1998). *The Stock Price Reaction to the Challenger Crash: Information Disclosure in an Efficient Market*.
- McCarthy, J., Minsky, M., Rochester, N., Shannon, (1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. AI Magazine Volume 27 Number 4 (2006).
- Purdy M., Daugherty P., (2017). *AI Research: How AI boosts industry profits and innovation*, Accenture.
- Ranganatham M., Subramanian V. (1993). *Weak Form of Efficient Markets Hypothesis: A Spectral Analytic Investigation*. Vol.18, No.2, April-June 1993
- Russell, Stuart J., (2010). *Artificial intelligence: a modern approach*. Upper Saddle River, N.J.: Prentice Hall.
- Simon, H.A. (1982). “*Models of Bounded Rationality*”. Vol. 2, Behavioral Economics and Business Organization, Cambridge: The MIT Press.
- Statista, *Artificial Intelligence: in-depth market analysis*.
- Wei J., Junbo W., Baozhong Y., (2021). *From Man vs. Machine to Man + Machine: The Art and AI of Stock Analyses*, National Bureau of Economic Research Working Paper Series No. 28800.
- Wu Yu-chen, Feng Jun-wen, (2018). *Development and Application of Artificial Neural Network*. Wireless Pers Commun 102:1645–1656.
- The European Commission’s High-Level Expert Group on Artificial Intelligence. *A Definition of Ai: Main Capabilities and Scientific Disciplines*, Brussels, 18 December 2018.
- <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>
- <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>
- <https://www.edge.org/response-detail/26084>

<https://www.ibm.com/blog/supervised-vs-unsupervised-learning/>

<https://www.ibm.com/it-it/topics/deep-learning>

<https://www.oracle.com/vn/artificial-intelligence/what-is-natural-language-processing/>

<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review>

<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year>

<https://newsroom.accenture.com/news/2017/accenture-report-artificial-intelligence-has-potential-to-increase-corporate-profitability-in-16-industries-by-an-average-of-38-percent-by-2035>

<https://www.goldmansachs.com/intelligence/pages/how-much-could-ai-boost-us-stocks.html>

<https://www.goldmansachs.com/intelligence/pages/why-ai-stocks-arent-in-a-bubble.html>

<https://www.eetimes.eu/nvidias-michael-kagan-building-on-ais-iphone-moment-to-architect-data-processings-future/>

<https://news.microsoft.com/announcement/microsoft-is-born/>

<https://www.microsoft.com/en-us/about/values>

<https://news.microsoft.com/2016/09/29/microsoft-expands-artificial-intelligence-ai-efforts-with-creation-of-new-microsoft-ai-and-research-group/>

<https://news.microsoft.com/apac/2024/05/02/microsoft-announces-us2-2-billion-investment-to-fuel-malysias-cloud-and-ai-transformation/>

<https://news.microsoft.com/2024/01/11/microsoft-unveils-new-generative-ai-and-data-solutions-across-the-shopper-journey-offering-copilot-experiences-through-microsoft-cloud-for-retail/>

<https://nvidianews.nvidia.com/news/generative-ai-rtx-pcs-and-workstations>