Detecting Manipulation in the Stock Market: Historical Overview and Comparative Analysis of Traditional Methods and Artificial Intelligence

Strategic Management and Competitiveness

National and Kapodistrian University of Athens, University of Thesally

Abstract Excerpt

In this paper we aim to showcase the need for an augmented detection approach where laws, regulations, qualitative and quantitative methods, and artificial intelligence are interconnected, taking into consideration the dynamics of the market, of the attempted fraud types, and of the detection methods, as they evolved since the very first incident.

Supervisor: Ioannis Leventidis

Author: Ntanou Eleni, 734118200053

February 2025

Contents

Table of figures

Detecting Manipulation in the Stock Market: Historical Overview and Comparative Analysis of Traditional Methods and Artificial Intelligence

Acknowledgements

I would like to express my deepest appreciation to Dr Ioannis Leventidis, professor of Mathematics in the Department of Economics of the National and Kapodistrian University of Athens, whose belief in me has kept my motivation high during this process.

I am also thankful to Dr Carayannis, professor of Entrepreneurship and Innovation in the National and Kapodistrian University of Athens, whose lectures and assignments helped me clarify my academic interests and navigate through resources and databases.

I would like to extend my sincere thanks to Dr Psalidopoulos and Dr Chortareas, for their patience and inspiration on the writing process during my initial postgraduate stages.

I would finally like to thank my friend, fellow postgraduate student, and PhD candidate, Manolioudaki Chariklia, for the moral support and the practical tips she has offered me at times of self-doubt.

Abstract

Stock market manipulation is a grave issue that undermines the integrity of the markets as well as investors' trust. Despite regulatory effort, manipulation continues to evolve, a fact that makes it essential to develop effective detection methods. This thesis aims to provide a comprehensive understanding of stock market manipulation by exploring its definitions, types, and detection methods. After manipulation definitions are provided, there is an overview of stock market manipulation types that are classified and commented on under a selection of criteria, such as the suitability of qualitative or quantitative methods for its detection, or the necessity for Order Book access. Then follows a brief historical overview which showcases manipulative instances and the response from the authorities and the scientists. We read about the role of regulations, the evolution of statistics and its implication in the stock market, and about the birth of artificially intelligent systems. Afterwards, we proceed to compare the traditional methods of statistics and econometrics to AI-augmented monitoring systems, and then we explore an illustration of both approaches in a programme that was created in python, which employs and ARIMA and a Random Forest model and uses the Gamestop stock and data from a short-squeezing scheme, focusing on the events of the first half of 2021. Lastly we read about the benefits and limitations of artificial intelligence in the field of stock market monitoring, and, after mentioning several concerns, we arrive at the conclusions. Here, we deduce the complimentary nature of statistics, econometrics and artificial intelligence, the inevitability of AI adoption in market surveillance, and the seriousness of the concerns regarding competition in the field, including the possible necessity of government support schemes to safeguard it.

Chapter 1 – Introduction

1.1 Importance and objectives

The detection of manipulative incidents in the stock market has been a tandalising pursuit and one not so recent endeavour. Each era's approach was characterised by different aspects; firstly the necessity for institutional regulations became evident, then quantitative methods were incorporated in the process, later data driven approaches came into spotlight. With the rise of artificial intelligence though, and with its incorporation in the surveillance procedures, old methods became more polished and more effective, and new concerns arose. In this paper we aim to showcase the need for an augmented detection approach where laws, regulations, qualitative and quantitative methods, and artificial intelligence are interconnected, taking into consideration the dynamics of the market, of the attempted fraud types, and of the detection methods, as they evolved since the very first incident. The historical overview is not the main objective, and thus it is not presented extensively, but several turning points are included to show the evolution of the phenomenon through time and get a better insight of its dynamics.

One of the objectives of comparative analysis, and the most influential in writing this paper, is to help understand the risks of a selection of options, and assist in the decision-making process. According to Pickvance, before a comparative analysis can begin, an "accurate, descriptive grasp of the specificities of cases" is a pre-requisitive, and, consequently, historical context and manipulation types must be presented before the comparison itself appears. The manipulation types themselves have many variations and sub-types, and for this reason in their detailed presentation a few individual aspects that are crucial during detection should be highlighted, thus this part will follow a variation-finding comparison rationale. Lastly, the overall approach of the paper resembles a SWOT analysis, as each manipulation detection method is analysed with strengths, weaknesses, opportunities, and threats in mind.

The attempts to manipulate the stock market to one's own benefit have been around for a few centuries, and they date back to the first stock market bubble and subsequent crash of the Dutch Market during the Tulip Mania craze that took place in the 17th century. Since then, the manipulation methods have evolved and become more nuanced, but at the same direction moves the battle of detecting such fraudulent activities and the restoration of a healthy, safe, and transparent environment that protects the law-abiding investors. The literature is very rich in studies of separate cases, specific manipulation schemes, different approaches of categorisation, and more generalised lay-outs of the phenomenon, as we see in the works of Cumming, Dannhauser and Johan (2015), Aggarwal and Wu (2006), Cumming, Johan and Li (2011), and many more. At the same time, the topic is not of the interest of the scholars alone; the industry and the governments have a deep interest into retaining the stock market at an appropriate state, safeguarding it from illegal practices in order to avoid the costs of its failure and the spill into the real economy with all the well-known devastating consequences. Characteristic examples we see in the work of Dyck, Morse and Zingales, (2010) & (2014), among others. And Academia has responded to this need with notable works on exploration of new detection methods and the evaluation of existing ones, on public policies and the role of constitutions, on laws and punishment, and many more, in an ongoing effort to help heal this great necessity.

This paper will attempt to follow the historic evolution of manipulative incidents and attempts of its detection, focusing on the tools and methods that have been employed through the years. We shall have a taste of institutional regulations, statistics, econometrics, and artificial intelligence, as well as evaluating and comparing their outcomes. We shall also explore an illustrative example of a well-known stock market misconduct that took place in recent years, which will be analysed computationally, both with traditional econometrics and with machine learning, a form of artificial intelligence method. Finally, we shall take into account the concerns of regulatory bodies as well as their proposals.

1.2 Methodology

For the theoretical part of this paper there has been a combination of sources, some of them being printed material, others available online. These sources were both primary and secondary, the former being legislation and regulations from the US, the UK, and the EU supervisory authorities, the rest belonging to the latter type of sources. The historical overview deals with facts from the 17th century onwards, and the material that was taken into consideration ranges from translations of original publications of past centuries to recent works that deal with the manipulation of the stock market in retrospect. For definitions of manipulation, its types, and some notable examples, the material used includes articles from electronic journals, academic databases, regulation announcements, guidelines and directives from the European Union and the United States of

America, as well as specific commissions on transparency and regulation of the market. Facts regarding the evolution of statistics that were later used in finance and subsequently in fraud detection were retrieved mainly from the Royal Statistical Society (RSS) and were enhanced with details from papers that dealt with manipulation in the stock market. For the comparison of traditional and AI-augmented methods information was retrieved by economic forums, banks publications, and the tech community. It should be clarified that the above explanation is indicative of what was the main source for which part of this paper, as in many cases information contained in a document overlapped, and were therefore used in several parts of this paper as additional specifications.

The computational part of the paper was materialised in python, a programming language that, although powerful, it is characterised by relative simplicity in coding and a wealth of libraries with open access. The ARIMA and Random Forest (RF) in the example program were directly retrieved from these libraries, and made use of stock market data regarding the Gamestop stock (ticker name: GME) from yahoo finance. The training of the machine learning model was conducted by python itself with no pre-set initial seed, and it was neither optimised nor fine-tuned in order to offer a better perspective of the raw potentials and pitfalls of AI methods. Employed only once, it still gave better results than the ARIMA in this particular case, even though it should not have, as we will explain in the related chapter. Although we can only speculate as to why or how it managed it, this fact is one of the major concerns of the employment of artificial intelligence in the surveillance process, as we shall read in greater detail in the particular chapter.

Chapter 2 – Definitions and types of manipulation in the stock market

2.1 Definition and categorisations

The term “manipulation of the stock market”, although globally recognised as the distortion of the stock market for mischievous purposes, it does not have one exact global definition. Zulkifley, Munir, Sukor, and Shafiai (2023) hypothesise that this is due to “unclear legal meanings and inconsistent terms that are used in economics and financial literature”. Furthermore, Fischel and Ross (1991) agrees that such term is “ambiguous”, and many others, such as Jarrow (1992), Charian and Jarrow (1995) or Kyle and Viswanathan (2008) provide different definitions as well, some of them regarding manipulation of the market in such a way that it can never be illegal at all. To avoid confusion of the lack of standardisation, we shall accept the definitions of two major regulatory and surveillance bodies, and elaborate from there. According to the United States Securities and Exchange Commission (SEC), stock market manipulation is the engagement in conduct that creates an artificial price for a security, thus interfering with the free and fair operations of the market. In the European Union, Market Abuse Regulation (MAR) regards stock market manipulation to encompass activities intended to distort market prices.

The literature also disagrees as to the taxonomy of manipulation types, and many categorisations have been proposed, again without a universal system. Interestingly enough, private SupTech firms, whose services include supervision and compliance checks and reports, seem to be more consistent, although not identical. After thorough research as to the different manipulation types and detection methods, both in literature and within the private business ecosystem, as well as tech forums and communities, some common denominators were spotted that exhibited particular interest. As a result, although there are several types and sub-types of stock market manipulation, they can be grouped into two major categories, as it is proposed in the next paragraphs.

The first category includes those types where a solely quantitative approach is not sufficient, and the combination of qualitative methods and/or inside knowledge are required; the second category contains the manipulations forms whose detection can be accomplished by relying heavily on quantitative approaches, such as statistics and econometrics.

An alternative division and categorisation of the various manipulation types can be according to whether access to the Order Book is needed or not for their detection. The Order Book is an electronic list which contains buy and sell orders of a specific security, it is organised by price levels, and it contains invaluable information about price and availability, about market depth, which is the number of shares that are being bid on or offered at each price point, and lastly it contains a list of the market participants. Also known as the Continuous Book, it is a dynamic book which is updated in real time throughout a trading day, and it is divided into three parts, the buy orders, the sell orders, and the orders history, but it does not contain the order that specify execution at market open or market close, which are maintained separately in the Opening (Order) Book and the Closing (Order) Book respectively. However, although the end result of the existence of the Order Book is that improves market transparency via the sharing of information for a fee, hidden order, also known as dark pools, which are placed by large players are not available. This omission reduces the utility of the Book to some extent, since there is no way of knowing whether the orders one is viewing are representative of true supply and demand for a stock, and also it makes it more difficult to spot manipulation types that require access to the book data in order to be detected.

On the other hand, if the available information is the typical result of downloading stock market real-time and/or historical data, in other words a dataset including Date, Open, High, Low, Close, and Volume, then quantitative methods could be employed to detect manipulation schemes such as ramping, bull/bear raids, pumps and dumps, wash trades, manipulation of transaction-based fixes, and front-running, without the need to gain access to the Order Book. The analytical techniques which could be applied on the above dataset traditionally include Time Series Analysis, Moving Average Calculation, Volume Analysis, Candlestick Patterns, and Volatility Analysis. Time series analysis involves the study of the price and the trading volume over time, in order to identify trends, patterns, and anomalies. The moving average calculation will help to smooth out price data, so that trends can be identified more clearly. Volume analysis involves the evaluation of the trading volume against price movements in order to spot instances where the high volume does not correspond with the expected price changes. When focusing on candlestick patterns, price data (Open, High, Low, and Close) are used in order to identify specific candlestick patterns like spikes or shooting stars, which may be indicative of manipulation. Lastly, volatility analysis measures the price volatility (standard deviation, average true range) and analyses extreme volatility which does not align with historical patterns.

There are six large families of manipulation types which have distinct characteristics; price manipulation, circular trading, misuse of inside knowledge, price influencing, improper order handling, and misleading conduct. While we explore these manipulation types and their sub-divisions further, we shall indicate whether qualitative approaches are needed to assist the quantitative analysis, and whether access to the order book is necessary for their detection.

2.2 Price manipulation

The first manipulation practice we shall explore is price manipulation, which is an umbrella term that encompasses several sub-types such as spoofing, layering, rambling, pools, cornering the market, squeeze, Bull/Bear trade, and Pumps & Dumps.

The first sub-type of price manipulation, spoofing, which literary means 'to hoax or trick somebody' according to the Oxford Dictionary, is a form of stock market fraud also referred to as crypto currency spoofing where traders place one or more highly visible orders but with no intention to fulfil them, as they cancel the bids right before execution. The purpose for doing this is to artificially move the price of the relevant security by creating a false impression that it is being traded in the market, and ultimately benefit the trader's own position. According to a study by E. J. Lee et al on the Allen and Gale model on transaction-based manipulation (1992) and its more recent modifications by Aggarwal and Wu (2006), the stocks that were targeted for spoofing demonstrated increase volatility in returns, and, at the same time, the market capitalisation, price level and managerial transparency were significantly lower than expected. Nowadays, spoofing can also refer to algorithmic trading activity which traders employ in order to outpace other market participants and manipulate the markets. Quantitative methods alone can detect spoofing fairly easily, as long as there is access to the order book, as statistical analysis of the order book data can detect patterns of placing and cancelling orders. Specific signs to look out for, according to the US Commodity Futures Trading Commission (CFTC), are "manual and automated trading schemes that place and quickly cancel buds and offers in futures contracts in order to benefit other orders and/or positions, orders being quickly placed and cancelled at or near the best bid or offer, especially if opposite-side orders are filled, and multiple orders of the same size repeatedly and simultaneously being placed and cancelled".

The next sub-type is layering, which is a strategy in high-frequency trading (Jiang, Mahoney & Mei, 2005) where a trader places and then cancels orders that they never intended to execute in the first place, in an attempt to fluctuate the stock price. Layering is a variation of spoofing, but rather than a cluster of orders all at the same price point, it uses a series of orders at different price points, and the market price becomes the middle of that spread of orders. For example, in order to buy a stock at a lower price, the trader initially places orders to sell at, or even bellow, the market price. Such a practice lead to Swift Trade's fine of 8 million pounds by British regulators in 2011, making it the first foreign firm to have received a fine in Britain, and also led to the subsequent closure of the Canadian Firm, as we read to the Guardian's news article in January 2013. The detection of such a practice is similar to that of spoofing, where data of the order book are statistically analysed in order of patterns of ordering placing and subsequent cancelling to be spotted.

Another price manipulation technique, which is very popular, is ramping. This refers to cases where a company's share price is driven up in order for one to gain financial advantage. In specific, there is buying and selling of a financial instrument at such a substantial portion of the overall market activity that the usual demand for it is exceeded, resulting in the creation of a trend which the market then follows. Ultimately, the market is influenced in a way that it moves to the trader's desired direction, and he takes advantage of this price movement by selling financial instrument in discussion to make a quick profit. This manipulative practice is usually performed either on instruments with smaller market capitalisation or during periods when trading liquidity is relatively low, because in these instances the manipulation will have a maximum impact. This strategy is also employed in algorithmic trading, where it is also known as momentum ignition. Here, the algorithm detects initial unusual volume and price movement of a certain financial instrument, and then it builds on it by conducting a series of transactions which aim at accelerating the price move, so that the firm can cash in by selling its position. The detection of ramping, similarly to spoofing and layering, requires statistical analysis of the order book data. Analysts explore these data in an attempt to find an increase in trading volume which is followed by significant price changes of a financial instrument. The regulatory landscape is clear about the prohibition of such practice, as the EU Article 12 1.(a).(i) in the Market Abuse Regulation dictates, with the Section 1.6 of the UK MAR to cover this activity in the Great Britain as well, and its counterpart US prohibition against market price ramping which is contained in Section 9(a).(2).(1) of the Securities Exchange Act of

1934. More recently, US FINRA's article to investors on market price ramping was issued in 2023, following a 2022 warning about ramping schemes which were associated with small-cap IPOs.

