

UNIVERSITÀ POLITECNICA DELLE MARCHE
FACOLTÁ DI ECONOMIA "GIORGIO FUÁ"

BACHELOR'S DEGREE PROGRAM IN

DIGITAL ECONOMICS AND BUSINESS

Artificial Intelligence and Online Trading: Impacts on Financial Market Analysis

Thesis Supervisor

Prof. Francesco James Mazzocchini

Student

Francesco Carbonetti

Academic year 2025/2026

TABLE OF CONTENTS

Table of Contents…… i

List of Figures and Tables …… iii

Introduction 1

Chapter 1 – Financial Markets and Online Trading 3

1.1 Definition of financial markets 3
1.2 Functions of financial markets 5
1.3 The evolution of online trading 6
1.4 Digital trading platforms 8
1.5 Financial instruments 9

Chapter 2 – Artificial Intelligence in Finance 10

2.1 Definition and development of artificial intelligence in finance 10
2.2 Main applications of AI in financial services 11
2.3 Machine learning and predictive algorithms 13
2.4 Automation and financial data analysis 14

Chapter 3 – AI-Driven Trading 16

3.1 Algorithmic trading and automated trading 16
3.2 The use of AI in market data analysis 17
3.3 Sentiment analysis and trend prediction 18
3.4 Advantages, limits, and risks of AI in trading 24

Chapter 4 – Impacts on Financial Markets and Regulatory Issues.....26

4.1 Effects on market efficiency and execution speed.....26

4.2 Impact on volatility and liquidity.....31

4.3 Changes in investor behaviour.....38

4.4 Ethical and regulatory aspects.....42

Conclusions 48

Bibliography 51

LIST OF FIGURES AND TABLES

Chapter III 17

Table 3.1: Dual Sentiment Index: historical backtest performance by signal type 22

Table 3.2: Dual Sentiment Index: individual indicator readings as of 2 April 2026 23

Table 3.3: Dual Sentiment Index: composite index status as of 2 April 2026 23

Chapter IV 27

Table 4.1: Intraday liquidity measures for CAC40 stocks, opening vs. closing half-hour 33

Table 4.2: HFT intensity (quote-to-trade ratio) and volatility, stable vs. crash periods, CAC40 large-cap stocks 34

Table 4.3: Selected regulatory instruments addressing algorithmic and high-frequency trading 46

INTRODUCTION

Over the past two decades, financial markets have undergone a profound technological transformation. Trading that once relied on floor-based intermediation and telephone-based order placement has been replaced, in large part, by electronic systems capable of executing transactions within microseconds. This shift has been accompanied, more recently, by the growing integration of artificial intelligence into the analysis of financial data and into the execution of trading strategies themselves. Machine-learning techniques are now used to process market data, extract signals from unstructured information such as news and social media and generate trading decisions with limited or no direct human intervention. This thesis examines how these developments have reshaped online trading and, more broadly, the analysis of financial markets.

The growing prominence of artificial intelligence in trading raises questions that are simultaneously technical, economic, and regulatory in nature. From a technical standpoint, it is necessary to understand which AI techniques are actually used in financial markets and how they differ from earlier forms of algorithmic and automated trading. From an economic standpoint, it is important to assess whether AI-driven and high-frequency trading improve the functioning of financial markets — for instance by narrowing bid–ask spreads, deepening order books, and accelerating the incorporation of information into prices — or whether they introduce new sources of instability, particularly during periods of market stress. From a regulatory standpoint, the increasing reliance on automated and AI-based systems raises questions of transparency, accountability, and fairness among market participants with unequal technological capacities. This thesis is motivated by the observation that these three dimensions — technical, economic, and regulatory — cannot be fully understood in isolation from one another.

This interest was reinforced by my curricular internship at Cube Investimenti, during which I was able to observe directly how a sentiment-based, rule-driven trading model is designed, backtested, and monitored in practice. This experience motivated the inclusion of a concrete case study illustrating how the principles discussed in the academic literature on AI-driven trading are translated into an operational tool used by a financial intermediary.

The objective of this thesis is therefore twofold. First, it aims to provide a structured overview of how artificial intelligence has been integrated into financial-market analysis and trading, tracing this development from the earlier rise of electronic and online trading platforms. Second, it aims to assess, on the basis of the empirical literature, the consequences of AI-driven and high-

frequency trading for market efficiency, volatility, liquidity, investor behaviour, and financial regulation, with particular attention to episodes of market stress in which the effects of automated trading appear to diverge most sharply from those observed under ordinary conditions.

The thesis is organized as follows. The First Chapter defines financial markets and situates them within the broader financial system, before tracing the historical shift from floor-based trading to electronic and online trading platforms. The Second Chapter examines the development of artificial intelligence and its main applications within the financial sector, including machine learning, natural language processing, and automated data analysis. The Third Chapter turns to AI-driven trading specifically, covering algorithmic execution, the use of AI in market-data analysis, and sentiment-based trend prediction, and includes a case study based on a proprietary sentiment index. Last Chapter, the fourth one, assesses the consequences of AI-driven and high-frequency trading for market efficiency, volatility and liquidity, investor behaviour, and the ethical and regulatory issues these technologies raise. The thesis concludes by drawing together the main findings and considering their broader implications for financial-market analysis and online trading.

CHAPTER 1
FINANCIAL MARKETS AND ONLINE TRADING

1.1 Definition of Financial Markets

Financial markets arise to reduce the information costs of borrowing and lending and of making transactions, and in so doing, financial systems serve a number of functions that are essential in a modern economy: they provide payment services that facilitate the exchange of goods and services, mobilize savings, allocate credit, and monitor borrowers (World Bank, 2002). From a functional perspective, this definition can be extended: financial markets are best understood not as a fixed set of institutions but as part of a broader financial system in which financial functions act as the stable "anchors" or "givens," while the institutional structure built to perform them — financial institutions, financial markets, products, services, and supporting infrastructure such as regulatory rules and the accounting system — is determined endogenously and evolves in response to changes in technology, transaction costs, and behaviour (Merton & Bodie, 2004). Because the particular institutions and organizational forms that arise within the financial system develop as an endogenous response to minimize the costs of transaction frictions and behavioral distortions in executing the financial functions common to every economy, financial markets should not be equated with any single, stable structure; rather, they represent a continuously adapting architecture built to execute the same essential functions more efficiently over time (Merton & Bodie, 2004). Beyond this functional reading, financial markets have also been characterized in heterodox terms: in a world of nonergodic uncertainty, the primary function of organized financial markets is to provide liquidity, understood as the ability to buy and resell assets in a well-organized, orderly market in order to obtain the medium of contractual settlement needed to meet one's nominal contractual liabilities when they come due (Davidson, 2000). These perspectives converge on a single conclusion: financial markets should be understood not merely as neutral mechanisms for allocating capital to its most productive uses, but as evolving institutional systems whose structure reflects both the functions they are designed to serve and the degree of confidence and orderliness that market participants require to operate within them.

A useful way to further specify this definition is by considering the broader financial system within which markets operate, alongside financial intermediaries and the regulatory infrastructure that

supports them. By evaluating alternative investments and monitoring the activities of borrowers, financial intermediaries — such as banks, insurance companies, provident and pension funds, and mutual funds — overcome the information problems that individual agents would otherwise face if they traded directly with one another, thereby increasing the efficiency of resource use (World Bank, 2002). While financial markets provide the venues in which claims are priced and exchanged, it is often intermediaries that originate and structure these claims in the first place, meaning that markets and intermediaries should be viewed as complementary rather than competing components of the same financial system (World Bank, 2002). A complementary distinction concerns the overall orientation of a financial system: although banks dominate most formal financial systems, the relative importance of the stock market tends to increase with the level of development, even though far more finance is typically raised from bank loans than from selling equity, even in industrial countries (World Bank, 2002). Taken together, these distinctions show that "financial markets" cannot be fully understood in isolation from the intermediaries and institutional arrangements that surround them, since it is the interaction between markets, intermediaries, and regulation that determines how efficiently a financial system performs its core functions.