On a different time horizon, pools, which are dealing rings that collude so as to engage is pre-arranged transactions in order to create false impressions of market activity and drive prices up, usually extend over long time periods and can stretch over months. This type of price manipulation has also drawn much attention of the academia world. Leffler and Farwell very eloquently describe such a practice as follows: "A final device of the pool was artificial market activity. This consisted of a 'heavy churning' of the stock in the market; it was bought and sold by the pool in heavy volume… Its purpose was obvious to all familiar with pool operations. The public must be attracted to the stock; few things attract speculators more quickly than a rising volume. The public's attitude became wetted in anticipation pf 'something big going on'. It rushed in to buy before it was 'too late'. As the stock rose under increased activity, the public entered the market in ever increasing numbers; this was exactly the purpose of the operation". On the other hand, Huebrer and Pratt both argue that pool participants seek out undervalued stocks and purchase them quietly in order to profit from a later price rise. By making use of trading data such as trading volume and price changes, analysts can identify patterns of concentrated buying and selling that artificially inflate or deflate prices of stocks without the need to retrieve data from the order book (Mahoney, 1999), (Jiang, Mahoney & Mei, 2005). According to the works of Allen and Gale (1992) and Aggarwal and Wu (2004), there are possible pooling equilibriums where the price is significantly influenced, but an ordinary investor may be unable to distinguish between manipulators and well-informed investors.

A sub-type of price manipulation that stands out from the rest in this category is a technique known as cornering the market. Typically, one is introduced to this idea with the words of the ancient Greek philosopher, Aristotle, in The Politics, where he narrates of Thales from Miletus: "Thales, so the story goes, because of his poverty was taunted with the uselessness of philosophy; but from his knowledge of astronomy he had observed while it was still winter that there was going to be a large crop of olives, so he raised a small sum of money and paid round deposits for the whole of the olive presses in Miletus and Chios, which he hired at a low rent as nobody was running him up; and when the season arrived, there was a sudden demand for a number of presses at the same time, and by letting them out on what terms he liked he realised a large sum of money, so proving that it is easy for philosophers to be rich if they choose". The same principle is transferred to the stock market. According to this strategy, and as the name infers, he who practices it will acquire as many shares as needed in order to be able to manipulate its price. The common practice would be to buy and stockpile a massive percentage of a certain commodity, usually of firms of a niche industry and especially in the area of future trading, where it is significantly easier to corner the market. The end result is to inflate the price and sell right afterwards, so that a huge amount of profit will be created for this type of manipulators. Of course, nowadays there are laws against anti-monopoly practices and regulators may need to intervene to restore the market, but agreements to corner the market are not always illegal. According to the Federal Trade Commission (FTC) and the US Department of Justice, the agreements that are per se illegal are those which cause either an increasing price or a declining output. Apart from such scenarios, agreements are "analysed under the rule of reason to determine their overall competitive effect. […] Rule of reason analysis focuses on the state of the state of competition with, as compared to without, the relevant agreement". In this potentially manipulative case, pure quantitative data may not be sufficient to conclusively decide whether one is manipulating the market or not, as the statistical methods would be limited to the simulation and comparison of the two market states. However, according to the same guidelines, qualitative insights into the trader's market strategy and his intentions would also be required in order to safely say that a certain trader is engaging in cornering the market. In particular, the Agencies will analyse the nature of such an agreement and inquire the purposes behind it, as well as the market power of the participants, and "where the likelihood of anticompetitive harm is evident from the nature of the agreement […] the Agencies challenge such agreements without a detailed market analysis".

One of the most powerful short-term market catalysts that can aggregate surges in share prices is short-squeeze. Squeezes can happen very quickly and can move the stock more dramatically than a normal rally, with the prices rising in a parabolic way, causing huge losses if someone is not paying close attention. They are often fuelled by unusually high interest in a specific security, and when the price of the stock is driven up, the short sellers rush for the exits all at once, attempting to buy to close their position. This adds fuel to the force, pushing the stock prices even higher, until irrationality sets in. There are several variations of squeeze, with MOASS and Gamma squeezes being the most prominent. MOASS, which stands for Mother Of All Short Squeezes, is a trading strategy in which a very high volume of buyers drives up share stocks that were being shorted by other investors. On the other hand, Gamma squeeze occurs when heavy buying forces the option market participants to buy more shares of the underlying stock to hedge their position. The higher price causes more calls for buying which requires further buying by market participants, creating a snowball effect. The detection of squeeze relies on pure quantitative methods, and no access to order book data is required to spot it.

Another well-recognise form of stock market fraud is bull and bear raid. This practise involves the dissemination of false or misleading information about a security or its issuer which can have a significant impact on the price of the security itself. The names refer to terminology coined in the field of market sentiment, and they are used to communicate the message that investors expect prices to move upwards, described as bull or bullish, or downwards, characterised as bear or bearish. The conduct of bear raid is typically considered a form of securities fraud, and it involves the spread of negative rumours about a targeted firm in order for its stock prices to move downwards. Alternatively, traders may take on large short positions at first, that is, they invest in such a way that they will make profit if the market value of the asset falls, and then they sell at large volumes, thus manipulating the price, usually creating a self-perpetual motion. On the other hand, a bullish market is a market where prices are generally expected to rise, and, based on this assumption, a bull investor buys a security with the intention to sell it later at a higher price. The hard part is to identify the bottom and peak prices and sell in time. A commonly accepted threshold for the start of a bull market is when prices of stocks are rising at least by 20 per cent. Also, during a bull market – which can last for months or years – there is significant increase in the trading volume and in liquidity levels, mainly because there is more demand for securities (which in turn tend to have rising valuations) but at the same time there are fewer investors willing to sell. Additionally, there is certain volatility and several corrections during this period. To spot a bull/bear raid, analysts look for sudden price movements which are correlated with abnormal trading volumes, and the trading data are efficient for the application of quantitative methods without the need for order book data.

Pump and dump is another widespread manipulative technique that involves the artificial inflation of the price of an owned stock through misleading or false positive statements, which is the pump phase, aiming at selling the cheaply purchased stocks at a higher price, which is the dump phase. These schemes are often based on hype and speculation instead of sound investment practices. For example, the sentiment behind a particular stock may not make much sense due to various reasons, like the company might be in the red or have only minimum revenue, but the price of its stock shoots up regardless of these facts. Empirical evidence of the effects of spam e-mail in pump-and-dump schemes can be seen by the works of Bohme and Holz (2006), Frieder and Zittrain (2007), and Hanke and Hauser (2008), where small and illiquid stocks which are traded on Pink Sheets is included as well. This technique may take place over the internet via spam email campaigns, through media channels where fake press releases are posted, or even over the phone, through telemarketing from 'boiler room' brokerage houses. Nowadays, the vast concentration is on media like Telegram and Discord, and, especially in the case of crypto currency pumps and dumps, where the landscape is not as highly regulated as the more conventional counterparts, over the counter transactions (OTC) are a cause of additional worry to regulatory authorities. The SEC has expressed concerns of the increase in OTC, as in this case the two trading parties engage in direct trading without the inclusion of a stock exchange, and thus OTCs are frequently targeted by manipulators, as we also read in the works of Cumming, Dannhauser and Johan (2015) and Aitken, Cumming and Zhan (2015). Again, as with the cases of pools, squeeze, and bull/bear raids, no access to the order room is required for the analysts to spot pump and dump schemes, as the statistical analysis of the trading data alone can reveal patterns of rapid price increases that are followed by sell-offs, being evidenced with trading volumes and price changes.

2.3 Circular trading

A different approach to stock market manipulation is circular trading, as we read in Cumming, Dannhauser and Johan (2015). This manipulation type has its own sub-divisions, with the most prominent being wash trades, churning, compensation trades, and parking, also known as warehousing.

Wash trades can be part of the first stage of a pump and dump scheme, as described above, and they describe a form of manipulation where the same financial instruments are bought and sold in order to create artificial market activity. This creates misleading signals about an asset's demand and, subsequently, on its market health, an act that has the potential to influence other investors' decisions, but on false pretences. Prior to its prohibition in the 1930s, wash trading was a popular way for stock manipulators to falsely signal interest in a stock to pump its value, so that they can earn money by shorting the stock. Its detection can be easily conducted by retrieving data from the order book trade logs and looking for immediate buy and sell orders that cancel one another, trading volumes, and change of ownership.

A similar practice to wash trades is compensation trades, with the difference that in compensation trades there may be more than just one party involved in the creation of artificial market activity. Again the real price and volume of the asset are distorted, and false signals are created to lure other investors. What is more, compensation trades can be a part of a really broad strategy which is split into parts and scatter over a network of practitioners. Subsequently, the detection of such a practice requires more sophisticated approaches. It is imperative to look at the patterns of trades among different parties in order to spot accounts with correlated trading behaviours and timing, as well as conducting cross-market analysis to spot co-ordinated efforts not only among different accounts but also across different platforms.

Another way of conducting circular trading is churning, which refers to the case when a broker engages in excessive buying and selling of securities in a customer's account in order to generate commissions to his own benefit, without considering the customer's investment goals. The main indications of churning at the expense of the broker's customer are unauthorised or frequent trading, and/or high fees, whereas for the analysts who monitor the stock market conducts its main indication is a high level of trading activity that results in very small price movements. The trading data alone can be sufficient for statistical detection of the practice, whereas order book data is needed in order to pinpoint the culprit.

A last well-known form of circular trading as called parking or warehousing. In this case, a party sells shares to another party under the agreement that the original owner of the shares will buy them back after a specific, usually short, period of time. This practice usually takes place near disclosure periods in the case that traders need to reduce their positions in order to appear compliant to certain regulatory guidelines. An outstanding work on the topic is provided by Karpoff, Lee and Martin in the "Consequences to managers for cooking the books". A relatively recent example is a 2014 incident, when two brokers from Wall Street had several securities placed in one another's trading books in an attempt to avoid penalties that would affect their end-of-year bonus, but to no avail, as they were finally exposed. Regulatory authorities uncovered this violation by identifying frequent trades without meaningful change in beneficial ownership, an approach that goes beyond strictly quantitative analysis.

2.4 Misuse of inside knowledge

This category of stock market manipulation is rather self-explanatory to start with, and it can be sub-categorised into two main instances, insider trading and unlawful disclosure of inside information, as we read in the work of Cumming, Dannhauser and Johan (2015).

Insider trading, a straightforward term at first glance, engulfs several different practices in reality. The ones mostly encountered in literature are the Classic Inside Trading, which is the buying and selling of assets based on important non-publically available pieces of information, the Tipper – Tippee Trading that describes an insider who gives others access to confidential information that they use to their advantage when trading, Trading during Black-Out periods, that is, inside trading during times when particular people are barred from trading, Front-Running, which is trading on behalf of customers or corporations ahead of significant orders, and Misappropriations, where trading is done using confidential information that is stolen or misused. However, inside trading is not always illegal, on the contrary, there are specific circumstances decided by the authorities that dictate whether a certain deed violates the law.

In general, inside trading is legal when it adheres to all applicable rules and laws. For instance, it is legally allowed when it involves pre-planned transactions. In specific, traders can carry out transactions following rules which are known as the Rule 10b5-1, in order to avoid accusations of insider trading. This rule was constructed in order to clarify the Rule 10b-5 that was created under the Securities and Exchange Act of 1934. Another legally abiding scenario is when the information had been appropriately released to the public, for instance when it is publically disclosed through regulatory filings or other official statements. Last but not least, inside trading is not illegal when it refers to non-material information, like information that would not significantly impact the stock price or the investor choices. On the other hand, insider trading is regarded as a violation when it is conducted based on material non-public information that could significantly impact the stock prices and thus the investor choices. Another illegal scenario is the case of fiduciary duty breach, during which insiders, be it company executives or workers, violate their obligation to behave in the best interest of share-holders by misusing or disclosing secret information for their benefit. Finally, insider trading is strictly forbidden in cases of misappropriation, which is when trading is based on confidential information that has been improperly obtained or stolen.

What is more, the unlawful disclosure of inside information is also an umbrella terms for a series of actions where information is released without correct permissions, and the result is the distortion of market performance. The only exception to this is when the disclosure is a regular aspect of the person's employment or professional duties.

In any case, this type of manipulation has significant effects on the market. In a 2004 study, Chakraborty and Yilmaz show that "when the market faces uncertainty about the existence of the insider in the market and when there is a large number of trading periods before all private information is revealed, long-lived informed traders will manipulate in every equilibrium". The sooner this situation is spotted and healed, the better. Although there are several studies where regression and accounting principles have given outstanding results, like the work of Aggrawal and Cooper (2015), both of the above manipulation types cannot be detected by quantitative methods alone. There is need for understanding the intentions behind certain trades, as well as the timing and context of the chunks of information that were released, and possibly other factors as well. This very nature of that particular stock market fraud renders both statistics and the order book insufficient, and call for qualitative approaches and careful research, if we are to conclusively decide whether there was a violation or not.

2.5 Price influencing

Although easily confused with price manipulation, and the truth is that in some cases the signs in the market of both types are overlapping, price influencing refers to methods and techniques such as the manipulation of submission and transaction-based fixes, portfolio price manipulation, and triggering or protecting barriers, according to the work of Cumming, Dannhauser and Johan (2015). This shift of focus from price to benchmark can significantly impact the market as a whole. According to Andrew Verstein, Assistant Professor of Law at Wake Forest University, "benchmarks such as LIBOR or the S&P500 summarize market prices and […] are written directly into industrial contracts, financial derivatives, statutes, and regulations, and so their accuracy affects the economy every bit as much as the prices themselves." In an illustrative example on exchange rates of Euros and Yen that he includes in his 2015 article, he describes how easy it is to bias a benchmark, and consequently the market as a whole, in only a 2-minute time frame by aggressively buy on one venue and aggressively sell on another.

Manipulation of submission-based fixes is a term that describes the transmission or false or inaccurate information which will be used to calculate a closing price, a reference price or an index. This case typically demands certain knowledge of the mechanisms which lay behind fixes, a fact that extends beyond quantitative approaches. On the other hand, market manipulation of transaction-based fixes is the act of buying or selling a large volume of securities and/or derivative contracts, with the intention of illicitly influencing a benchmark. Quantitative analysts can analyse the transactions during specific time frames so as to reveal discrepancies from normal trading patterns.

Portfolio price manipulation, which is also known as window dressing, has as its objective to enhance the performance of a portfolio shortly before the deadline of a reporting period, thus falls into audit regulations and policies.

2.6 Improper order handling

In the work of Cumming, Dannhauser and Johan we read that improper order handling involves disclosing client order information, front-running, cherry-picking, and triggering or protecting stop-loss orders.

Disclosing client order information, also known as dissemination of client order information, is the act of giving third parties an information advantage over the market at large.

On the other hand, cherry picking is the practice where the allocation of a security to a client is withheld pending assessment as to whether the execution order is a winning or a losing trade. If the market moves adversely, the trade is allocated to the client. If the market moves positively, the trade is taken by the defrauding party. Alternatively, it can be the fraudulent practice of allocating winning trades to favoured clients at the expense of the rest of the clientele. Both cherry-picking and client order dissemination require a certain understanding of the interactions between brokers and clients as well as the ability to detect subjunctive choices during the trade execution, which is the reason why a quantitative approach on each own would not bear much fruit in their detection.

Front-running is the act of trading stocks just before a large market shifting transaction occurs. These trades are based on non-public information and will result in an immediate profit once the stocks are traded after the large transaction. It is similar to inside trading, but the major distinction is that inside trading happens when a single investor profits from internal company information that is not available to the public, whereas front-running involves using knowledge of a client's trades to place personal trades ahead of them, profiting from the price movement caused by the client's orders. Its detection is based on identifying statistical anomalies that showcase trades to be executed based on prior knowledge of pending orders, thus access to the order book data is also a necessity.

Triggering or protecting stop-loss orders seems to be a highly refined example of how one can make market practices work for his or her favour. A stop-loss order is a passive order which needs a trigger price in order to be activated. Above or below the stop-loss price, a trigger price acts as a threshold and only after crossing this price does the stop-loss order change from a passive to an active order. This type of order can help limit a trader's losses if a stock he/she owns falls more than the trader would like. When triggered, it becomes an active order, and the share are sold at the current market price. This is a common practice among professional traders, and many rules have derived from this technique of trading which are common knowledge, namely, the 7% stop-loss rule, which dictates that one should always sell a stock if it falls 7 per cent below what one had paid for it. However, it does not come without its drawbacks. Stop-loss orders may result in unnecessary selling or buying if there are temporary fluctuations in the stock price, especially with short-term intraday price moves. What is worse, is that market makers are keenly aware of any stop-loss orders one places with their broker, and can force a whipsaw in the price, thereby triggering the passive order to turn into an active one, by pushing the market, thus bumping you out of your position, and then running the price right back up again and making a quick profit. Quantitative analysts look at price fluctuations at surrounding key levels, which can be indicative of manipulative activities relating to stop-loss orders.

2.7 Misleading conduct

This type of manipulation refers to disseminating inaccurate or false information according to Cumming, Dannhauser and Johan. A rumour game model presented by Van Bommel in 2003 demonstrates the spread of inaccurate information among stock market audience which results into the changing of the market price to the benefit of him who started the rumour. Also, Fishman and Haggerty (1995), in their equilibrium model of information- based manipulation, show that if uninformed insiders behave as if they were informed and disclose their trading, they are able to influence a significant amount of followers and make a profit. At the dawn of the 21st century, this particular manipulative practice has escalated and is also referred to as mass-misinformation. Statistical methods can track inconsistent price movements related to news dissemination, but qualitative analysis would also be required in order for the content to be understood.

All the above information and categorisation according to detection requirements is summarised in the following table for a clear overview. In the first two columns we can see the manipulation types and sub-categories, whereas on the last two it is demonstrated whether quantitative analysis alone is enough signal suspicious activity, and whether the data from the Order Book is necessary for the detection of possible manipulative practices. The sub-categories that are marked in bold letters show whether a quantitative analysis of stock market data, without access to the data in the order book, can be sufficient on its own, making use solely of the traditional statistical and econometric methods.

Manipulation Types	Subcategories	Quantitative Analysis	Order Book Necessary
Price Manipulation	1. Spoofing 2. Layering 3. Rambling **4. Pools** 5. Cornering the Market **6. Squeeze** **7. Bull/Bear Trade** 8. **Pumps & Dumps**	Yes Yes Yes **Yes** Not enough **Yes** **Yes** **Yes**	Yes Yes Yes **No** No **No** **No** **No**
Circular Trading	1. Wash Trades **2. Churning** **3. Compensation Trades** 4. Parking (aka: warehousing)	Yes **Yes** **Yes** Not enough	Yes **No** **No** No
Misuse of inside knowledge	1. Insider Trading 2. Unlawful disclosure of insider information	Not enough Not enough	No No
Price influencing	1. Manipulation of Submission-based fixes 2. Manipulation of Transaction-based fixes **3. Portfolio price manipulation**	Not enough Yes **Yes**	No Yes **No**
Improper Order Handling	1. Disclosing client order information 2. Front-Running 3. Cherry Picking 4. Triggering / Protecting Stop-Loss orders	Not enough Yes Yes Not enough	No Yes Yes No
Misleading Conduct	1. Disseminating inaccurate or false information	Not enough	No

Table 1: Categorisation of stock market manipulation types

Chapter 3 – Manipulation and detection methods: a brief historical overview

3.1 Regulations

The existence of the stock market and the attempts to manipulate it have evolved side by side. In this chapter a few significant historical instances of stock market manipulation that ended up in bubbles and crises will be presented, as well as the measures taken in the aftermath to detect and prevent such practices.

The first bubble ever documented is also referred to as Tulip Mania, and it took place in the Dutch market as back as the 17^{th} century. In brief, tulip bulbs became the object of extreme speculation and hype, driving the prices of a single bulb, Semper Augustus Bulb, to as high as the price of a spacious house. Apart from the manipulation techniques that require a computer or try to circumnavigate regulations, all the other types which were described in chapter two were present: spoofing, shorting, front-running, pools, bull/bear raids (these phrases were first coined at that time), and so on and so forth. As De La Vega narrates, 'the speculators excel in tricks, they do business and find excuses wherein hiding-places, concealment of facts, quarrels, provocations, mockery, idle talk, violent desires, collusion, artful deceptions, betrayals, cheatings, and even the tragic end are to be found'. In Confusion de Confusiones, De La Vega describes in detail many of the practices that are listed in chapter two, and are still employed today. We read about insider trading, price movements and trading volume among others, but of course not in the structured way we analyse them today. He also describes empirical approaches that one can use to detect another party's manipulative tactics, and either how to take advantage of them or warn a close circle of friends, but the idea of institutional regulation does not exist yet, as Galbraith confirms. On 3 February 1637, at an all-time high and after the first future contracts were not fulfilled, the Dutch stock market plummeted.

In the aftermath of the crisis, the Dutch government stepped forward and first and foremost decided that the debts which had been created were gambling debts, and the stock holders would receive only a fair share of the value. Indeed, even today, the law defines a gambling contract as 'an agreement between two or more parties to wager or bet on the outcome of a future event or contest over which the bettors have no control and which typically involves chance. These contracts are generally considered illegal due to a variety of legal and public policy concerns. Some of the concerns about these contracts include gambling addiction, financial hardship, the potential for fraud, and other social problems like potentially attracting criminal elements who are lured by the prospect of making easy money.' But, 'Despite the general rule against enforcing wagers, there are exceptions, most statutory but some rooted in the common law. The common law permits the sale or purchase of securities: One invests an X amount of money in a company's stock, hoping the stock will increase in value, though he/she has no control over the firm's management. It is not called gambling; it is considered respectable risk-taking in the capitalist system, or "entrepreneurialism". Insurance contracts are also speculative, but unless one party has no insurable interest (a concern for the person or thing insured) in the insured, the contract is not considered a wager' according to the RV Community College. As we see, the first intervention in the stock market was made on behalf of the institutions in an attempt to contain the damage of a market failure.

Apart from the tulip mania, which is characterized by private credit and informal economy setting –although at the time the Dutch and the British markets were the most sophisticated- it was the Mississippi Bubble in France in the 1720s that raised awareness to the necessity for regulations against speculative practices and for protection of the investors. This instance is characterised of a commercial bank, the General Bank, founded by John Law, which held national debt, used paper –fiat- money and showcased intense leverage and sophistication, as we read in the work of Buchan (2018). In the aftermath, the need for a different monetary policy, and strong central bank, transparency and a regulation system were evident. It should be noted that the Bubble Act of 1720, enacted by the English Parliament, was falsely considered one such measure, as modern research has shown that its purpose was to impose a monopoly over British trade. Nevertheless, the Mississippi Bubble is considered by many to have led the foundations for the later formalization of the stock market institution, starting with the London Stock Exchange 3 March 1801.

The London Stock Exchange (LSE) set a global paradigm and paved the way for the modern world of trading. The brokers who were allowed to trade had to be official members and pay a fee, in contrast to Tulip Mania, where anyone could participate, but also deceive their clients. What is more, LSE set specific standards on trading practices, such as the disclosure of financial information (Companies Act 1862 & Securities Act 1948), aiming at bringing in the spotlight the fundamentals and avoid speculative practices, as was the case in the two previous examples, and adapted sets of rules for the protection of the integrity of the market.

3.2 Statistics

Apart from the regulatory approach to market monitoring, there were also different attempts to fully understand the market and its movements; in the 17th century, De La Vega encompasses the significance of market sentiment, the psychological component of the shareholders, when writing the following: ‘It is not the price of the stock that enlivens the stock market, but the rumours that circulate around them’. At the same period, we see the emerge of descriptive statistics in the work of John Graud (London, 17th century), and so many others to follow at an accelerative pace during the 18th century, like Thomas Bayes, Gottfried Achenwall –who is also believed to have coined the very term ‘statistics’, Pascal, Bernoulli, Laplace, Gauss, and Quetelet, to name the most prominent who appear in the timeline of the Royal Statistical Society. However, it is only until the 19th century that we have a shift of focus towards quantitative approaches in a more structured way; in 1885, while Francis Galton was working on hereditary and statistics, he developed was soon to be known as regression analysis, and in 1896 Karl Pearson introduced the Pearson correlation coefficient, which calculates the how and the how much one variable is related to another. Then came Spearman, Mitchel, Gosset (mostly known under the pseudonym Student), Kirstine Smith, Ronald Fischer, Jeffreys, Wilcoxon, Mahalanobis, Neyman, Rose, Yates, Wald, Kolmogorov, Kendall, Wolfowitz, Sukhatme, and Wallis. All the above works did not concentrate on finance exclusively, but they paved the way for others to follow. In the 20th century we see works from Tukey, Finney, Kruskal, Rao, Kiefer, Cox, Srivastava, Ghosh, and Efron.

Societies of quantitative methods were also formed. The oldest one seems to be the Manchester Statistical Society, which preceded the formal launch of the London Statistical Society in 1824. The latter became the Royal Statistical Society (RSS) on 15th March 1834. The first Econometric Society was established in Ohio, US, on 29th December 1930.

In a more strictly financial focus, we should note M.G. Kendal’s work on randomness and time series, M.F.M. Osborne on Geometric Brownian Motion on which Bachelier had also worked on his dissertation on stochastic processes (29/03/1900), Harry V. Roberts on the random walk hypothesis, and Horbrook Working on future prices and market maker behaviour. Later we see

Grauger's studies on the inferences of tastes and beliefs from bank and stock market data, and Morgenstern's spectral analysis of price movements, as well as predictability of stock market prices. More recent are the works of Fama on behaviour of stock market prices and the reason why the underlying distribution of price changes has an infinite variance, and of Mandelbrot on randomness and fractals in financial markets.

3.3 Artificial intelligence

In recent years we have come to witness the rise of artificial intelligence and its incorporation in several sectors. For reasons of clarity, the term 'artificial intelligence' should be defined and examples of the available variations ought to be provided. Instead of a standard definition, AI is mostly used as an umbrella term for methods which combine statistics and computer science in order to solve problems. The European Commission (2021) considers AI to be software that "can generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with" and it is created with the use of Machine Learning (ML), logic and knowledge-based approaches, as well as statistical procedures. The definition of the European Council (2022) is in accord with that of the European Commission, adding that an AI system "is designed to operate with elements of autonomy" and "infers how to achieve a given set of objectives". The International Organization for Securities and Commissions, IOSCO, (2021) defines AI as "the science and engineering of making intelligent machines, or simply, the study of methods for making computers mimic human decisions to solve problems. The Organisation for Economic Co-operation and Development (OECD), in 2021 defined Ai as "machine-based systems with varying levels of autonomy that can, for a given set of human-defined objectives, make predictions, recommendations, and decisions". The Financial Stability Board, FSB, (2017) considers it as "the theory and development of computer systems able to perform tasks that traditionally have required human intelligence". From a more general perspective though, a large proportions of the systems used in finance that are recently called AI have pre-existed the term in the form of statistics and econometric modelling, as we read previously in this chapter, and, due to developments in computational power and in the materials used in the hardware of computers, a wider range of more complex problems can be now addressed. In a way, what we are going to read in the next chapter is the evolution of what we saw in this one.

Chapter 4: Comparing Traditional and AI-assisted monitoring systems

In this chapter a comparison between traditional and AI-assisted methods to detect stock market manipulation shall be attempted. Both the strong aspects and the weaknesses of these approaches will be described, with emphasis placed on the way they detect manipulative incidents, and they will also be analysed, with the focus to be on the final result, which is for a system to be not only efficient, but cost-effective as well.

While it is true that traditional and AI-assisted monitoring systems share several common features and characteristics, they also have distinct differences. To start with definitions, programming glossaries define a feature as 'a part of a piece of software that performs a certain function, for example the capacity to generate and modify documents. More complicated software may have features that aren't available in rival programmes. ' What is more, the users can see the features of a programme or a method and interact directly with them, whereas characteristics have a more abstract nature. A characteristic is an inherent quality that defines a programme or a method as a whole, and can include aspects like cross-platform collaboration. In statistics and econometry, a feature can be hypothesis testing like a t-test, and the inclusion of fixed effects in a regression analysis respectively. On the other hand, a characteristic could be the non-parametric nature of a statistical test, or the Endogeneity of a variable in an econometric model. In this chapter we shall explore the features and characteristics of both traditional approaches and methods that are assisted by artificial intelligence of various forms.

4.1 Artificial Intelligence

Artificial intelligence has so many forms that some specifications should be provided regarding its main concepts. Machine Learning (ML) is a very broad system that has the ability to learn and optimise its behaviour when solving problems. It is not explicitly programmed what to do, but instead it uses statistical models and algorithms to analyse sets of data and draw inferences from their patterns. When talking about machine learning models, a model is actually a programme that contains parameters, whereas learning means the process of fine-tuning these parameters in order to fit the data, thus the same principles that apply to statistics and econometrics also apply in artificial intelligence and machine learning (AI/ML). Specifically, the model should neither overfit nor underfit the data, and it should be able to generalize what it leant and apply it to new, unseen sets of data.

According to its way of training, it can be further divided into supervised learning, unsupervised learning, and reinforcement learning. In the first case, which according to seminars conducted by Wolfram is the most common one, the model learns from labelled data, utilising both data and previous results, with the aim to predict outcomes when it is fed new unlabelled data. In other words, the model is presented with a set of known classes, the labelled data, and is asked to identify the class of input data which has not been encountered before. To achieve this, the first stage is the preparation of the data, which mainly refers to splitting them into a training set and a testing set. Usually the training set is up to 70 or 80 per cent of the data and the remaining percentage is kept aside for the testing process, but the exact ratio of the two sets of data that would be optimal varies depending on the model, the data, or other factors. The second stage is to create the model, which at heart it is a classifier, and the next steps are to train and test the classifier. The last step, and a very crucial one, is to evaluate the accuracy of the classifier over the testing data. As one can easily understand, this is again a similar procedure of evaluating traditional models in econometrics, where we use R^2 and other metrics for the same purpose.