A third relevant distinction concerns the theoretical lens through which the functioning of financial markets is interpreted, since this shapes how their stability and efficiency are assessed. According to efficient-market theory, the ultimate social functions of financial markets are spreading risks, guiding the investment of scarce capital, and processing and disseminating the information possessed by diverse traders, so that prices will always reflect fundamental values and any short-run instability merely calls for limited corrective adjustments.

By contrast, liquidity preference theory holds that the primary function of financial markets is not efficiency but the provision of liquidity in a fundamentally uncertain world, meaning that market orderliness must be actively maintained — for example through market makers or capital-flow regulations — rather than assumed to emerge spontaneously (Davidson, 2000). Because sudden shifts in market sentiment can overwhelm this orderliness and generate instability that has little to do with changes in economic fundamentals, this perspective implies that observed financial crises are evidence of a structural need for supporting institutions rather than mere frictions on an otherwise self-correcting system (Davidson, 2000). While efficient-market theory has historically underpinned the liberalization of financial markets worldwide, the recurrence of financial crises

since the 1970s has lent renewed weight to the liquidity-based view, a debate that becomes particularly relevant when assessing the resilience of contemporary financial markets, including those increasingly mediated by digital trading platforms.

With financial markets now defined and situated within the broader financial system — distinguishing between the functions they perform, their relationship with financial intermediaries, and the competing theoretical perspectives on their stability — the next step is to examine the specific functions that financial markets perform within the economic system.

1.2 Functions of Financial Markets

Beyond serving as venues for the issuance and exchange of financial claims, financial markets are best interpreted in relation to the broader functions of the financial system. Levine (2005) groups these functions into five areas: reducing the informational barriers that affect investment decisions and the allocation of funds; overseeing financed firms and supporting corporate governance; enabling risk transfer, diversification, and risk management; collecting dispersed savings and making them available for investment; and lowering the costs of exchanging goods and services. In this perspective, financial development does not imply the emergence of entirely new purposes. Rather, it consists in the improved capacity of financial instruments, markets, and intermediaries to perform these enduring functions by mitigating information, enforcement, and transaction costs (Levine, 2005).

This approach is further formalized in the Functional and Structural Finance (FSF) framework of Merton and Bodie (2004). The framework regards the underlying functions of finance, such as transferring resources over time and across locations, supporting payment and settlement processes, facilitating the pooling and management of risk, and contributing to price discovery, as the stable reference points of a financial system. By contrast, the institutional arrangements used to carry out those functions are not permanent. Financial institutions, markets, instruments, services, and supporting infrastructure are endogenously determined and may evolve as technological conditions, transaction costs, behavioral patterns, and other features of the economic environment change (Merton & Bodie, 2004). Financial markets should therefore be viewed as one among several institutional mechanisms through which a relatively stable set of economic functions can be performed, rather than as a fixed or inherently necessary organizational structure.

Financial contracts constitute the practical instruments through which these functions are implemented. Chami, Fullenkamp, and Sharma (2009) define financial instruments as contracts between borrowers and lenders governing the transfer, use, and repayment of funds. Their framework also highlights the role of liquidity providers and regulators, whose incentives, constraints, and actions influence whether markets can develop and function effectively. The form of a financial contract consequently reflects the terms on which participants can reconcile their different interests within the prevailing economic, legal, social, technical, and regulatory setting (Chami, Fullenkamp, & Sharma, 2009). From this perspective, the characteristics of financial instruments and market structures are shaped by negotiated compromises under real-world constraints, rather than by the frictionless allocation mechanisms assumed in purely theoretical models. Overall, these contributions suggest that the core purposes of finance — improving information, directing funds toward investment, managing and distributing risk, mobilizing savings, and facilitating exchange — are persistent, whereas the institutional and contractual arrangements that deliver them continuously adapt to technological progress, regulation, and changing market conditions. This dynamic is particularly relevant to the analysis of the increasing role of digital and algorithmically driven trading platforms developed in the following chapters.

1.3 The Evolution of Online Trading

The organization of stock trading has undergone a fundamental transformation over the past several decades, shifting from a system centered on physical exchange floors and human intermediaries to one increasingly conducted through electronic systems. Historically, exchanges such as the New York Stock Exchange operated as member organizations in which trading took place through face-to-face negotiation at a physical location. Brokers carried orders to designated trading posts, where specialists helped match buyers and sellers and supported market liquidity (Stoll, 2006). From the 1970s onward, and particularly during the 1990s and early 2000s, this model was increasingly challenged by fully electronic markets, including electronic communications networks (ECNs). These systems automated the processing of orders, from submission and routing to execution and confirmation, thereby reducing the need for direct human intervention and lowering execution times and several components of trading costs (Stoll, 2006).
This transformation in market infrastructure was accompanied by a parallel change on the retail side, driven by the emergence of the Internet as a channel for both order execution and investment

information. Barber and Odean (2001) argue that online brokerage reduced the fixed and marginal costs of providing financial services, enabling smaller and newer firms to compete with established providers. Online brokerages replaced many telephone-based interactions with computerized systems, substantially lowering the resources required to process trades and allowing investors to place orders without relying on a personal broker. Between 1995 and mid-2000, investors opened 12.5 million online brokerage accounts, while online trading represented approximately 37 percent of U.S. retail trading volume in equities and options in 1998 (Barber & Odean, 2001). These developments were accompanied by a marked reduction in commissions and by broader access to market information and execution services.

The movement toward online trading did not simply reduce the cost and time required to place orders; it also changed the environment in which individual investment decisions were made. Barber and Odean (2002) examine 1,607 investors who moved from telephone-based to online trading between 1992 and 1995. They find that, after the switch, these investors traded more actively and more speculatively, while their net performance deteriorated relative to the period before they began trading online and relative to comparable investors who continued to trade by telephone. The authors argue that lower trading costs, faster execution, and easier access alone cannot explain this outcome, since these factors would not ordinarily be expected to reduce investment performance. Instead, they associate the pattern with overconfidence, reinforced by self-attribution, an illusion of knowledge, and an illusion of control (Barber & Odean, 2002). The online environment provided investors with extensive information and direct control over order placement, but access to more information did not necessarily improve their ability to select securities or time trades effectively (Barber & Odean, 2001).

Read together, these developments reveal two related but analytically distinct dimensions of the transition to online trading. The first was a structural transformation in market infrastructure, as order handling and execution moved from physical trading floors and predominantly human intermediation toward electronic and increasingly automated systems. The second was a behavioral transformation among retail investors, whose trading decisions were influenced by lower barriers to participation, faster execution, greater autonomy, and extensive online information. Electronic trading did not eliminate the economic functions performed by brokers, dealers, specialists, and markets; rather, it changed the institutional and technological means through which those functions were provided. These developments created important foundations

for the subsequent expansion of digital trading platforms and for the later integration of algorithmic and artificial-intelligence-based tools into financial-market activity, examined in the chapters that follow.

1.4 Digital Trading Platforms

In contemporary financial markets, a digital trading platform — more broadly referred to in the literature as an electronic trading (ET) system — can be defined as a facility that provides some or all of the following services: electronic order routing, automated trade execution, and the electronic dissemination of pre-trade information, such as bid and offer quotes and market depth, together with post-trade information, such as transaction prices and volumes (Committee on the Global Financial System, 2001). ET systems differ from traditional trading arrangements in that they can be location-neutral, support continuous multilateral interaction, scale to accommodate additional transactions and participants, and integrate different stages of the trading process through straight-through processing (STP). From an architectural perspective, two broad models can be distinguished. Centralized electronic order books use matching algorithms to combine orders submitted by multiple market participants according to predetermined priority rules. Decentralized electronic dealer systems, by contrast, automate bilateral interactions that were traditionally conducted by telephone and may take the form of single-dealer platforms, in which one intermediary provides quotes to customers, or multi-dealer platforms, in which customers can request quotes from several dealers simultaneously (Committee on the Global Financial System, 2001). In practice, electronic systems may combine features of these models, and the boundary between them is not always absolute. Their suitability also depends on the market segment, the characteristics of the instrument, the size of the transaction, and the management of counterparty credit risk. ET can therefore be understood not merely as the electronic replacement of a physical trading floor, but as a form of market infrastructure that changes how orders, information, liquidity, execution, clearing, and settlement are organized (Committee on the Global Financial System, 2001).

1.5 Financial Instruments

The suitability of a financial instrument for electronic trading depends primarily on its degree of standardization, liquidity, transaction size, and counterparty credit risk. Standardized and relatively liquid instruments, particularly those with limited counterparty risk, are generally more suitable for centralized electronic order books because they can be traded against standing orders under predefined rules. Examples include actively traded foreign-exchange instruments, government securities, and exchange-traded index futures. By contrast, heterogeneous or less liquid instruments often require greater pre-trade discussion, bilateral negotiation, and more extensive arrangements for managing counterparty exposure. The CGFS notes that centralized electronic trading is less suitable where instruments cannot be settled through effective delivery-versus-payment mechanisms or where market movements create substantial counterparty risk. Such risk may nevertheless be addressed through system-based credit limits or central counterparties, allowing some instruments that would otherwise remain dependent on bilateral dealing to be traded electronically (Committee on the Global Financial System, 2001).

A further distinction concerns the way different market participants use electronic platforms. Retail and individual investors may use online platforms to access smaller-sized and more affordable contracts, including E-mini equity-index futures, which are designed to attract participants unable or unwilling to trade full-sized contracts. Arbitrageurs and informed traders, by contrast, may value the speed, anonymity, and information efficiency of electronic systems because these characteristics can help them respond to temporary price discrepancies. Chung, Sheu, and Hsu (2010) find that E-mini index futures generally exhibit smaller pricing errors than comparable floor-traded contracts, although this advantage is weakened during periods of high volatility, when noise-trader and arbitrage risks become more significant. Institutional investors may use multi-dealer platforms to request and compare quotes from several dealers, particularly in fixed-income markets, thereby reducing search costs while preserving dealer intermediation. Dealers and financial intermediaries may use automated pricing engines to generate and update quotes for smaller or more liquid transactions, while retaining human intervention for unusually large orders or periods of exceptional volatility (Committee on the Global Financial System, 2001). These demonstrate that electronic trading is not a uniform activity: its operation and effects depend on the participant's resources, information, objectives, and relationship with intermediaries.

CHAPTER 2

ARTIFICIAL INTELLIGENCE IN FINANCE

2.1 Definition and Development of Artificial Intelligence in Finance

Artificial intelligence (AI) is a field of computer science concerned with creating intelligent machines capable of performing cognitive tasks traditionally associated with human capabilities, including reasoning, learning, taking action, and speech recognition (Bahoo et al., 2024). In the financial domain, AI refers to the use of computational techniques to support, improve, or automate financial activities and decision-making. The relevant literature encompasses a broad range of approaches, including expert systems, machine learning, neural networks, genetic algorithms, fuzzy logic, data mining, natural language processing, algorithmic trading, and robo-advisory applications (Milana & Ashta, 2021; Bahoo et al., 2024).

From a historical perspective, the development of AI dates back to the early history of electronic computing, when the concept was associated with the capacity of machines to perform correct logical operations. As electronic computing developed during the late 1940s, expectations emerged that computers would move beyond rapid calculation toward forms of performance comparable to human intelligence. These expectations, however, were not fulfilled in a linear way. The history of AI has instead been characterized by alternating phases of enthusiasm, technological advancement, disappointment, and reduced investment, often referred to as “AI winters” (Milana & Ashta, 2021; Bahoo et al., 2024).

Milana and Ashta (2021) describe this development through the logic of a technological hype cycle, in which initial innovation is followed by inflated expectations, subsequent disillusionment, gradual refinement, and eventually a more stable stage of productivity. The development of AI illustrates this pattern: early ambitions were constrained by limited computing capacity, insufficient data, and technological barriers, while later advances in expert systems, neural networks, personal computing, cloud computing, and machine learning made broader applications possible. In finance, research on AI-related techniques has existed since at least the 1980s and 1990s, but the scale and visibility of the field have expanded substantially in recent years (Milana & Ashta, 2021).

Bibliometric evidence confirms that the academic literature on AI in finance has grown markedly over time, with a particularly strong upward trend from approximately 2015 onward. This expansion is associated with the increasing availability of computational power, large and unstructured datasets, machine-learning techniques, and digital technologies that can be applied to financial problems (Bahoo et al., 2024). The literature now covers a wide range of uses, including predictive and forecasting systems, classification and early-warning models, credit-risk assessment, fraud detection, portfolio management, algorithmic trading, natural language processing, sentiment analysis, and large-scale data mining (Bahoo et al., 2024).

Taken together, these perspectives indicate that AI in finance should not be interpreted as a single, static technology. Rather, it represents an evolving family of computational methods whose relevance in financial services and markets has increased alongside advances in data availability, computing capacity, and algorithmic design. This historical and conceptual background provides the basis for examining the specific applications of AI in financial services in the following sections.

2.2 Main Applications of AI in Financial Services

Building on the definitional and historical foundations established in the previous section, the applications of AI within financial services can be organized according to the specific capabilities they exploit. Cao (2021) identifies three primary functional categories through which AI generates value in finance: prediction, optimization, and natural language understanding. Prediction applications — including demand forecasting, credit risk modeling, and fraud probability assessment — leverage machine learning to generate probabilistic estimates of future financial states with a degree of accuracy that traditional statistical approaches often struggle to match (Cao, 2021). Optimization applications, by contrast, are directed at improving the efficiency of decision-making processes such as portfolio allocation and resource management, while natural language understanding applications extract insights from unstructured textual sources, such as financial news, regulatory filings, and analyst reports, that would otherwise remain inaccessible to conventional quantitative models (Cao, 2021).

This functional taxonomy is broadly corroborated by bibliometric evidence on where AI research in finance has actually concentrated. Goodell, Kumar, Lim, and Pattnaik (2021), analyzing the

thematic structure of AI and machine learning scholarship in finance between 1986 and 2021, identify three overarching clusters of research: portfolio construction, valuation, and investor behaviour; financial fraud and distress; and sentiment inference, forecasting, and planning. The convergence between this bibliometric mapping and Cao's functional categories suggests that AI's practical uptake in financial services has closely tracked the areas in which the underlying techniques offer the clearest advantage over traditional methods — namely, tasks involving large volumes of structured or unstructured data, complex nonlinear relationships, or the need for rapid, adaptive decision-making.

One of the clearest illustrations of how these capabilities translate into a concrete, customer-facing financial service is the emergence of robo-advisory platforms. Robo-advisors use AI-driven algorithms to assess a customer's risk tolerance and financial objectives through an initial questionnaire and subsequently generate and manage a personalized investment portfolio with limited or no human intervention, a process that echoes the traditional function of a human financial advisor while substantially reducing the associated cost and time (Belanche, Casaló, & Flavián, 2019). The growing academic and commercial interest in robo-advisory services reflects a broader pattern in the adoption of AI within financial services: rather than replacing entire categories of financial professionals outright, AI applications tend to automate specific, well-defined sub-tasks — risk profiling, portfolio rebalancing, or credit scoring — that were previously performed manually, while decisions requiring greater contextual judgment continue to involve human oversight (Belanche et al., 2019).

Viewed as a whole, these perspectives show that the main applications of AI in financial services span a continuum from back-office, prediction-oriented tasks such as fraud detection and credit risk assessment, to front-office, customer-facing applications such as robo-advisory, all of which are underpinned by the same core capabilities — prediction, optimization, and language understanding — identified in the broader AI-in-finance literature. Having outlined this general landscape of applications, it is now useful to examine in greater depth the specific machine learning techniques and predictive algorithms that make these applications possible.