The second variety of ML, unsupervised learning, is trained with unlabelled data only, and relies on algorithms in order to form clusters of data with similar properties and find patterns in the data without any human intervention. The procedure here is rather different than the previous one, and also we start to realise what a black box is. While creating the code for such programmes, we make use of certain readily-available functions from the libraries of each programming language, which is a rather straightforward process. For example, in order to employ the Random Forest technique which is contained in the SciKitLearn library in Python programming language, one has to simply call it with commands like the following:

“from sklearn.ensemble import RandomForestRegressor” for import,

“model_rf = RandomForestRegressor(n_estimators=100, random_state=42)” to call the function, and

“model_rf.fit(X_train, y_train)” to put it to work.

However, what one cannot see is the processes that take place behind the scenes. A machine learning model needs vectors of numbers to work, thus a map of the data, called encoder or feature extractor, needs to be created. During the training process, the data are organised in vectors through the encoder, and then both these vectors and the selected model are fed to the trainer. During the evaluation process, the new, unseen data go through the encoder to form vectors, and then both the outcome of the trainer of the previous process and the vectors of the new data are fed into the trained model, they are again transformed into vectors, and finally, with the help of a decoder, the programme provides the output, the answer to the problem we asked to be solved. For a clearer understanding of this process a flowchart-like schema is provided in Figure 1. The blue line represents the training process, the green line stands for the evaluation process, whereas the orange line corresponds to the production of outcomes.

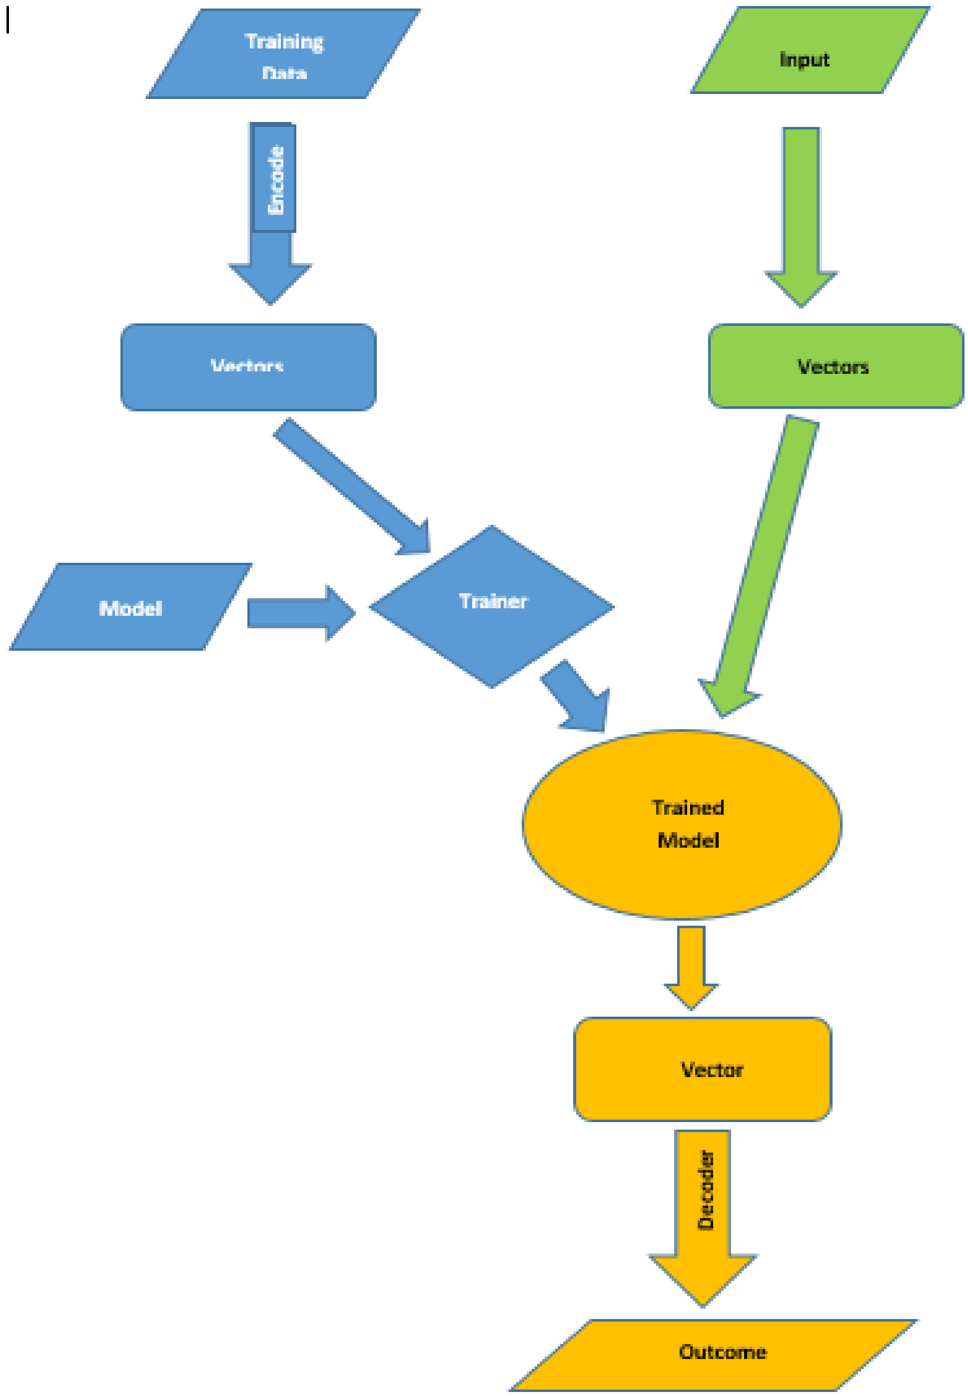

Figure 1: Machine Learning Workflow

A typical introductory example regards dog breeds. A model is fed with pictures of dogs of various breeds. Each picture is comprised by numerous pixels, and is pixel is perceived by the computer as a sequence of numbers for the corresponding colour, hue, shade, texture etcetera. The numbers behind the pixels are what is vectorised and follows the process we explained above. Unfortunately, one cannot surely tell what makes the model give one answer instead of another, and why some results are hilarious, like the one where a model identifies a cupcake as a small dog of a certain breed, a problem known as the "puppy or muffin" problem. In the article 'Puppy vs Muffin Problem' (Kondeth, 2020) an illustrative example depicts this dilemma, as we can see in Figure 2. In this particular example, one can see the similarities, and, with the way computers function in mind, can understand the confusion. Of course potential solutions are being explored by the programming community as well, and many of them are available online for everyone to try in cyber places like medium.com, freecodecamp.org and github.com, to name but a few.

The problem is that in more serious instances where confusion incidents are not so apparent, or where the real world situation might not be that obvious, such results raise serious concerns. In the world of stock market surveillance, a legitimate trade activity may be flagged as manipulative, a situation known as false positive. This is why although the model should create an alert, human analysts need to intervene and examine the flagged case before any allegations take place.

Figure 2: (retrieved from https://www.freecodecamp.org/news/chihuahua-or-muffin-my-search-for-the-best-computer-vision-api-cbda4d6b425d/)

The last of the three, reinforcement learning, unlike the two previous types, does not focus on data patterns, labels or properties. Instead, it tries to find the best policy that optimises a reinforcement signal provided by the user, such as a reward function.

Additionally, in all three of the previously discussed methods, a deep learning method can also been applied. This method also employs machine learning principles, but in its core it is based on neural networks, which are composed of layer, the number of which increases along with complexity. Each layer in its core is a regression, and each and layer feeds the others in a way the human brain and its cognitive centres are constructed and function. The parallelism between actual natural networks, the human brain and the computer neural networks are very illustrative described in Newman's book "Networks: An Introduction". In neural networks therefore there are many layers of fitting, and there is increasing abstraction, and at the same time the layers are alternating from linear to non-linear to linear and so on. Each layer is designated to target a specific feature such as convolutional, recurrent, LSTM, and many more. This is an extremely flexible approach that can learn from its own mistakes and improve itself, but at the expense of transparency and explainability, what programmers call a black box, and its need for data during the training process is vast. A common representation of a neural network is of a graph with arranged nodes that represent layers, and the interconnections are represented by lines, as we can see in Figure 3. Each layer contains information about its type, its input and output, and its trainable weights. We manipulate the neural networks like any other symbolic object, and when a neural network is trained with particular data, what really happens is an optimisation process of the weights inside the object.

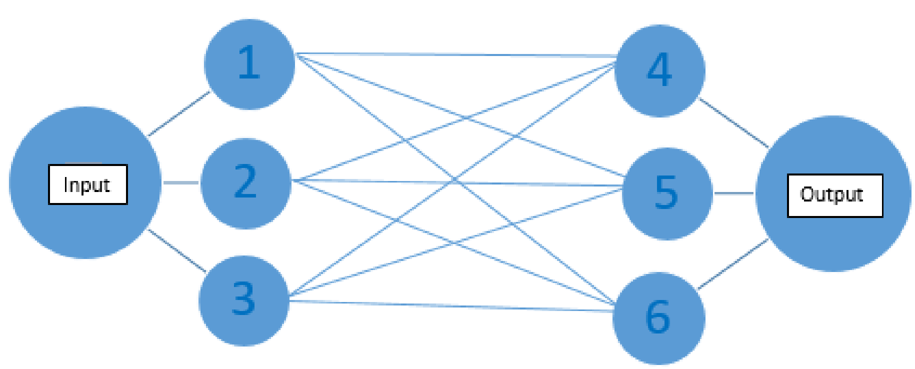

Figure 3: Graph of a simple neural network

On a different tone, another branch of AI is natural language processing (NLP) which receives words and texts as input data. It combines computational linguistics with statistics and machine learning in order to vectorise chunks of speech before they are fed to a statistical model, much like the usual data preparation of numbers. A final distinction in artificial intelligence models in between parametric and non-parametric. The former have a finite number of parameters, and the knowledge of the parameter set leads to the assumption of a completely determined probability distribution, whereas the latter are distribution agnostic and thus need a much broader set of data in order to generate reliable output, as we read in the works of Chandola, Banerjee and Kumar (2009).

All the above have several real-world applications in general, and in the world of finance and stock market surveillance in specific. Their particular features and characteristics can prove extremely useful in market analysis and fraud detection. Several techniques have been developed, and already widespread methods have been transformed into more nuanced, elegant procedures with the help of the many types of artificial intelligence.

To start with, entity recognition is a technique that machine learning algorithms use for the detection of stock market fraud. This techniques can identify entities like companies or investors who are associated with abnormal trading patterns, and at the same time it can assess the context in which their transactions are taking place. For instance, if many different stakeholders are related to a specific piece of news and at the same time have an unusual trading behaviour, this system will recognise these relationships and flag them as suspicious, so that they can undergo further investigation. What is more, sentiment analysis is another area of near-excellence of AI systems according to Bollen, Mao and Zeng (2011), as they utile natural language processing, also referred to as NLP, and they can access press releases, social media, news articles and the like, in order to harvest data that are indicative of the public sentiment. In this way, they evaluate the public sentiment trends, and then they use it to spot trading signs that are pointing to potential manipulation practices. For example, if the public sentiment towards a stock has been evaluated as negative and the data exhibit sudden spikes, that scenario might be flagged as suspicious, so the analysts know to allocate resources into the deeper investigation of the incident.

Moreover, alert generation in AI-augmented monitoring systems can be done in real-time, as such algorithms have the ability to analyse vast amounts of data at very high speeds. The core idea remains the same as with the traditional statistics and econometrics, though. When a value that the analysts have set as a threshold of normal activity is crossed, like when there is positive sentiment and at the same time excess volatility, an alert is created and the teams of analysts know where to take a closer look.

What is more, expert systems, which are also called decision support systems, and, according to Turban and Aronson (2018), they provide analysts with deeper understanding of complex patterns, so that they can make more informed decisions regarding the compliance to industry standards, and they can prepare more accurate reports for the regulatory bodies. To do so, these systems utilise features of regulatory compliance and regulatory reporting, which are also assessed by the algorithms themselves, as we will see in the example of the next chapter. What is more, audit trails within an artificial intelligent system can log every action taken, and as a result they enable analysts or regulators to trace back the decision making process and check whether actions of manipulation were carried out. Thus, the audit trail feature adds another layer to the model, one of compliance and accountability, a fact that ultimately enhances transparency.

4.2 Statistics and Econometrics

On the other hand, the traditional methods for detecting instances in the stock market that are suspicious for manipulative practices are based on strict rules which are determined in advance, like a 10% deviation from normal behaviour. Econometry and Statistics use several procedures to spot those deviations from the norm which neither market related news nor fundamentals can fully explain. In the core lies regression analysis, which can either be employed on its own, or be a basic element of more nuanced methods and approaches. Such methods include several models of great importance, with the most distinguished to be ARIMA, GARCH, and many others, while various different approaches can be event study analysis, network analysis, and more.

The core methodology of regression analysis is of fundamental importance and has a very wide range of applications, but in this paper we shall focus only on forecasting and detecting suspicious incidents that can be indicative of manipulation. The general aim of regression is to find the way in which a phenomenon that we are studying is influenced by certain factors of our choice. The phenomenon in question is called the dependent variable, usually symbolized by Y, whereas the factors we assume that might influence it are called the independent variables, usually symbolized by X. As models may increase in complexity, Y and X can be vectors and contain several variables. What regression does is to find the values of the independent variable co-efficients, beta, which tell us if and by how much a particular factor X influences the phenomenon we study, Y, and towards which direction. Additionally, it provides very useful metrics, such as R^2 or p-value, which are indicative of the model's fit and robustness. Regressions analysis is also very flexible, as it can be used with different types of data, such as time series, cross-sectional data, or panel data, and apart from the relationship definition that is described above, it can be used for hypothesis testing regarding the above parameters, for forecasting in several different models, or for the close study of a particular incident, a practice which is known as event study analysis, and regression lies in its very heart, also. (Greene,)

To start with the ARIMA model, which stands for autoregressive integrated moving average and can collapse to ARMA, AR, or MA when some of its components are set to zero, this model uses the past values of a time series in order to predict its future ones. This prediction of stock prices which was based on historical trends can then be used by analysts to spot significant or unexplained deviations from what is considered expected behaviour and may signal manipulation. Furthermore, if the focus is shift to the error variance of an ARMA model, then a model named GARCH is employed. GARCH models, whose initials mean generalised autoregressive conditional heteroscedasticity, and they can also collapse to an ARCH process if its beta component is set to zero, are used to retrieve volatility metrics. GARCH models also have many variations, like NGARCH to accommodate for non-linear cases and leverage effects, IGARCH to integrate unit root, or Gaussian process-driven GARCH, which is more robust to overfitting and captures non-linear dependencies without a severe increase in complexity. According to literature however, the claim that GARCH models can also measure the impact of news on the volatility of a stock, thus incorporating aspects of sentiment analysis in the traditional statistical methods is not conclusive yet, as Engle and Victor (1993) state. By using this model, analysts look for abnormally large residuals, shown in the plots as spikes, could be the sign of manipulative practices, such as pump and dump schemes, and as a result it is worth to investigate these instances further. (Meidanis)

In event study analysis, as said above, regression lies at its core. According to MacKinlay, in order to study a specific event, we need a specific design of regression in order to isolate the effect of the event in question from other factors, and then study its impact on the stock price, or other financial variables. The output of this procedure is two metrics, the abnormal return, AR, and the cumulative abnormal return, CAR, over a strictly specified time frame, called a window. These metrics calculate the difference between the stock's expected performance, usually calculated by a regression to define what is normal and make the forecast, and its actual performance that is seen in the markets around the event date. This comparison may raise a flag as to suspicious activity according to the amount of deviation from the norm. To be able to capture the event's impact, the data are divided into event windows and estimation windows. The event window is a time around a specific event, say a suspicious one for manipulation practices, when we wish to analyse a stock's performance, and there is a split into the pre-event and post-event windows, which centres around the date of the event. The estimation window on the other hand is the period before the event window, which we consider as the normal behaviour and we use it as a benchmark for our comparison. Then, the difference between the actual return of the stock in the event window and the expected return that was calculated in the estimation window are being compared to find the abnormal returns, and the sum of all these differences over a window is the cumulative abnormal return, which shows the event's total impact. All ARs and CARs undergo statistical tests like a t-test for their significance to be identified. In other words, the methods ensures that the metrics do not make use of incidents that can be put down to mere chance. A potential sign for manipulation that an analyst would look for would be the existence of statistically significant abnormal returns around a spike, or if many events or stocks exhibit similar patterns and, at the same time, their abnormal return metrics are statistically significant around a time frame. The former can be a signal for a pump and dump scheme, or spoofing and layering in the order book, whereas the latter is more indicative of co-ordinated manipulation practices on behalf of many individuals. However, we should note that the selection of the correct window can be a rather tricky procedure, and the model specification as well as the efficiency of the market are added limitations to the procedure.