2.3 Machine Learning and Predictive Algorithms

Building on this broad landscape of AI applications in financial services, the discussion now turns to look more closely at the specific machine learning techniques that underpin many of these applications, particularly those concerned with prediction. Machine learning techniques applied to finance can be organized according to the learning paradigm they employ: supervised learning methods, such as regression and classification models, are used to predict prices and classify market patterns based on labeled historical data; unsupervised learning techniques, including clustering and anomaly detection, are employed to identify hidden market structures and unusual trading patterns without relying on labeled outcomes; deep learning models, particularly recurrent neural networks and long short-term memory (LSTM) networks, are commonly used for time-series forecasting given their capacity to capture temporal dependencies in sequential data; and reinforcement learning, an increasingly popular approach, enables adaptive trading strategies that continuously adjust to changing market conditions (Rohan et al., 2026).

This taxonomy of techniques is complemented by a broader methodological literature that has sought to systematize the growing body of research applying computational intelligence to financial markets. Cavalcante, Brasileiro, Souza, Nobrega, and Oliveira (2016), in one of the most widely cited surveys in this area, review the range of soft computing techniques applied to financial markets and propose a general framework for building intelligent trading systems, while also highlighting that the large volume of financial information available to traders has made the case for computational, rather than purely manual, approaches to decision-making increasingly compelling. Their survey underscores that the application of machine learning to finance is not merely a matter of adopting individual algorithms, but of integrating multiple techniques — feature selection, model training, and validation — into coherent systems capable of supporting real-world trading decisions.

The practical value of these techniques has been most rigorously demonstrated in the domain of empirical asset pricing. In an influential comparative study, Gu, Kelly, and Xiu (2020) evaluate a wide range of machine learning methods applied to the canonical problem of measuring asset risk premiums, demonstrating that machine learning forecasts generate substantial economic gains for investors relative to traditional linear models. The authors find that tree-based methods and neural

networks perform best among the techniques considered, and trace this superior performance to their ability to capture nonlinear interactions among predictive variables — interactions that conventional linear regression approaches are structurally unable to detect (Gu, Kelly, & Xiu, 2020). Notably, despite the diversity of methods examined, the authors find that nearly all approaches converge on the same set of dominant predictive signals, centered on variations in momentum, liquidity, and volatility, suggesting that the advantage of machine learning lies less in discovering entirely new predictive information than in extracting more value from already-known signals through more flexible functional forms (Gu, Kelly, & Xiu, 2020).

Taken together, these perspectives suggest that machine learning and predictive algorithms constitute a genuinely diverse toolkit rather than a single unified method, encompassing supervised, unsupervised, deep learning, and reinforcement learning approaches, each suited to different types of financial prediction problems, and that their growing adoption in both academic research and industry practice reflects demonstrable improvements in the accuracy and economic value of financial forecasts relative to traditional statistical techniques. This foundation in predictive modeling provides the necessary basis for examining, in the next section, how these techniques are embedded within broader systems of automation for financial data analysis.

2.4 Automation and Financial Data Analysis

The techniques and applications discussed in the preceding sections are ultimately made possible by, and have in turn accelerated, a broader transformation in the scale and character of the data available to financial institutions and researchers. Goldstein, Spatt, and Ye (2021) define the "big data" phenomenon in finance as a combination of three distinguishing features: large size, high dimensionality, and complex structure, and argue that this convergence of characteristics is revolutionizing the finance industry and reshaping the direction of future financial research. Unlike the relatively small, well-structured datasets that characterized traditional empirical finance, the data now available to institutions spans structured transaction records, high-frequency trading data, and unstructured sources such as text and images, a shift that necessitates automated analytical tools capable of processing volumes and forms of information that exceed the practical limits of manual analysis (Goldstein, Spatt, & Ye, 2021).

This transformation has had concrete implications across several core areas of finance. Goldstein, Spatt, and Ye (2021) note that the growing availability of large, high-dimensional, and complex datasets has pushed the frontier of research questions in corporate finance, market microstructure, and asset pricing, enabling analyses that would have been infeasible under earlier data constraints. In a more recent follow-up assessment of the field, Goldstein, Spatt, and Ye (2025) revisit this three-pronged definition and observe that subsequent research has extended it in new directions, including the evaluation of market microstructure reforms, the study of medium-frequency trading, the development of improved methods for imputing missing data, and a deeper focus on the valuation of data itself as a financial asset. This evolution suggests that automation in financial data analysis is not a static capability but a continuously advancing one, as new statistical and computational methods are developed to keep pace with the growing complexity of available data.

At the level of individual institutions, this shift toward automated, large-scale data analysis has translated into concrete operational systems for financial intelligence and data management. A case in point is offered by Liu and Fu (2024), who develop a computer-intelligence-based financial control system for a corporate group, combining a neural network with a particle-swarm optimization algorithm to evaluate financial management effectiveness and to support risk early-warning; although their setting is corporate financial control rather than market analysis in the strict sense, the study illustrates the same broader logic of using automated, algorithm-based tools to give decision-makers a clearer and more current understanding of an organization's financial position. Such systems illustrate how the abstract properties of big data identified by Goldstein, Spatt, and Ye — scale, dimensionality, and structural complexity — are addressed in practice through automated pipelines that continuously ingest, process, and analyze financial data, reducing the reliance on manual, periodic reporting in favor of more immediate, data-driven decision-making.

This evidence points to a single conclusion: automation and financial data analysis are mutually reinforcing developments: the sheer scale and complexity of contemporary financial data make automated analytical methods indispensable, while advances in automation, in turn, make it possible to extract meaningful insight from data that would otherwise remain unusable.

CHAPTER 3

AI-DRIVEN TRADING

3.1 Algorithmic Trading and Automated Trading

Building on the definition of digital trading platforms established in previous chapters, it is useful to examine how these platforms enable a specific form of market participation: algorithmic trading. Algorithmic trading can be understood as the use of computer-based systems to execute trading strategies rapidly and efficiently, often on the basis of real-time market data and predefined instructions (Khan et al., 2024). Although it is frequently associated with recent advances in artificial intelligence, algorithmic trading is not a new phenomenon. Khan et al. (2024) trace its development to the 1970s, when computerized trading systems began to emerge in U.S. markets. This historical continuity suggests that AI-driven trading should be viewed as a later stage in the longer development of computerized and automated execution, rather than as an entirely new phenomenon.

The widespread adoption of algorithmic trading can be explained by several complementary advantages over manual execution. Empirical evidence from the Deutsche Börse suggests that algorithmic traders can reduce monitoring and search frictions and contribute to the efficiency of liquidity provision. They tend to demand liquidity when bid–ask spreads are narrow and supply liquidity when spreads are wider, thereby improving the functioning of the market for liquidity (Hendershott & Riordan, 2013). At the individual level, algorithmic systems are associated with faster execution, the ability to monitor several markets through large-scale data analysis, and a lower likelihood of mistakes than human traders. They may also reduce the influence of emotions during the execution process (Khan et al., 2024). Although banks and hedge funds remain important users, corporations may also employ high-frequency trading systems, indicating that automated trading is not confined exclusively to traditional financial institutions (Khan et al., 2024).

Within the broader field of algorithmic trading, high-frequency trading (HFT) refers to systems capable of executing transactions within milliseconds and exploiting small, short-lived market variations (Khan et al., 2024). Some recent contributions describe AI-driven trading as an increasingly important and growing component of overall market activity, even though the precise

share of AI-driven or HFT volume in total trading is difficult to pin down and varies across markets and measurement approaches (Joby & Vandana, 2024). Common algorithmic strategies discussed in the literature include the automated analysis of news events, the identification of recurring market patterns, and the processing of trading signals to detect potentially profitable opportunities (Khan et al., 2024).

This overview suggests, in turn, that algorithmic and automated trading provide an important operational basis for applying more advanced AI techniques to financial-market analysis and execution. Having established what algorithmic trading is and why it has become increasingly important in contemporary markets, the following section examines how artificial intelligence contributes to the analysis of market data that supports these trading systems.

3.2 The Use of AI in Market Data Analysis

Whereas the preceding section addressed the execution of trades, AI's contribution to financial markets begins earlier in the process, at the stage of data analysis itself. AI-driven data processing and analysis are considered essential to contemporary investment decision-making, since they provide real-time insights and predictions that enhance both market forecasting and portfolio management (Khan et al., 2024). Techniques such as machine learning, natural language processing, and deep learning allow financial analysts to manage volumes of structured and unstructured data — ranging from transaction records to news articles and regulatory filings — that would be impractical to process through manual analysis alone (Khan et al., 2024). In high-frequency trading environments specifically, AI systems are capable of analyzing extensive datasets within milliseconds, allowing market participants to identify and act upon inefficiencies at a speed that traditional, human-driven approaches cannot replicate (Khan et al., 2024).