Another popular method, networks (Liu and Xiao, 2012), is defined according to Newman as "the understanding and modelling the behaviour of real-world networked systems and observational data". To understand its basic application in the stock market surveillance, we shall use a representation of a directed network from the book of Newman "Networks: An Introduction", from page 104, Figure 6.2 of the online version. In this picture we can see six numbered nodes, which can represents six different stock market participants. The lines that connect the edges inform us about the interaction of these participants, and the arrows on each line indicate who contacts or places orders to whom. We see that node number 3 is the recipient of orders from nodes one and 5. This can represent two investors giving their orders to a broker, who then places both on node number two. Now, if the trading data show that a particular stock has a significant trading volume but its price is relative the same, and at the same time we observe multiple orders in the Order Book placed from nodes one and five but of the opposite direction, that is, one cancels the other, that should be a sign for further investigation for potential manipulation practices. In an alternative scenario, if we observe node six to place opposing bids on nodes four and five, where nodes four and five represent different markets, and, for the example's shake we assume that the trader in node six exchanges yen and dollars, then perhaps we are witnessing an attack on the exchange rates.

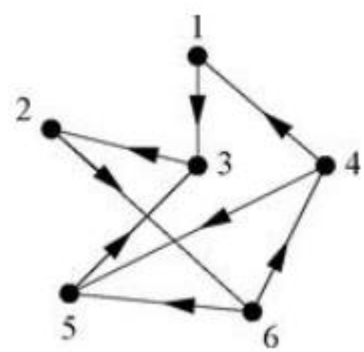

Figure 4: (retrieved from Newman's book Networks: An Introduction, page 104, Figure 6.2)

This is an extremely simplified scenario with a very basic network graph, aiming to just give an idea of the potential of network analysis and cross-market analysis in the battle against misconduct practices in the markets. Network analysis can incorporate many different mathematical and statistical or econometric methods to tackle specific problems. In the same book Newman analyses random walks for networks (page 155), as well as dynamic systems of networks, both discrete and continuous, steady and dynamic states, attractors, oscillations and many more on chapter 18.

4.3 Hybrid Models

Apart from pure statistics, pure econometrics, or pure artificial approaches, hybrid models can be employed. With the term hybrid we mean a combination of traditional and AI approaches to harness the best results of the two. Several programming languages and applications already have functions in their libraries for this very purpose. For example Mathematica, which is coded in Wolfram language, in its anomaly detection functions can utilise multinormal methods to find anomalies in a given set of data. What lies behind the scenes is the well-known normal distribution, and a selected acceptance threshold, the equivalent confidence level of the usual statistical procedures, but instead of a usual fitting process a neural network is employed. Similarly, other functions can be utilised to handle anomaly detection in the stock market and sere surveillance purposes by combining the best aspects of the two approaches, statistics and artificial intelligence.

All in all, the traditional methods of statistics and econometrics are clear, well-established, and have been tried and tested for decades. Apart from the fact that these methods are reliable themselves, the scientists who employ them have a very wide literature at their disposal to scrutinise every small detail to answer a question that may arise, and they also have years of experience to assist them, human intuition supported by cumulative experience. However, as we saw, these techniques are not easily adaptable to new, more sophisticated and more complex manipulation methods. Not only that, but also the new era in trading with its many platforms and the innumerous data makes real-time analysis and scalability to be really concerning issues regarding the effectiveness of the traditional monitoring methods. In contrast, systems that incorporate artificial intelligence of any kind, like machine learning and natural language processing, can handle real-time analysis and scalability much faster and more efficiently regarding the results. The cost of employment of such methods however may be deterring for smaller businesses, delaying their adoption at a large scale, as we discuss more extensively in the following chapters.

Chapter 5: A demonstration

5.1 Gamestop: the story

Gamestop is a company that sells video games and gadgets, and its stock was heavily shorted by Wall Street investors and hedge funds, such as Melvin Capital. Keith Gill, an analyst of Mass Mutual who was also a youtuber that posted his investment strategies online, posted his decision to buy and hold onto the Camestop stock –a practice known as diamond hands. He gradually had a deep influence on an increasing crowd of people, who followed his footsteps and started buying massively the stock, driving its price up and causing the hedge fund to lose significant amounts of money. Citadel and Point72 infused 3 billion dollars to Melvin Capital in a bailout attempt, and –allegedly- forced RobinHood CEO, Vlad Tenev, to de-activate the 'buy' option of the Gamestop stock in his platform in order to continue the operations for Robin Hood to go public. All four main participants of this instance, the Robin Hood CEO, the Citadel CEO, the Melvin Capital CEO, and Keith Gill were called into Congressional hearing, but the above accusation was never proven, as the lawsuit against Robin Hood and Citadel Securities was dismissed in court and SEC filed no charges. Nonetheless, the moment the buy-option button was re-activated, the crowds started buying again, driving the Gamestop stock price back up. Some of them earned money, some of them lost, and Melvin Capital filed for bankruptcy with a total loss of 6.8 billion dollars. Keith Gill's final post was dropped on 16th April, 2021 before he retired from public with a net worth of 35 million dollars.

5.2 Comparing model forecasts

To demonstrate a comparison between traditional (statistical and econometric) techniques and AI-assisted ones, an ARIMA (5.1.0) model and a Random Forest model were employed respectively, and they were fitted or trained on the stock market data of Gamestop (ticker name: GME), which was caught in a manipulative incident back in 2021. The ARIMA model includes five previous closing values (in other words, five lags), and the time series needed to be differenced once to have stationarity. The Random Forest utilised 80 per cent of the data retrieved from yahoo finance for its training, and the rest 20 per cent for the forecast. The code was written in python and several libraries were utilised for the fitting part, the machine learning part, the plots and the results. In the following pages we will see the graphs and analyse the results in order to compare both approaches and decide which performs better. The code is provided in the Appendix.

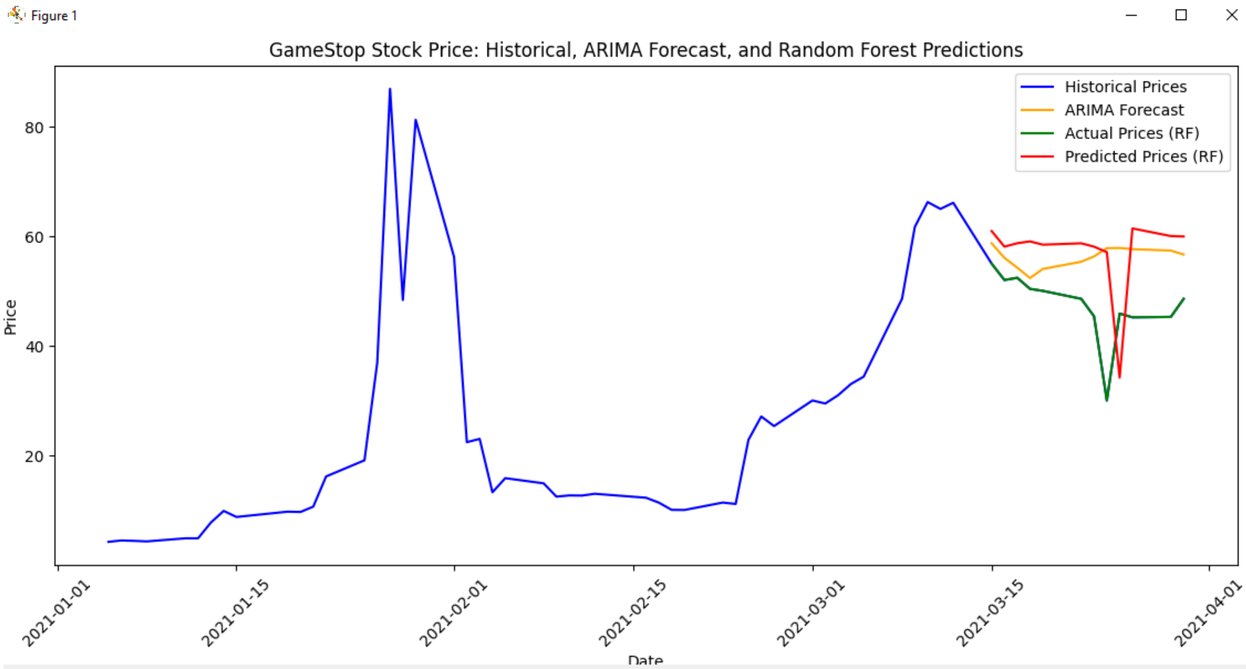

Figure 5: GameStop Stock Price: Historical, ARIMA Forecast, and Random Forest Predictions

Actual prices min: 30.4 at 23/03/2021

RF Predicted prices min: 34.4 at 24/03/2021

ARIMA Forecast min: 52.4 at 17/03/2021

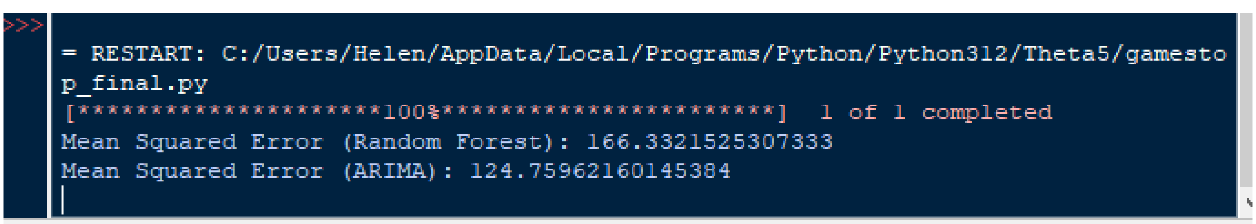

Figure 6: Evaluation of the two models

In the above figure we see the final evaluation that python does to both models. We read that the Random Forest model has a mean square error of 166.33 whereas the same metric for the ARIMA is 124.759, suggesting that the latter is a better model than the former.

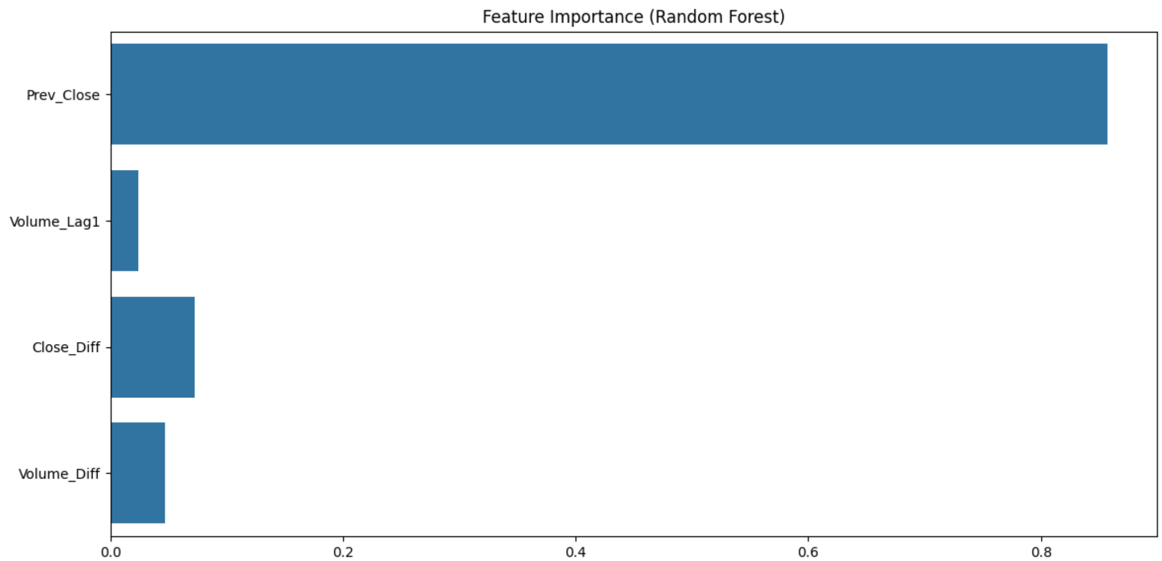

Figure 7: Feature Importance in the Random Forest Model

The above bar plot shows the feature importance of the Random Forest procedure. The prev. close bar has a value that exceeds 0.8 and also by far exceeds the rest of the metrics, indicating that the closure price of the previous day is the most important feature in this model, and this is the one the model heavily relies on when making predictions. This result is confirmed by the theory of autocorrelation in the stock markets, that is, today's price is strongly influenced by yesterday's price, and we observe the model to learn the same exact thing during its training. On the other hand, yesterday's trading volume, named volume lag on the graph, proves itself to be rather insignificant, with a value less than 0.05. As a result, the model does not rely on this feature when creating the predictions. What is more, the random forest model finds the short-term differences of the price and the volume to carry some significance, with values of around 0.05 and 0.1 respectively, so in the forecasting process these features of the trading activity are being taken into account, although they do not weigh as much as the closure price.