The practical value of applying machine learning to financial data analysis has also been demonstrated rigorously in the academic asset-pricing literature. As discussed in Chapter 2, Gu, Kelly, and Xiu (2020) show that machine learning methods generate substantial economic gains over traditional linear approaches when applied to the analysis of asset risk premiums, owing largely to their capacity to detect nonlinear relationships within financial data that conventional statistical models are structurally unable to capture. This finding is echoed in the broader applied literature on portfolio management, where machine learning algorithms are used to enable continuous, responsive portfolio optimization by uncovering patterns and relationships in

historical data that human analysts would be likely to overlook, while deep learning techniques are considered particularly valuable in volatile markets for their ability to model complex, nonlinear relationships across multiple assets simultaneously (Khan et al., 2024).

Beyond structured price and volume data, AI-driven analysis increasingly extends to alternative data sources that were traditionally inaccessible to systematic financial analysis. Financial institutions and hedge funds have begun to incorporate unconventional data streams, such as consumer transaction records and satellite imagery, into their analytical processes in order to extract additional market intelligence not reflected in conventional financial statements (Khan et al., 2024). This expansion of the analytical scope of AI illustrates a broader pattern already noted in Chapter 2: as the scale, dimensionality, and structural complexity of available financial data continue to grow, automated analytical tools become correspondingly more central to extracting usable insight from that data (Goldstein, Spatt, & Ye, 2021).

These applications, considered jointly, show that AI's role in market data analysis is not confined to accelerating existing analytical tasks, but extends to enabling entirely new categories of analysis — from nonlinear asset pricing models to the incorporation of unconventional data sources — that would be largely infeasible under traditional, manually driven approaches. With this broader role of AI in market-data analysis established, attention now turns to one of the most prominent and rapidly growing applications of this capability: the analysis of sentiment for the purpose of predicting market trends.

3.3 Sentiment Analysis and Trend Prediction

Among the various applications of AI to financial data analysis discussed in the previous section, sentiment analysis occupies a particularly prominent position, owing to its capacity to extract quantifiable signals from sources of information — news articles, social media posts, and investor surveys — that are inherently unstructured and qualitative in nature. Financial market forecasting is widely regarded as one of the most attractive practical applications of sentiment analysis, precisely because it offers the possibility of translating diffuse, text-based expressions of investor opinion into concrete, actionable predictive signals (Mudinas, Zhang, & Levene, 2019). However, the empirical relationship between sentiment and subsequent price movements is more nuanced than is often assumed: in a rigorous test using Granger-causality analysis, Mudinas, Zhang, and Levene (2019) find that sentiment attitudes, meaning simple positive-versus-negative polarity,

generally do not appear to Granger-cause stock price changes, whereas more granular sentiment emotions, such as joy or sadness, do so only in specific cases rather than as a universal pattern. Despite this nuance, the same study finds that incorporating sentiment emotions as additional features within a broader machine learning-based trend prediction model can improve its predictive accuracy for at least certain stocks (Mudinas et al., 2019), suggesting that sentiment is most useful not as a standalone predictor, but as one input among several within a larger analytical framework.

This more integrative use of sentiment is consistent with the broader literature on AI applications in investment decision-making, which finds that sentiment analysis is strongly linked to market fluctuations and, when properly incorporated into predictive models, allows investors to make more timely, better-informed decisions (Costola, Iacopini, & Santagiustina, 2023, as cited in Khan et al., 2024). Natural language processing techniques are increasingly used to extract sentiment signals from financial news, social media activity, and company filings, automating a task that would otherwise require impractical amounts of manual reading and interpretation, and thereby allowing investors to remain updated on shifts in market mood in close to real time (Khan et al., 2024).

During my internship at Cube Investimenti, I had the opportunity to deepen my understanding of these concepts through direct exposure to the firm's Dual Sentiment Index (DSI), a proprietary quantitative system designed to identify statistically anomalous levels of aggregate investor pessimism (Cube Investimenti, 2026). The DSI illustrates, in a concrete institutional setting, several of the principles discussed in the academic literature above: rather than relying on a single sentiment proxy, the model combines multiple indicators into a composite score, following a contrarian logic whereby episodes of extreme fear are historically associated with a reduced potential for further downside and a higher likelihood of subsequent positive returns (Cube Investimenti, 2026). The system is structured around two complementary components operating on distinct time horizons: a Tactical Index, built on three equally weighted market-based indicators — implied volatility (VIX), the put/call ratio, and the term structure of volatility (VIX/VXV) — designed to capture short, acute episodes of fear over a one-to-three-month horizon; and a Structural Index, which aggregates high-yield credit spreads (HY OAS), the net positioning of commercial operators on S&P 500 futures (COT Net), the average equity exposure of active U.S. managers (NAAIM), and the share of bullish retail investors (AAII Bull%), together with a

conditional VIX component that activates only above a threshold of 30, in order to detect more prolonged deteriorations in sentiment over a six-to-twelve-month horizon (Cube Investimenti, 2026). Each indicator is converted into a rolling z-score computed over 156 weeks (three years), a normalization technique that allows the model to adapt automatically to the prevailing market regime rather than comparing structurally different volatility environments on a fixed historical baseline (Cube Investimenti, 2026).

A particularly instructive methodological feature of the DSI is its signal-activation mechanism: rather than generating a signal as soon as a composite index enters "fear" territory — defined as a composite z-score below -0.75σ — the model waits for an upward inversion of a short-term moving average (a three-week moving average for the Tactical Index and a four-week moving average for the Structural Index) while the index remains within that anomalous zone, a design choice explicitly intended to reduce false positives by filtering out indices that are still deteriorating (Cube Investimenti, 2026). The system's most powerful configuration, referred to internally as a "double signal," occurs when both the Tactical and Structural components generate a signal within a window of eight weeks of one another.

The internal backtest of the DSI, covering the period from January 2006 to April 2026 (1,057 weeks) and including the 2008–2009 financial crisis, the COVID-19 shock, the 2022 bear market, and the 2025 correction, reports a clear hierarchy among the three signal types in terms of subsequent returns, hit rates, and drawdowns. Table 3.1 reproduces the main results of this internal backtest.

Signal type	N	Avg. return +3M	Hit rate +3M	Avg. return +12M	Hit rate +12M	Max drawdown
Double signal (Tactical + Structural)	3	+15.3%	100%	+34.3%	100%	-2.3%
Structural signal (MA4), standalone	6	+2.5%	67%	+18.8%	100%	-7.3%
Tactical signal (MA3), standalone	12	+9.1%	100%	+24.0%	91%	-3.3%

Table 3.1 - Dual Sentiment Index: historical backtest performance by signal type (Cube Investimenti, 2026; backtest period: January 2006-April 2026, 1,057 weeks).

Note: N = number of historical occurrences. Hit rate = share of signals followed by a positive return over the indicated horizon. Max drawdown = maximum drawdown recorded in the 12 months following the signal. Source: Cube Investimenti internal documentation (Bloomberg, CFTC, NAAIM, AAII data).

As Table 3.1 shows, the double signal is by far the rarest configuration, having occurred only three times over the twenty-year sample, but it is associated with the highest average returns at both the three-month and twelve-month horizons and with the smallest maximum drawdown, consistent with the idea that the simultaneous confirmation of a short-term and a longer-term sentiment extreme identifies especially reliable turning points. The standalone Tactical signal displays a 100% historical hit rate at the three-month horizon across all twelve occurrences, while the standalone Structural signal shows a comparatively weaker three-month hit rate of 67%, although its twelve-month hit rate matches that of the other signal types. This pattern is consistent with the academic literature's caution that sentiment indicators tend to perform best when combined into a broader, multi-indicator composite rather than relied upon in isolation (Mudinas et al., 2019): the Structural Index alone appears less informative in the short run, but it materially improves the reliability of the double signal when it coincides with a Tactical trigger.

A further illustration of how the DSI operates in practice is provided by its status on 2 April 2026, the most recent observation available during the internship. Table 3.2 reports the individual indicator readings underlying the two composite indices at that date, while Table 3.3 reports the resulting status of the composite indices themselves.

Indicator	Value	Z-score	Regime
VIX	23.87	-1.46 sigma	Fear
Put/Call Ratio	1.16	-0.18 sigma	Neutral
VIX/VXV	0.960	-1.06 sigma	Fear
HY OAS (high-yield credit spread)	299 bps	+0.43 sigma	Neutral
COT Net (commercial futures positioning)	-42.5k	+0.72 sigma	Neutral
NAAIM (active manager equity exposure)	68.36	-0.62 sigma	Neutral
AAII Bull% (retail bullish sentiment)	33.6%	-0.61 sigma	Neutral

Table 3.2 - Dual Sentiment Index: individual indicator readings as of 2 April 2026 (Cube Investimenti, 2026)

Composite index	Z-score	Fear threshold	Status
Tactical Index	-0.90 sigma	< -0.75 sigma	ACTIVE SIGNAL
Structural Index	+0.08 sigma	< -0.75 sigma	Neutral

Table 3.3 - Dual Sentiment Index: composite index status as of 2 April 2026 (Cube Investimenti, 2026)

At that date, the VIX and the VIX/VXV term-structure ratio were both in "fear" territory, while the put/call ratio remained neutral; on the Structural side, credit spreads, commercial positioning, professional manager exposure, and retail sentiment were all within neutral ranges. As a result, the Tactical Index had crossed below the −0.75σ threshold and generated an active signal, following the upward inversion of its three-week moving average, whereas the Structural Index remained neutral at +0.08σ. This configuration corresponds to a standalone Tactical signal without Structural

confirmation — historically associated, per Table 3.1, with an average three-month return of +9.1% and a 100% hit rate, but without the additional resilience that a confirmed double signal would provide. This example illustrates concretely how the two-tier architecture of the DSI is designed to distinguish between short-lived volatility spikes and more structural deteriorations in sentiment, and how the absence of Structural confirmation is treated by the model as a meaningfully different — and historically less robust — market condition.

These figures should nevertheless be interpreted with the same caution that applies to most backtested trading signals. The historical sample underlying Table 3.1 is limited to eighteen Tactical signals, six Structural signals, and three double signals over twenty years, a sample size that constrains the statistical reliability of the reported hit rates and average returns, particularly for the rarer Structural and double-signal categories. The developers of the DSI themselves note that the model is designed for episodes of acute panic and market corrections rather than for anticipating structural recessions or systemic solvency crises, and that in a prolonged structural bear market the Tactical component could generate multiple signals during intermediate rebounds, some of which might prove premature (Cube Investimenti, 2026). This real-world implementation offers a concrete illustration of two points emphasized in the academic literature: first, that sentiment indicators tend to function best as part of a broader, multi-indicator composite rather than in isolation, consistent with the findings of Mudinas et al. (2019); and second, that the practical value of AI- and rule-based sentiment systems depends heavily on careful signal design aimed at minimizing false positives, a concern that becomes especially important when such systems are deployed to inform real investment decisions rather than purely academic forecasting exercises.

Taken together, these perspectives suggest that sentiment analysis and trend prediction constitute a genuinely valuable, but methodologically delicate, application of AI to financial markets: valuable because they allow previously inaccessible qualitative and behavioral information to be systematically incorporated into investment decisions, but delicate because the empirical link between raw sentiment and subsequent price movements is neither universal nor mechanically reliable, and depends heavily on how sentiment signals are combined, normalized, and filtered before being translated into an operational trading signal. Having examined the specific applications of AI to trading execution, data analysis, and sentiment-based prediction, it is now

useful to step back and assess, more critically, the advantages, limitations, and risks that accompany the use of AI in trading more broadly.

3.4 Advantages, Limits, and Risks of AI in Trading

The applications examined in the preceding sections point to clear advantages of AI-driven trading over traditional, human-mediated approaches. AI-based trading systems are widely reported to reduce emotional biases in decision-making, resulting in investment choices that are more systematic and less prone to the overconfidence, loss aversion, and recency effects that frequently distort human judgment (Khan et al., 2024). By continuously learning from new data and adapting to evolving market conditions, AI systems can also respond to rapid developments in financial markets with a degree of agility that traditional, less adaptive approaches often struggle to match (Khan et al., 2024). These advantages, combined with the capacity to process vast quantities of structured and unstructured data discussed in Section 3.2, help explain the sustained growth in the adoption of AI-driven trading across institutional and, increasingly, retail contexts.

Notwithstanding these advantages, the literature identifies a number of significant limitations. A first concern relates to the quality and representativeness of the data on which AI models are trained: the accuracy of AI predictions is heavily contingent on the completeness and quality of the underlying data, and gaps or biases in historical data can result in flawed models and correspondingly defective investment strategies (Khan et al., 2024). This concern is compounded by the risk of algorithmic bias, whereby AI models may unintentionally reproduce and reinforce distortions present in their historical training data, generating skewed predictions and inequitable investment outcomes (Khan et al., 2024). A further, structurally important limitation is the reliance of most AI models on historical patterns, which does not guarantee reliable forecasts of future market behaviour, particularly during atypical episodes such as financial crises, when historical relationships between variables may break down entirely (Khan et al., 2024).

Beyond these methodological limitations, the growing role of AI in trading raises risks that extend to the stability of financial markets as a whole. Perhaps the most frequently cited illustration of this systemic dimension is the 2010 Flash Crash, an episode in which excessive automation is considered to have contributed to a sudden and severe, if short-lived, market disruption (Kirilenko, Kyle, Samadi, & Tuzun, 2017; Joby & Vandana, 2024). Episodes of this kind illustrate that while individual AI-driven trading systems may each behave in a manner consistent with their

programmed objectives, their simultaneous and largely uncoordinated interaction can generate emergent, system-level instability that no single system was designed to produce — a concern that becomes especially relevant when considering the aggregate impact of AI-driven trading on market efficiency, volatility, and liquidity, discussed in the next chapter. A further, closely related risk concerns the interpretability of AI models themselves: many AI algorithms, particularly those based on deep learning, are frequently criticized as "black boxes," in which the internal decision-making process remains opaque even to the institutions deploying them, raising concerns about accountability, investor trust, and the ability of regulators to adequately oversee automated trading activity (Khan et al., 2024).

On balance, these advantages, limitations, and risks suggest that AI-driven trading should not be understood as a simple, unambiguous improvement over traditional approaches, but rather as a technology whose benefits in speed, scale, and consistency are accompanied by genuine and, in some cases, systemic trade-offs. This assessment provides a natural bridge to Chapter 4, which examines these trade-offs at the level of the financial system as a whole, considering in turn the effects of AI-driven trading on market efficiency and execution speed, volatility and liquidity, investor behaviour, and the regulatory and ethical frameworks that seek to govern its use.

CHAPTER 4

IMPACTS ON FINANCIAL MARKETS AND REGULATORY ISSUES

4.1 Effects on Market Efficiency and Execution Speed

The increasing use of algorithmic and high-frequency trading has transformed the mechanisms through which information is processed, orders are executed, and prices are formed in modern securities markets. High-frequency trading is a particular form of algorithmic trading characterized by infrastructure designed to minimize latency, limited or no human intervention in order initiation, generation, and routing, and high intraday volumes of orders, quotes, and cancellations (European Parliament and Council, 2014, as reported in Clapham et al., 2023). Its principal technological advantage is the ability to process market information and react within extremely short time intervals. This development has made execution speed an important dimension of market participation and has raised the broader question of whether investments in low-latency technology improve the operational efficiency of securities markets or instead intensify a costly technological arms race.