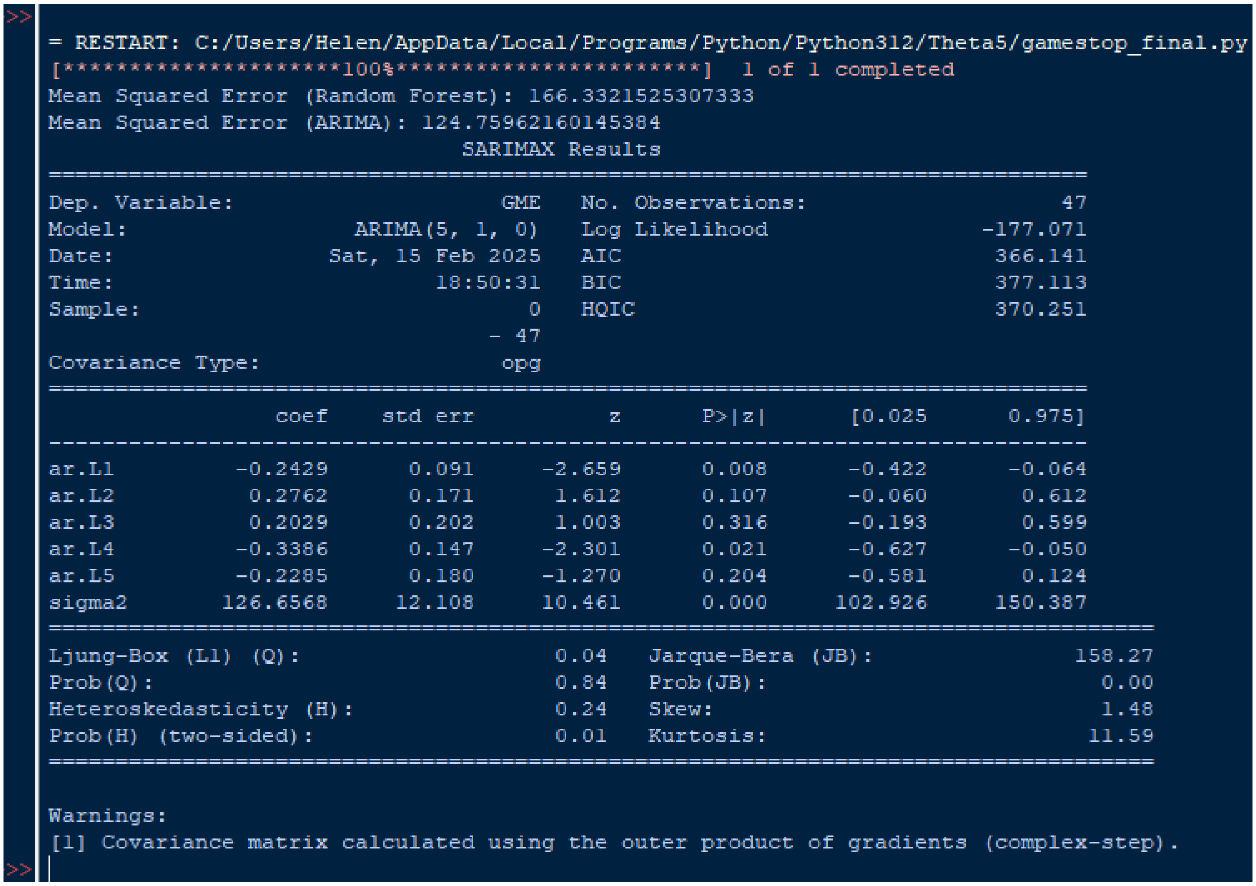

```
>>
= RESTART: C:/Users/Helen/AppData/Local/Programs/Python/Python312/Theta5/gamestop_final.py
[*********************100%***********************]  1 of 1 completed
Mean Squared Error (Random Forest): 166.3321525307333
Mean Squared Error (ARIMA): 124.75962160145384
                               SARIMAX Results
==============================================================================
Dep. Variable:                    GME   No. Observations:                   47
Model:                 ARIMA(5, 1, 0)   Log Likelihood                -177.071
Date:                Sat, 15 Feb 2025   AIC                            366.141
Time:                        18:50:31   BIC                            377.113
Sample:                             0   HQIC                           370.251
                                 - 47
Covariance Type:                  opg
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
ar.L1         -0.2429      0.091     -2.659      0.008      -0.422      -0.064
ar.L2          0.2762      0.171      1.612      0.107      -0.060       0.612
ar.L3          0.2029      0.202      1.003      0.316      -0.193       0.599
ar.L4         -0.3386      0.147     -2.301      0.021      -0.627      -0.050
ar.L5         -0.2285      0.180     -1.270      0.204      -0.581       0.124
sigma2       126.6568     12.108     10.461      0.000     102.926     150.387
===================================================================================
Ljung-Box (L1) (Q):                   0.04   Jarque-Bera (JB):               158.27
Prob(Q):                              0.84   Prob(JB):                         0.00
Heteroskedasticity (H):               0.24   Skew:                             1.48
Prob(H) (two-sided):                  0.01   Kurtosis:                        11.59
===================================================================================

Warnings:
[1] Covariance matrix calculated using the outer product of gradients (complex-step).
>>
```

Figure 8: ARIMA Analytical Results

To analyse the co-efficients, we start with the order of the ARIMA model (p,d,q) which is (5,1,0). The component p represents the lags, and p=5 means that the model uses 5 previous closing prices, the d part shows how many times we had to difference the data to make the model stationary, which in this case is only once, and finally the q is the number of the moving average terms, but no such component was used in the final code. The p-values of the co-efficients of the autoregressive part, the AR, inform us about the significance that the lagged terms have in the predictions. Only ar.L1 and ar.L4 are below the p-value threshold of 0.05, thus they are the most significant in the forecasting. This means that the future price of the stock is affected primarily by the closing price of the previous day with an ar.L1 p-values of 0.08, and, to an extend, by the closing price 4 days ago, with an ar.L4 p-value that equals 0.021, and this fact shows that the model was able to discover some pattern in the movement of the prices. What is more, both coefficients of these two lags have a negative impact on the current price, as their negative values of -0.2429 and -0.3386 indicate.

The sigma2 stand for the s^2, or, in other words, the estimated variance of the residuals. Then the residuals are undergone standard procedure tests for normality, autocorrelation, and heteroscedasticity. To test normality, a Jarque-Bera test was employed, but its probability was calculated to ne null, meaning that the residuals are not normally distributed, which is very unfortunate regarding the robustness of the model. The Ljung-Box test for autocorrelation took a value of 0.84, which means that the residuals are not significantly autocorrelated, thus the model can capture patterns in the data very well. Finally, the heteroscedasticity test has a probability value as low as 0.01, showing that the presence of heteroscedasticity in the data is statistically significant. In other words, the variance of the residuals is not constant, like the same s2, but varies over time, like $s^2{}_1$, $s^2{}_2$, …, $s^2{}_k$, a fact that is rather concerning as well. Overall, the results of the above diagnostic tests uncover the fact that some of the ARIMA model assumptions were violated, and there should be further tests to check whether this violation has a significant impact on the predictions. The remedy for this, if there is no AI-assisted option, could be to transform the same ARIMA model, for example by using the logarithms, or it could be to repeat the same procedure by varying the model components (p,d,q), or even try a different model, like a GARCH, as this model can handle volatility much better. However, even in this case we once again have to test and fine-tune its components until we end up not necessarily with the optimal, but at least with an acceptable solution.

In the table of results we also see the Log Likelihood at a value of -177.071, and the values of the Akaike, The Bayes and the HQIC criteria to equal 366.141, 377.113, and 370.251 respectively. In case of multiple iterations at different p, d, and/or q order we would compare these values, and opt for that ARIMA variation with the highest Log Likelihood and the lowest of the remaining criteria. Alternatively, we could make use of the 'auto-ARIMA' option from the python libraries. This is an algorithm that iterates itself while using various combinations for the p, d, and q, evaluates the results, and suggests which ARIMA order is the most suitable for a particular set of data. However, this particular option would require a bit more than an average personal computer, especially if the data set is rather large and the ordinary computers for personal use have a somewhat limited capacity.

5.3 Overview

Overall, when we compare a traditional approach, like the econometric ARIMA model, and an AI model, like the Random Forest which employs machine learning techniques, we can more clearly observe the advantages and the limitations of the two. Both models agree that the previous day's closing value is of paramount importance for the prediction of the next day's values. This fact confirms a short-term memory in the movements of the price, and it should definitely be taken into consideration when the stock market is being monitored for unexpected changes that could mean an incident of manipulation. For example, if ARIMA or Random Forest have been used to establish a particular threshold of a stock's path which would be considered the normal case, then a deviation from it, like an unexpected increase or decrease in the price that is happening fast and it is not supported by market news or trends might be flagged as suspicious, and further investigation might need to be done. The same goes for spikes in volume, especially if they take place right before a significant fluctuation of the price.

However, ARIMA suggests that we take into consideration the price of four days ago as well, even with a smaller weigh, whereas the Random Forests recommends taking the differences in price and volume of the previous day as relevant factors to the forecast, although not that crucial. On the down side, ARIMA's assumptions of normality and homoscedasticity of the residuals are violated, thus the forecasting may not be very effective, and the Random Forest process is what is called a black box in its core, meaning that we cannot directly observe why it gives the results we see, and that we have to rely on the metrics it creates about its own effectiveness.

Apart from the metrics, we should pay special attention to the graph itself. We see that the short memory of 4 days in the ARIMA model prevents it from seeing longer into the past and notice that the price of the stock does not correspond to what is considered normal in the graph, thus it cannot predict the eminent drop that comes a little later. On the other hand, the Random Forest has used all the lower prices at its training, and therefore has an understanding of the usual levels of the price, and seems to be able to predict a sort of correction soon enough. It indeed predicted the downturn of the price in retrospect, marked by the spike in the graph. What seems as a mystery however, is that it could foresee that the price will drop to its average levels pretty much around the time it actually happen, and then also forecasted its subsequent upward movement, a fact that also happened in reality. To put these forecasts into context, the drop of the price occurred when the platform RobinHood de-activated the buy button, creating anxiety to investors, who had only two options, to sell or hold. There have been allegations that RobinHood was forced to disrupt the trading activity of the stock in order to assist the hedge fund which was shorting the GME stock for quite a while and earning vast amounts of money, and was now threatened with bankruptcy. After the immediate outcry of the public, RobinHood re-activated the buy button and the stock rise again, a fact that the Random Forest model was able to predict. However, Random Forest is not a model that incorporates news and the sentiment of the market in its forecast.

The reason that the Random Forest could see the drop and rise of the stock price is partially due to its training data, which indicated a different average price levels than the few days of skyrocketed prices, as we said above, but there are more things to consider. One of them is the way the data were split into training and testing. The training data, being at an 80% level of the entire set, provided the model with a glimpse of price fluctuations and upwards trends. If the data were split at a different ratio, say 70% - 30%, or 60% - 40%, then the results could have been really different and the model might not be able to foresee the rise above the average, the subsequent drop, and the final rise again of the stock price.

A second, and equally important thing to consider, is the way Random Forests function in general. Random Forests utilise supervised learning, which, as we mentioned above, works with label data during the training process. The model learnt to map the data it was fed by creating subsets of these training data and by constructing decision trees in order to give the output we observed on the plot. In this particular example, the algorithm wasn't given a specific seed to initiate from, meaning that if we run the programme again, the results will be slightly different, and the spike may or may not appear again, it might be shallower or deeper, and it can appear at around the same date or a little earlier or later than the date we see it in the plot, depending on the way the algorithm splits the data into trees. To be fair, the results will not be radically different, as random forests are a rather robust methods of estimation, but they would not replicate exactly. What is more, although random forests can be combined with neural networks or sentiment analysis, neither of them was applied in the example code, meaning that the model had no way of knowing the news, the investors' feeling, or any other relevant information to conclude that there will be any kind of interference in the market. The same is true about the ARIMA model, and since any kind of interference or correction cannot be predicted by past data, the ARIMA could not produce the spike that the random forest did. The fact that the latter was successful to capture the reality can be partly coincidental, due to the way the training data –including the lower prices before the pump- were formed into decision trees. Hence the comparison of the mean square error of the two methods in Figure X; the ARIMA model is considered to fit the data better, although the plot and the unexpected reality say otherwise. In bottom line, the random forest was able to predict something that it should not be able to, partly by coincidence, and it should not be trusted for long-term predictions before more refinement and fine-tuning.

The main reason that this example was created is to highlight several things that concerns artificial intelligent practitioners. First of all, we see the extend of the black box, as we have no way to tell how this prediction came about and why it is accurate. Secondly, we can understand that AI methods need a human supervisor to put the results into context. Moreover, in a case where unexpected events took place, strict quantitative methods may not be sufficient on their own. Hybrid methods, that is, the combination of traditional approaches with artificial intelligence, could possibly fetch better results, but not is such scenarios. Only augmented by qualitative aspects, sentiment analysis, the keen eye of the experienced analyst and human intuition, these methods would perform adequately enough to be trusted. In all, one should bear in mind that these are only tools at the analyst's disposal and they should be primarily utilised to save him from mundane work, and not replace him whatsoever.

Chapter 6: Benefits, Limitations, and Concerns regarding AI

The adoption of AI in the financial sector in general and in stock market surveillance in particular has transformed the status quo greatly. In this chapter we shall discuss the benefits and limitations of this artificial intelligence reality, as well as attempt to discuss about other significant impacts which have occurred due to this technology. Before we delve any deeper, we should introduce two terms used by ESMA, RegTech and SupTech. The former stands for Regulation Technology and refers mainly to information technology that is used in regulatory compliance and it conducts activities such as risk management, while the latter is about Supervision Technology which is mainly used by supervision authorities. What is more, the observations of the 2019 CRO Forum on machine decisions has been taken into account while presenting the following, regarding both the advantages and the possible pitfalls of AI.

6.1 Benefits

6.1.1. Scalaility, Sentiment Analysis, Network Analysis

Perhaps one of the main advantages of AI-assisted stock market surveillance systems is that they can monitor the market in real time. This is due to the advances in technology, both of the hardware/software and of the architecture, as Burns, Oppenheimer and Brewster (2016) explain, enabling computers to scale up their usual tasks and carry out far more complicated, heavily data-driven tasks. Their scalability, that is, the ability to process vast amounts of data very quickly, in combination with their ability to detect patterns and anomalies, can provide analysts with immediate alerts when suspicious transactions are taking place, and it is then the analyst's job to investigate further and decide whether there is indeed a manipulation scheme or not. What is more, these systems are capable of integrating information from various sources such as news or social media data into their processes, thus enhancing the results they create.

This sentiment analysis incorporation can reveal very useful insights into what the movements of the stock market mean. However, sentiment analysis is more of an extra tool at the hands of the analysts rather than a method that can stand on its own. It is not a matter of efficiency, as this type of analysis can be a very rapid process and can be conducted at a very large scale, but a matter of accuracy mainly. The heavy dependence on the quality of the data – which, especially in the case of social media, might not be the best- and the model's limitations to understanding irony, sarcasm, or even humour, may have a severe impact on the sentiment indicators that the model will produce. In addition, in the real world there is some lagging between the spread of a piece of news and the reaction of the stock market, a fact that may result in sentiment analysis models not being very efficient at short-period forecasts. Of course this can have its upside, too, as AI processes and analyses data much faster than humans, and can show us a market snapshot right before it happens – if it is not misinterpreting sarcasm, of course.

Additionally, although artificial intelligence enables financial surveillance models to handle data at extremely large scale and at an extremely large speed, they lack true understanding of the context behind the data. This means that market conditions and news may have a certain sentimental impact on investors, which might escape the model's realm, it is beyond it understanding. Sentiment analysis is evolving of course, but the human touch in the final interpretation remains necessary. In such cases, the fact that many of the model's processes are a black box to the analysts can be highly problematic, and the more traditional counterparts of statistics and econometrics can be more trustworthy because scientists have more experience on what to trust and when their models need further refinement. In the case of AI, experience will become a reality eventually, and we will be discussing of equivalent years of experience on artificially intelligent systems, so the aspect of the necessity of the final human touch should be reviewed deeper in the years to come, and, at the same time, Natural Language Process (NLP), that is, systems that combine AI and linguistics, will inevitably improve as well.

On a different tone, artificial intelligence has greatly benefited network analysis, where, nowadays, a smaller number of expert personnel and a larger number of more complicated relationships can be analysed. The process involves the analysis of the Order Book data in an attempt to detect manipulation types that involve different market agents or even entirely different markets, as we saw in detail in previous chapters. This enhanced type of network analysis can use several years of raw data of the Order Book and combine them with stock market data, and is able to reveal cases of manipulation that employed a more sophisticated scheme and took place over longer times and among an increased number of participants, in comparison to the capabilities of the traditional network analysis. In reference to Table 1, where we can see that access to the Order Book is a pre-requisitive in order for almost half the manipulation types to be detected, we can get a glimpse of how important this breakthrough is and what potential it holds for the future.

6.1.2 False Positives / False Negatives

Another one of the major advantages of the AI systems in manipulation detection is false positive and negative outcomes. The term false positive means that a legitimate trading activity was incorrectly flagged as suspicious, whereas false negative means that the model failed to detect an actual manipulative incident. In order to prevent real manipulation from escape detection, models have inbuilt characteristics that make them cling towards false positives. In plain English, they will create more false alarms in order to minimise the danger of missing a true case of fraud. Even in this scenario however, artificially intelligent systems have managed to outperform their more traditional, rule-based counterparts of statistics and econometrics. Machine-learning based surveillance systems achieved this result by using mathematical optimization methods in experiments which were conducted by the UK Financial Conduct Authority, while similar tests are being carried out by the Bank of Italy as well, reposts ESMA in 2019.

6.2 Limitations

6.2.1 Accountability

One of the greatest issues in the use of artificially intelligent models to detect market manipulation is perhaps accountability. According to Prenio and Yong (2021), the key difference between regulatory requirements for traditional and AI models is the stronger emphasis for the latter on human responsibilities in order to prevent discrimination and other non-ethical decisions. To make matters more clear, when an artificially intelligent model is characterised as discriminatory, the general concept is that it produces biased results despite having been trained by unbiased data. An example will explain the situation further. Say we are training an AI model to distinguish the average high of the residents of a country, and create an alert for the extreme values, and say we have chosen supervised learning as the training method. During this method of training, we have fed the model with data regarding the citizens, like: “female, 28 years old, 1.78m” and label this as “tall”, or “male, 31 years old, 1.78m” labelled as “average”, according to the reality in Mediterranean countries, where a height of 1.78m is considered as tall for a woman but average for a man. When the model is fully trained and it is supposed to present results, if it is met with a height of 1.85 it will probably give as a result that this is a man, because the possibilities of being a man –that is, taller – is higher. On the other hand, if the population of the country in question consists say of 60 per cent women and 40 per cent men, then the model will probably assume that a 1.85m tall person is a woman, due to the increased probability of being one. The problem here is not the data, nor the labels or the training of the model, but rather the way it makes inferences and produces answers. To avoid such a scenario in the example, there should be requests on behalf of the model as to the gender of the person and an additional if-loop in the inference process, or there should be results presenting both scenarios and their respective probabilities, like “The person is X% a tall woman and Y% an average man”.

However, when the models become more and more complex, the circumnavigation around biases becomes more and more demanding, and certainly the models themselves cannot be held responsible for this. Here is that accountability comes into question. Somebody must be responsible for this, it should be a human, and it should not be only one, as in the creation, training, back-testing, and overall evaluation of models a whole chain of experts is involved. But it is not only them, as in the case of manipulation detection in the stock market, the IT experts carry out the plans of and interact with financial experts, who also share responsibility in communicating the reality correctly and to its full extend, so a potential bias can be spotted and avoided. On top of this vital collaboration which already occurs, are the institutions. Their own responsibility is to give clear descriptions of what constitutes market fraud and clear instructions regarding the regulations that should be met. As Prenio and Yong put it: “while the emerging AI principles are useful, there are growing calls for financial regulators to provide more concrete practical guidance given the challenges in implementing the principles”.

6.2.2 Data Quality and Data Collection

In quantitative methods, be it traditional or modern, data quality is key in order to receive reliable results. However, what seems to be the main difference between statistics and econometrics on the one hand, and artificially intelligent systems on the other, is that the latter seem to be more resilient to not so good quality data. In other words, though it is definite that poor quality data will lead to poor quality statistics and econometrics outcomes, the same is not necessarily true for AI systems. Biased data will always bring biased results in the traditional methods, but AI, not being rule-based, can produce acceptable results even with biased data. After all, as we read in the previous section, bias and discrimination in AI differs significantly from what we are used to understand when working with more traditional models.

Data collection on the other hand is a different thing. For one, vast amounts of data need to be harnessed to train AI models, and concerns have been express on the Press that in the not so distant future the more advanced counterparts of today's AI systems will not have sufficient data to be trained. This being primarily the field and a concern of IT experts, and beyond the scope of this paper, we shall not discuss it deeper here. Instead, we shall focus on the flip side of the coin, which is data collection for supervisory purposes. According to a 2019 report, ESMA stated that the common practice is for the regulatory authorities to collect data in the form of standard reporting templates at regular intervals. In order for data-driven, more automated processes to be the norm, the focus should shift from the creation of such templates towards the primary data that are used to construct such reports and towards the creation of reporting utilities, making data collection a matter of RegTech. In specific, it is advisable to create centralised structures which will serve both "as a common database of reported granular data but also as a repository of the interpretation of reporting rules in a format that is readable by computers", thus providing an alternative to template creation and to manual labour in the form of clean and accessible data. Furthermore, "machine-readable regulations, in particular in the field of regulatory reporting" should be created. To do so, the information received from market participants should be standarised and codified, so that the data will be easier to manage, retrieve, and handle. This case once more illustrates the need for regulators to specify their instructions in a more practical way, as said in the previous section, too.

6.2.3 Inevitability and Adoption Rate

In the 2019 report on Trends, Risks, and Vulnerabilities, ESMA finds that there is increased demand for automated software that will be able to make decisions without the need for human intervention, and at the same time increased supply due to the evolution in computer capacity and architecture. The fact that the financial sector is becoming increasingly digitalised, and also the computation power being such that enables experts to handle big data, are both contributing to a data-driven norm, making it inevitable to avoid the use of artificially intelligent systems for regulation and surveillance (Broaders and Prenio, 2018). According to ESMA, "As we move towards this more intense, data-driven supervisory process, supervisors and regulators need to adapt. Failure to do so risks undermining years of work involved in implementing regulations".

Moreover, the rate of adoption of artificially intelligent systems for market supervision is crucial as well. On the one hand, if it occurs too fast, smaller firms will not be able to adapt due to increased upfront costs at the time of the adoption of AI systems as well as significant expenses while maintaining it on the following years. ESMA reports of 2021 and 2023 reveal that there are fears of smaller firms missing out the AI-adoption race, as we will read in more depth in the next section, but so far the rate of penetration of artificially intelligent systems is not alarming yet. On the other hand, if the adoption rate is too slow, ESMA expresses concerns about regulators being left behind manipulators, thus jeopardising years of attempts to make the stock market a safe place for participants, where transparency prevails.

6.2.4 Adaptability

Another equally important matter is that of adaptability to an ever-changing market. As markets transform with the passing of time, so do the tools and processes of manipulating them, and AI models need constant updates and training to the new ways in order to remain effective. This very fact is an additional cost of the application of AI surveillance methods to the ones we read above, and could possibly delay further or even deter its adoption at large, at least on behalf of smaller firms. Unfortunately, from the viewpoint of upfront cost and further research and maintenance expenses, the feature of adaptability comes in direct contrast with inevitability of AI adoption. Apparently, one should not forget that the surveillance of the stock market is a complete ecosystem that it is currently expanding, and there is both the need to integrate existing methods and infrastructure with the new ones while we are heading to a highly digitalised transformation, as well as the necessity to constantly update and upgrade these systems to stay afoot of manipulators. Broeders and Prenio (2018) provide many details about recommended initiations that will enable regulators to "migrate to a digital based supervisory process" and also handle the increasing amounts of data and avoid being overflown. One such initiative the Data Science Hub, which was set up by the Netherlands Bank, and definitely more should follow.

Now that the last feature, adaptability, has been presented, it would be useful to recap a few selected aspects which, when put together, can pose a major concern for stock market regulators and participants alike. The facts that we have presented so far are these: Artificial intelligence will be inevitably adopted at a large scale in the stock market, including matters of surveillance; it should not be otherwise, because the integrity of the market will be at risk due to the parallel evolution and sophistication of manipulations tactics and practices; the various costs are too high for the average firms; however, the rate of adoption so far is not threatening for not so large firms, at least for now; several benefits of artificial intelligence, such as the feature of adaptability, may be turned into drawbacks due to budget limitations. This is a rather concerning future development, and it has not escaped the attention of the European regulation bodies, as we will read in the next section.

6.2.5 Reduced Competition and Concentration into the hands of the few

Another worrying aspect of the adoption of artificial intelligence at a large scale into the financial markets is expressed by ESMA in its 2023 report. The European Securities and Markets Authority is deeply concerned about a potential concentration of such systems and models “among a few big players”, which can impose a systemic risk. This is feared due to the possible scenario that, in order to deal with the cost of AI adoption, many will resort to outsourcing, becoming dependent on few large asset managers who can easily afford the expenses on the latest technology, vast amounts of data, the necessary infrastructure, maintenance, research, and talent acquisition. What is more, the European Supervisory Authorities has already pointed out the risks due to “dominance of certain providers” in the 2021 Advice on Digital Finance, which also applies to the sector of digital financial services in general, and in this broader context Prenio and Yong (2021) and Varoufakis (2023) warns that “over-reliance on third-party service providers could also lead to commercial capture and dependency risk”. According to ESMA’s report on Trends, Risks and Vulnerabilities in 2023, “A survey of finance professionals revealed that only large firms could afford to dedicate the resources necessary to implement fintech methodologies that have uncertain cost-benefit trade-offs at this stage of their development, such as those based on AI (CFA Institute, 2020). This may also be true of entities involved in post-trade processes, namely CSDs and CCPs. This not only reduces competition, but also leads to a concentration of risks in a small number of firms” (ESMA, 2021, paragraph 30).

Perhaps an answer to the above concerns is what is proposed to the Senate Select Committee, Parliament of Australia, after an inquiry towards Australian entrepreneurs and small businesses that operate in the RegTech and FinTech fields. In specific, in the request for suggestions to the Australian Government regarding the facilitation of the continuing growth of these industries, the business world responded: “The single, most valuable assistance that the Australian Government can provide and Australian RegTech business is to engage with the vendors of the technology. Ideally, an Australian provider should be given priority by the government when offering the same service as a non-Australian provider. A separate budget should be set to pay for the services provided by Australian providers of RegTech software and a quota (money and time) enforced to ensure that meaningful engagement occurs. Technology decision-makers in government departments appear to have an overriding concern about what might happen if they make the wrong decision about purchasing technology from smaller local providers. This is often enunciated of as "nobody gets fired for buying IBM". This leads to overspending and a lack of support for local innovative solutions". If we take into consideration the enormous investment that cutting-edge technology requires, as it was briefly but vividly outlined in previous chapters, it is inevitable for such firms to seek funding. However, although the majority of angel investors seems to understand the long-term horizon of the business, venture capitalists "impose strong views on the management of the business, and are only focused on an exit". In brief, if we are to inverse the possible trend towards monopolisation of the field of supervision and regulation technologies, governments may need to step up and assist smaller businesses, in order to help retain a certain degree of competition.

Chapter 7: Conclusions

Regarding the technical aspect of things, although the technological advances in artificial technology and its incorporation into the field of finance, and in specific in stock market monitoring systems, the traditional statistical and econometric methods still hold a prominent position in the analyst's set of tools. As we analytically saw in the previous chapters, the main reasons for this are their transparency and the cost effectiveness, as well as the fact that they have been tried and tested for quite a long time, thus they have proved to be credible tools. On the other hand, the breakthrough of AI can address limitations of these traditional surveillance methods, such as on matters of scalability and data quality, to name the most important ones. Probably, the optimal scenario would be the adoption of a system that utilises the advantages of both tools and can provide us with credible, effective, and affordable results. So far however, the facts point out that the life-time experience of an analyst, their creativity and insightfulness, the human touch in short, must have the last word in deciding whether a suspicious incident flagged by quantitative tools of any kind was manipulation or not.

A more important conclusion when comparing the two manipulation detection approaches might not be the effectiveness and precision of the tools, but the potential change of the status quo due to the inevitability of AI adoption. The concerns that are expressed by the European authorities regarding the decrease in competition are rather worrying and require immediate action. However, the rate of adoption of AI is such that provides valuable time for both the regulatory bodies and the governments as well, in order for a sustainable course of action to be constructed.

On the above grounds, perhaps future quantitative research should focus on developing methods which will improve the transparency and interpretability of artificial models, so that their major liability can be overcome, as well as the seamless integration of both approaches into a new, hybrid set models. Moreover, the field of computer programming could conduct further investigations regarding the quality of training data and/or the improvement of a model's robustness towards data quality. Furthermore, interdisciplinary research which combines law and information technology may also be necessary. The integration and transformation of the objectives and directions of regulatory bodies into functioning code would provide a better guidance, practical assistance, and improved results. Last, but equally vital as the above, might be a theoretical research, where the different definitions of manipulation techniques and detection methods can be unified and codified into a robust taxonomy system, so that there is solid reference for future researchers.

As a last word, the detection of manipulation in a fast-changing world is a matter of both normative and positive approaches. Fields such as mathematics, statistics, econometrics, computer programming, economics, finance and the law are only different sides of one goal, and they should co-ordinate and co-operate if the battle of transparent, competitive markets is to be won.

8 References

Aggarwal, R., & Wu, G. (2006). Stock market manipulation. *Journal of Business 79*, 1915-1953.

Aggrawal, A., & Cooper, T. (2015). Insider trading before accounting scandals. *Journal of Corporate Finance 34*, 169-190

Airfocus. (n.d.). *What is a software feature?* Airfocus. https://airfocus.com/glossary/what-is-software-feature/

Aitken, M., Cumming, D.J., & Zhan, F. (2015). High frequency trading and end-of-day price dislocation. *Journal of Banking and Finance 59*, 330-349

Allen, F., & Gale, D. (1992). Stock price manipulation. *Review of Financial Studies 5*, 503-529.

Aristotle, Politics, Book I, Section 1259a: https://www.perseus.tufts.edu/hopper/text?doc=Aristot.+Pol.+1.1259a&redirect=true

Armstrong, P., (2019). Financial Innovation: RegTech and SupTech-change for markets and authorities. *Trends, Risks and Vulnerabilities No1*. European Securities and Markets Authority (ESMA)

Bohme, R., & Holz, T. (2006). *The effect of Stock Spam on Financial Markets* (Working Paper). University of Mannheim.

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter Mood Predicts the Stock Market. *Journal of Computational Science, 2(1)*, 1-8

Broeders, D., & Prenio, J. (2018). Innovative technology in financial supervision (suptech) –the experience of early users. *FSI Insights No. 9*.

Buchan, J. (2018). *John Law: A Scottish adventurer of the eighteenth century*. Maclehose Press.