Market quality is commonly assessed through several related dimensions, including liquidity, volatility, and price discovery (Clapham et al., 2023). Price discovery refers to the process through which new information is incorporated into market prices, while execution speed determines how rapidly traders can respond to information and submit orders. High-frequency traders can contribute to price discovery because their technology-based information processing and speed advantages allow them to react quickly to new information, trade against transitory pricing errors, and coordinate prices across different trading venues. The literature reviewed by Clapham et al. (2023) indicates that HFT can improve the incorporation of information into prices and facilitate price coordination across markets by reducing the time and costs required to monitor and respond to information across venues (Clapham et al., 2023).

The relationship between algorithmic trading and market quality is also reflected in the effects of automated liquidity provision. Hendershott, Jones, and Menkveld (2011), as reported by Clapham et al. (2023), find that the introduction of algorithmic trading on the New York Stock Exchange improved market quality by reducing bid–ask spreads. The bid–ask spread represents the

difference between the best available selling price and the best available buying price. Lower spreads allow market participants to trade at more favourable prices, thereby reducing implicit transaction costs and increasing liquidity (Clapham et al., 2023). Hendershott and Riordan (2013) further show that algorithmic traders tend to consume liquidity when spreads are narrow and provide liquidity when spreads are wide. This finding suggests that automated trading can respond dynamically to market conditions rather than simply increasing the number of orders submitted to the market.

The same general mechanism applies to high-frequency trading, although HFT relies additionally on specialized low-latency infrastructure. The literature reviewed by Clapham et al. (2023) generally associates HFT with lower bid–ask spreads, greater order-book depth, and lower short-term price volatility, particularly in equity markets and under ordinary trading conditions. These effects are partly explained by the fact that many HFT firms employ market-making strategies and predominantly submit passive orders, thereby providing liquidity to other market participants. In addition, theoretical work reviewed by Clapham et al. (2023) suggests that speed advantages allow market makers to revise their quotes more efficiently and reduce inventory costs, which may enable them to quote tighter spreads (Clapham et al., 2023). The resulting reduction in implicit transaction costs can therefore generate a positive spillover from HFT firms' investments to participants that do not themselves invest in low-latency technology.

A particularly informative empirical case is provided by the technical interruption that occurred on Deutsche Börse's Xetra trading platform on 2 October 2017. A failure of the high-frequency gateway interrupted the specialized connections that enabled HFT firms to communicate with the exchange at low latency. The event made it possible to observe how an automated securities market reacted when low-latency HFT technology was suddenly unavailable, without relying on a proxy for HFT activity or on a gradual regulatory change (Clapham et al., 2023). The researchers compare DAX30 stocks traded on Xetra, which formed the treatment group, with CAC40 stocks traded on Euronext Paris, which formed the control group. The two groups were selected because of their comparability in terms of market design, trading hours, industry coverage, macroeconomic exposure, and HFT activity (Clapham et al., 2023).

The results indicate that the interruption of HFT substantially reduced several dimensions of market quality while leaving overall trading activity largely unchanged. During the interruption, the average relative spread of DAX30 stocks increased by 17.4%, from 4.90 to 5.75 basis points,

while the spread of CAC40 stocks decreased slightly. The descriptive difference-in-differences calculation therefore indicated that the spread of the affected DAX30 stocks was 21.1% higher than that of the control group (Clapham et al., 2023). Because the quoted spread is an indicator of implicit transaction costs for relatively small orders, this result suggests that trading became more expensive when low-latency HFT technology was unavailable.

The interruption also reduced several dimensions of order-book liquidity. Relative to the control group, L1-Volume decreased by 9.1%, Depth(10) decreased by 24.0%, and order imbalance increased by 32.5% (Clapham et al., 2023). These results indicate that the contribution of HFT to market quality cannot be assessed solely by examining the number of trades or total trading volume. The quantity, distribution, and balance of available liquidity in the order book are also important. The results therefore suggest that the absence of low-latency HFT can make markets less capable of absorbing orders, even when the overall amount of trading does not decline significantly.

The regression results reinforce this conclusion. The estimated coefficient indicates that the interruption led to an increase in the relative spread of 1.08 basis points, corresponding to a rise of approximately 22% relative to the non-event level of 4.90 basis points. Depth(10), in turn, declined by approximately €203,947, or 25%, relative to the baseline value of €813,850, while order imbalance increased by more than 31% (Clapham et al., 2023). The effect on L1-Volume was weaker and statistically less robust than the effects on the other liquidity measures. The authors also report that the interruption did not significantly change the number of trades, trading volume, quotes, or order submissions. This suggests that the deterioration in market quality was not merely the mechanical result of a reduction in trading activity. Instead, the evidence indicates that low-latency HFT infrastructure affected the conditions under which trading took place, particularly by supporting tighter spreads and deeper order books (Clapham et al., 2023).

HFT also affects price volatility, although the evidence is less uniform than for liquidity. Clapham et al. (2023) distinguish between trade-price volatility and midpoint volatility because rapid order submissions and cancellations can affect quoted prices differently from executed transaction prices. When HFT was interrupted, trade-price volatility increased significantly. After accounting for changes in trading activity, the net effect was approximately 3% when the number of trades was used as the reference and approximately 7% when the number of quotes was used as the

reference. By contrast, the effect on midpoint volatility was not statistically significant and was not robust across specifications (Clapham et al., 2023).

This difference can be explained by the dual role of HFT. Frequent order submissions and cancellations may increase the movement of quotes and the midpoint, while market-making activity can absorb short-term price changes by supplying liquidity. When HFT is interrupted, these opposing effects may offset one another in the case of midpoint volatility, whereas the reduction in liquidity provision can produce a clearer effect on executed trade prices. The evidence therefore supports a qualified interpretation of the relationship between execution speed and market quality.

On the one hand, HFT can provide liquidity, reduce implicit transaction costs, process information rapidly, and support price discovery. The Xetra interruption demonstrates that the disappearance of low-latency HFT can make markets less liquid and more costly even when trading activity itself does not significantly decline (Clapham et al., 2023). On the other hand, speed does not automatically guarantee long-term informational efficiency. Clapham et al. (2023) note that HFT may reduce the time available to informed traders to recover the costs of producing information through research. Prices may therefore become more efficient in the short run because private information is incorporated more rapidly, while the sharing of informational rents may weaken institutional investors' incentives to engage in costly research and information production in the long run (Clapham et al., 2023).

A complementary market-microstructure perspective is offered by Tse, whose study examines the interaction between market makers, investors, liquidity, transaction costs, and price impact using high-frequency data. Tse emphasizes that market makers may become reluctant to trade when adverse-selection risk is too high, potentially weakening market clearing and contributing to flash crashes or other sudden market disruptions (Tse, 2024). This perspective suggests that efficient execution depends not only on how quickly orders can be submitted, but also on whether liquidity providers remain willing to quote and whether the market can continue to match buying and selling interests under stressed conditions.

The evidence from Xetra also shows that the effect of HFT should not be reduced to a simple opposition between human and automated trading. HFT firms were not completely disconnected from Xetra. Rather, they were unable to use their specialized low-latency connections and could continue trading through slower connections (Clapham et al., 2023). The natural experiment

therefore provides evidence on the value of speed and infrastructure more specifically than on the effect of the complete absence of algorithmic traders. This distinction is important for the analysis of AI-driven trading because it shows that the technological architecture supporting automated decisions is itself a component of market quality. Algorithms may generate trading decisions, but their contribution to execution and price discovery also depends on connectivity, data processing, order-book access, and the ability to interact with other market participants at low latency.

Nevertheless, the positive effects of HFT must be interpreted within the limitations of the evidence. Clapham et al. (2023) analyze highly liquid blue-chip stocks and a one-hour interruption at a particular European trading venue. Their findings may therefore not be fully generalizable to other stock segments, less liquid assets, or different market structures. In addition, although the DAX30 and CAC40 groups display common trends before the event and the matched-sample analysis produces comparable results, the research design cannot entirely exclude every possible difference between German and French securities markets (Clapham et al., 2023).