Burns, B., Oppenheimer, D., & Brewster, C. (2016). *Kubernetes: Up and running: Dive into the Future of Infrastructure*. O'Reilly Media

CFA Institute. (2020). *Fintech Survey 2: Has AI and machine learning adoption advanced?* Charlottesville, VA

Chakraborty, A., & Yilmaz, B. (2004). Informed manipulation. *Journal of Economic Theory 114*, 132-152.

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41(3)*, 1-72

Cherian, J. A., & Jarrow, R. A. (1995). Market manipulation. In *Handbooks in Operations Research and Management Science* (Vol. 9, pp. 611-630). American Economic Association

Chief Risk Officers (CRO) Forum (2019). *Machine Decisions: Governance of AI and Big Data Analytics. Page 7*. https://www.thecroforum.org/wp-content/uploads/2019/05/CROF-Machine-Decisions-Governance-of-AI-and-Big-Data-Analytics.pdf

Cumming, D., Dannhauser, R., & Johan, S. (2015). Financial market misconduct and agency conflicts: A synthesis and future directions. *Journal of Corporate Finance, Volume 34*, Pages 150-168, ISSN 0929-1199, https://doi.org/10.1016/j.jcorpfin.2015.07.016

Cumming, D., Johan, S., & Li, D. (2011). Exchange trading rules and stock market liquidity. *Journal of Financial Economics 99*, 651-671

De La Vega, J. (1688). *Confusion de confusiones*. (H. Kellenbenz, Trans.). Kress Library of Business and Economics. (Kress Library publication No.13), 1957, [i] p.42

Dyck, A., Morse, A., & Zingales, L. (2010). Who blows the whistle in corporate fraud? *Journal of Finance 65*, 2063-2253

Dyck, A., Morse, A., & Zingales, L. (2014). *How pervasive is corporate fraud?* (Working Paper). University of Chicago.

Engle, R. F. & Ng, Victor K. (1993). Measuring and testing the impact of news on volatility. *Journal of Finance. 48* (5): *1749*–1778

European Securities and Markets Authority (ESMA). (2023). *Artificial intelligence in EU securities markets: ESMA Report on Trends, Risks and Vulnerabilities* (ESMA Report No. ESMA50-164-6247). https://www.esma.europa.eu/sites/default/files/library/ESMA50-164-6247-AI_in_securities_markets.pdf

European Securities and Markets Authority (ESMA). (2023). *ESMA TRV Risk Analysis: Artificial intelligence in EU securities markets* (ESMA Report No. ESMA50-164-6247). https://www.esma.europa.eu/sites/default/files/library/ESMA50-164-6247-AI_in_securities_markets.pdf

European Supervisory Authorities (ESAs). (2022). *Joint European Supervisory Authority response to the European Commission's February 2021 call for advice on digital finance and related issues: Regulation and supervision of more fragmented or non-integrated value chains, platforms and bundling of various financial services, and risks of groups combining different activities.*(ESA 2022 01). https://www.esma.europa.eu/press-news/esma-news/esas-recommend-actions-ensure-eu%E2%80%99s-regulatory-and-supervisory-framework-remains

European Union. (2014). *Regulation (EU) No 596/2014 of the European Parliament and of the Council Of 16 April 2014 on market abuse (Market Abuse Regulation).* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02014R0596-20210101

Federal Trade Commission & US Department of Justice. (2000). *Antitrust Guidelines for Collaborations Among Competitors*.https://www.ftc.gov/sites/default/files/documents/public_events/joint-venture-hearings-antitrust-guidelines-collaboration-among-competitors/ftcdojguidelines-2.pdf

Financial Conduct Authority. (n.d.). *MAR 1.6: Manipulating Transactions*. FCA Handbook. Retrieved [January 2025] from https://www.handbook.fca.org.uk/handbook/MAR/1/6.html

Financial Industry Regulatory Authority. (n.d.). Ramp and dump schemes. Retrieved [January 2025] from https://www.finra.org/investors/insights/ramp-and-dump-schemes

Fischel, D. R., & Ross, D. J. (1991). Should the law prohibit "manipulation" in financial markets? *Harvard Law Review, vol. 105*, no. 2, pp. 503–553

Fishman, M. J., & Haggerty, K. (1995). The mandatory disclosure of trades and market liquidity. *Reviews of Financial Studies 8*, 637-676.

Frankq, C., & Zakoian, J-M. (2019). *GARCH Models: Structure, statistical inference and financial applications* (2nd ed.). John Wiley & Sons

FreeCodeCamp. (2021, January 28). *Chihuahua or Muffin Problem: My search for the best computer vision API*. FreeCodeCamp. https://www.freecodecamp.org/news/chihuahua-or-muffin-my-search-for-the-best-computer-vision-api-cbda4d6b425d/

Frieder, L., & Zittrain, J. (2007). *Spam Works: Evidence from Stock Touts and Corresponding Market Activity*. (Working Paper). Purdue University.

Galton, F. (1885). Regression towards Mediocrity in Hereditary Stature. *Journal of the Anthropological Institute of Great Britain and Ireland*.

Government of British Columbia. (n.d.). *Securities, Chapter 348*. https://www.bclaws.gov.bc.ca/civix/document/id/hstats/hstats/61446744

Greene, W. H. (2012). *Econometric Analysis*. (7th ed.). Pearson Education

Hanke M., & Hauser, F. (2008). On the effect of stock-spam e-mails. *Journal of Financial Markets 11*, 57-83.

International Organization of Securities and Commissions. (2017). *Objectives and principles of securities regulation* (p. 7). Principles for cooperation in regulation https://www.iosco.org/library/pubdocs/pdf/ioscopd561.pdf

Jarrow, R. A. (1992). Market manipulation, bubbles, corners, and short squeezes. *Journal of Financial and Quantitative Analysis*, 27(3), 311–336

Jiang, G., Mahoney, P. G., & Mei, J. (2004). Market Manipulation: A comprehensive study of stock pools. *Journal of Financial Economics 77*, 147-170.

Jiang, G., Mahoney, P. G., & Mei, J. (2005). Market manipulation: A comprehensive study of stock pools. *Journal of Financial Economics, Volume 77* (1), 147-170, ISSN 0304-405X, https://doi.org/10.1016/j.jfineco.2004.10.005

Karpoff, J., Lee, D. S., & Martin, G.S. (2008a). The consequences to managers for cooking the books. *Journal of Financial Economics 88*, 193–215

Kondeth, A. (2020, July 1). *Solving the famous problem of Puppy vs Muffin by deep learning – Web App on Google App Engine/Fastai*. Medium. https://medium.com/@ahffank/solving-the-famous-problem-of-puppy-vs-muffin-using-a-simple-web-application-on-google-app-a3264e2be49c

Kyle, A. S., & Viswanathan, S. (2008). How to define illegal price manipulation. *American Economic Review, vol. 98*, no. 2, pp. 274–279, 2008.

Lee, E. J., Eom, K. S., & Park, K. S. (2013). Microstructure-based manipulation: Strategic behaviour and performance of spoofing traders. *Journal of Financial Markets, Volume 16.* Issue 2, 2013, Pages 227-252, ISSN 1386-4181, https://doi.org/10.1016/j.finmar.2012.05.004

Legality. (n.d.) In *Business Law I.* RV Community College. https://rvcc.pressbooks.pub/businesslaw131/chapter/chapter-9-legality/

Liu, X. F. (2012). Dynamics of network of Global Stock Markets. *Accounting and Finance Research*. https://www.sciedu.ca/journal/index.php/afr/article/view/1233/694

MacKinlay, A. C. (1997). Event Studies in Economics and Finance. *Journal of Economic Literature*, *35*(1), 13–39. http://www.jstor.org/stable/2729691

Mahoney, P. (1999). The stock pools and the securities exchange act. *Journal of Financial Economics 51*, 343-369

Newman, M. E. J. (2010). *Networks: An introduction.* Oxford University Press

Pearson, K. (1896). *Mathematical Contributions to the Theory of Evolution: On the Growth of the Stature of Man*.

Parliament of Australia. (n.d.). *Financial Technology and Regulation Technology: Submission 34* https://www.aph.gov.au/DocumentStore

Pickvance, G. C. (2001) Four Varieties of comparative analysis. *Journal of Housing and the Built Environment 16*. 7-28

Prenio, J., & Yong, J. (2021). *Humans keeping AI in check – emerging regulatory expectations in the financial sector* (p.34). Financial Stability Institute, Bank for International Settlements https://www.bis.org/fsi/publ/insights35.htm

Pulbrook, A. (1865). *The Companies Act, 1862, with analytical references and copious index* . Internet Archive. https://archive.org/details/companiesactwit00pulbgoog/page/n66/mode/2up

Royal Statistical Society. (n.d.). *History of statistics on timeline.* https://ssca.org.in/media/History_of_Statistics_on_Timeline.pdf

The Guardian. (2013, January 28). *Swift Trade fined over market abuse.* https://www.theguardian.com/business/2013/jan/28/swift-trade-fined-over-market-abuse

Turban, E., & Aronson, J. E. (2018). *Decision Support Systems and Intelligent Systems*. Pearson.

U.S. Commodity Futures Trading Commission. (n.d.). *CFTC whistleblower alert: Blow the whistle on spoofing in the commodities and derivatives markets*. https://www.whistleblower.gov/whistleblower-alerts/Spoofing_WBO_Alert.htm

US Securities and Exchange Commission. (1934). *Securities Exchange Act of 1934*. https://www.columbia.edu/~hcs14/SX9.htm

Van Bommel, J. (2003). Rumours. *Journal of Finance 58*, 171-218.

Varoufakis, Y. (2023). *Techno Feudalism. What killed capitalism*. Penguin Random House

Verstein, A. (2015). *Benchmark manipulation*. https://heinonline.org/HOL/LandingPage?handle=hein.journals/bclr56&div=7&id=&page=

Zulkifley, M. A., Munir, A. F., Sukor, M. E. A., & Shafiai, M. H. M., (2023). *A Survey on Stock Market Manipulation Detectors Using Artificial Intelligence*. Tech Science Press. https://doi.org/10.32604/cmc.2023.036094

9 Glossary

AI: Artificial Intelligence

AR: Abnormal Returns (in Event Analysis)

AR: Auto-Regressive (in econometrics)

ARCH: Auto-Regressive Conditional Heteroskedasticity

ARIMA: Auto-Regressive Integrated Moving Average

ARMA: Auto-Regressive Moving Average

CAR: Cumulative Abnormal Returns

CFA: Financial Conduct Authority (UK)

CFTC: Commodity Futures Trading Commission (US)

CRO: Chief Risk Officers

ESAs: European Supervisory Authorities

ESMA: European Securities and Markets Authority

FCA: Financial Conduct Authority (UK)

FINRA: Financial Industry Regulatory Authority (US)

FSB: Financial Stability Board

FTC: Federal Trade Commission (US)

GARCH: Generalised Auto-Regressive Conditional Heteroskedasticity

GMA: ticker name of the Gamestop stock

HQIC: Hannan-Quinn Information Criterion

IGARCH: Integrated Generalised Auto-Regressive Conditional Heteroskedasticity

IOSCO: International Organisation of Securities and Commissions

IPO: Initial Public Offering

IT: Information Technology

LSE: London Stock Exchange

LSTM: Long Short-Term Memory

MA: Moving Average

MAR: Market Abuse Regulation (EU)

ML: Machine Learning

MOASS: Mother of all short squeezes

NGARCH: Non-linear Generalised Auto-Regressive Conditional Heteroskedasticity

NLP: Natural Language Processing

OECD: Organisation for Economic Co-operation and Development

OTC: Over the counter transaction

RegTech: Regulation Technology

RF: Random Forest

RSS: Royal Statistics Society

S&P500: The Standard and Poor's 500 stock market index

SEC: Securities and Exchange Commission (US)

SupTech: Supervision Technology

SWOT: Strengths, Weaknesses, Opportunities, Threats

UK MAR: Market Abuse Regulation for the United Kingdom

10 Appendix

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
import yfinance as yf
from statsmodels.tsa.arima.model import ARIMA
from sklearn.ensemble import RandomForestRegressor
from sklearn.model_selection import train_test_split
from sklearn.metrics import mean_squared_error
import seaborn as sns
from sklearn.preprocessing import MinMaxScaler
import warnings

warnings.filterwarnings("ignore", category=FutureWarning)
warnings.filterwarnings("ignore", category=UserWarning)

try:
    data = yf.download('GME', start='2021-01-01', end='2021-03-31')
    if data.empty:
        raise ValueError("No data downloaded for GME within the specified date range.")
except Exception as e:
    print(f"Error downloading data: {e}")
    exit()

data.ffill(inplace=True)
data.bfill(inplace=True)

data['Prev_Close'] = data['Close'].shift(1)
data['Volume_Lag1'] = data['Volume'].shift(1)
data['Close_Diff'] = data['Close'].diff()
data['Volume_Diff'] = data['Volume'].diff()

data.dropna(inplace=True)

train_size = 0.8
train_data = data[:int(len(data) * train_size)]
test_data = data[int(len(data) * train_size):]


try:
    model_arima = ARIMA(train_data['Close'], order=(5, 1, 0))
```

```
    model_fit_arima = model_arima.fit()
    arima_forecast = model_fit_arima.predict(start=len(train_data), end=len(data)-1)
except Exception as e:
    print(f"ARIMA Model Error: {e}")
    arima_forecast = None


try:
    features = ['Prev_Close', 'Volume_Lag1', 'Close_Diff', 'Volume_Diff']
    target = 'Close'

    scaler = MinMaxScaler()
    X_train = scaler.fit_transform(train_data[features])
    X_test = scaler.transform(test_data[features])
    y_train = train_data[target]
    y_test = test_data[target]

    y_train = y_train.values.ravel()
    y_test = y_test.values.ravel()

    model_rf = RandomForestRegressor(n_estimators=100, random_state=42)
    model_rf.fit(X_train, y_train)
    y_pred_rf = model_rf.predict(X_test)

    mse_rf = mean_squared_error(y_test, y_pred_rf)
    print(f'Mean Squared Error (Random Forest): {mse_rf}')
except Exception as e:
    print(f"Random Forest Model Error: {e}")
    y_pred_rf = None
    mse_rf = None


try:
    if arima_forecast is not None:
        mse_arima = mean_squared_error(test_data['Close'], arima_forecast)
        print(f'Mean Squared Error (ARIMA): {mse_arima}')
    else:
        mse_arima = None
except Exception as e:
    print(f"ARIMA Evaluation Error: {e}")
    mse_arima = None
```

```
plt.figure(figsize=(14, 7))
plt.plot(data.index, data['Close'], label='Historical Prices', color='blue')

if arima_forecast is not None:
    plt.plot(test_data.index, arima_forecast, label='ARIMA Forecast', color='orange')


if y_pred_rf is not None:
    plt.plot(test_data.index, y_test, label='Actual Prices (RF)', color='green')
    plt.plot(test_data.index, y_pred_rf, label='Predicted Prices (RF)', color='red')


plt.title('GameStop Stock Price: Historical, ARIMA Forecast, and Random Forest Predictions')
plt.xlabel('Date')
plt.ylabel('Price')
plt.legend()
plt.xticks(rotation=45)
plt.tight_layout()
plt.show()


if y_pred_rf is not None:
    importances = model_rf.feature_importances_
    feature_names = features
    sns.barplot(x=importances, y=feature_names)
    plt.title('Feature Importance (Random Forest)')
    plt.show()


if model_fit_arima:
    print(model_fit_arima.summary())
```