Taking stock of this evidence, algorithmic and high-frequency trading appear to improve certain dimensions of market quality, particularly by reducing execution frictions, supporting liquidity provision, and accelerating the incorporation of information into prices. The strongest evidence concerns liquidity and transaction costs: when the low-latency infrastructure was interrupted, spreads widened, order-book depth declined, and order imbalance increased, while trading activity did not change significantly (Clapham et al., 2023). The effect on volatility is more ambivalent. HFT can absorb short-term price movements through market-making activity, but rapid and interconnected trading may also amplify extreme price movements when market conditions deteriorate. Execution speed can therefore be interpreted as a source of both efficiency and exposure: it can make markets faster and more liquid under ordinary conditions, while increasing the importance of resilient technological infrastructure, effective market surveillance, and safeguards capable of limiting the consequences of technical failures and sudden withdrawals of liquidity. The following section examines this dual relationship in greater detail by analyzing how high-frequency and AI-driven trading affect market volatility and liquidity, two principal dimensions of market quality.

4.2 Impact on Volatility and Liquidity

The effects of algorithmic and high-frequency trading on volatility and liquidity are among the most contested issues in the analysis of modern financial markets. Liquidity is generally understood as the ability to trade large volumes quickly and at low cost, while volatility reflects the intensity of fluctuations in asset prices over a given period (Ben Ammar, Hellara, & Ghadhab, 2020). These two dimensions of market quality are closely connected: wider spreads and lower order-book depth increase the cost of trading, while rapid price changes may increase the compensation required by liquidity providers and induce them to withdraw from the market. The available evidence does not support a uniform conclusion that HFT is either stabilizing or destabilizing. Rather, its effects depend significantly on market conditions, the trading strategies employed, and the interaction between liquidity provision and liquidity demand.

Ben Ammar, Hellara, and Ghadhab (2020) define liquidity as the ability to trade large volumes quickly and at low cost and emphasize that it is a multidimensional concept that cannot be measured by a single indicator. Their analysis, based on tick-by-tick data for CAC40 stocks traded on Euronext over the first half of 2016, considers quoted and effective spreads, realized spreads, adverse-selection costs, and quoted depth. The quoted spread measures the difference between the best ask and the best bid relative to their midpoint, whereas the effective spread measures the distance between the transaction price and the quote midpoint. Quoted depth captures the quantity of liquidity available at the best bid and ask prices, complementing the price dimension of the spread with the quantity dimension of the order book (Ben Ammar et al., 2020). A narrow spread indicates lower implicit transaction costs for immediate execution, whereas greater quoted depth indicates that larger orders can be executed without causing a substantial price impact. The effective spread can further be decomposed into a realized component, representing revenues to liquidity providers, and an adverse-selection component, representing losses that arise when prices move against the market maker after a transaction. A market maker facing informed or better-informed traders may respond by widening spreads, cancelling stale quotes, and replacing them at different price levels (Ben Ammar et al., 2020).

Ben Ammar, Hellara, and Ghadhab (2020) use the quote-to-trade ratio and a normalized electronic-message ratio as proxies for HFT intensity, since HFT activity cannot be directly observed in most public data feeds. The quote-to-trade ratio captures the high message traffic and relatively low number of executions characteristic of HFT, while the message-to-volume measure reflects

changes in electronic trading activity after normalizing message traffic by trading volume. Their descriptive results show that quote updates occur far more frequently than trades and that the average rate of order cancellations is approximately 70%, illustrating the intensity of order management associated with automated trading. The same study reports a negative relationship between HFT intensity and several spread measures and a positive relationship between HFT intensity and quoted depth: the Spearman correlations between the HFT proxies and quoted and effective spreads are negative, while the correlation with quoted depth is strongly positive. The econometric analysis, based on a GMM estimator designed to correct for the endogeneity between HFT activity and liquidity, confirms this pattern: greater HFT intensity is associated with lower spreads and higher depth after controlling for stock-specific effects, time-of-day variation, turnover, volatility, and the inverse of stock price. The authors attribute the positive effect on liquidity mainly to lower adverse-selection costs arising from asymmetric information among market participants (Ben Ammar et al., 2020).

The intraday pattern of liquidity provides further context for these findings. At the beginning of the trading session, spreads are relatively wide and quoted depth is comparatively low; as the session progresses, spread measures decline and depth increases. Table 4.1 summarizes the main intraday liquidity measures reported by Ben Ammar et al. (2020) for the opening and closing half-hour of the trading session.

Measure	**First half-hour (09:00-09:30)**	**Last half-hour (17:00-17:30)**
Quoted spread (basis points)	6.83	3.91
Effective spread (basis points)	8.94	4.55
Quoted depth (EUR)	40,274	60,106

Table 4.1 – Intraday liquidity measures for CAC40 stocks, opening vs. closing half-hour (Ben Ammar, Hellara, & Ghadhab, 2020)

This pattern is consistent with a relatively illiquid market opening, followed by an improvement in liquidity as private information is progressively incorporated into prices and market makers adjust their positions. It also suggests that HFTs manage both the price and the quantity dimensions of liquidity by adjusting spreads and quoted depth simultaneously. This finding complements the

Xetra natural experiment discussed in Section 4.1 (Clapham, Haferkorn, & Zimmermann, 2023), where the sudden unavailability of low-latency HFT technology produced the opposite pattern on the same indicators. Taken together, the intraday evidence from Euronext and the natural experiment on Xetra point in the same direction, namely that HFT activity is associated with narrower spreads and deeper order books under ordinary trading conditions.

The relationship between HFT and volatility is more complex than the relationship between HFT and liquidity, and it depends heavily on prevailing market conditions. Under stable conditions, market-making HFTs may reduce short-term volatility by absorbing order flow and mitigating the price impact of trades in the same direction. Ben Ammar and Hellara (2022) examine this issue using high-frequency data on CAC40 stocks and distinguish between stable market conditions, observed over the period from February to May 2016, and an intraday crash associated with the Brexit referendum announcement of 24 June 2016. Their results show that greater HFT intensity, measured by the quote-to-trade ratio, is associated with lower realized volatility and a narrower intraday price range during stable periods, and that this relationship is considerably stronger for large-capitalization stocks than for small-capitalization stocks. Table 4.2 summarizes the main coefficients reported by the authors for the stable and crash periods.

Indicator	**Stable conditions (Feb.-May 2016)**	**Brexit intraday crash (24 June 2016)**
HFT intensity vs. realized volatility	Negative and significant	Positive and significant
HFT intensity vs. intraday price range	Negative and significant	Positive and significant
Trading volume / HFT intensity relative to normal days	Baseline	Volume approx. 4x higher; HFT intensity approx. 1/3 of usual level; volatility approx. 3x higher

Table 4.2 – HFT intensity (quote-to-trade ratio) and volatility, stable vs. crash periods, CAC40 large-cap stocks (Ben Ammar & Hellara, 2022)

Note: A Wald test rejects the hypothesis that the effect of HFT intensity on volatility is equal across the two regimes.

On 24 June 2016, trading volume was approximately four times higher than during the preceding four trading days, HFT intensity fell to roughly a third of its usual level, and volatility rose to approximately three times its ordinary level. At the opening of the session, uncertainty led HFTs to reduce their presence in the limit order book, submit fewer passive orders, and cancel orders more frequently; until approximately 11:00, HFTs submitted more aggressive orders and consumed more liquidity than they provided. Ben Ammar and Hellara (2022) therefore caution against using trading volume alone as an indicator of liquidity, since a high number of transactions may coexist with a deterioration in the availability of passive liquidity. The interaction between HFT intensity and the crash period is positive and statistically significant for both realized volatility and price range, in contrast to the negative relationship observed during stable conditions, and a Wald test formally rejects the hypothesis that the effect of HFT intensity is the same during normal and stressed conditions.

The authors interpret this result as evidence that rapid interactions among HFT algorithms, together with simultaneous withdrawals from the limit order book, can generate intense price movements. HFTs react quickly to adverse conditions by cancelling orders before execution, which protects them from adverse selection but reduces the liquidity available to other market participants. This conditional pattern is consistent with a broader body of research: some studies find that HFT improves liquidity, reduces volatility, and contributes to more efficient price discovery, while others find that HFT can limit efficiency, increase volatility, and generate negative externalities. The difference between these findings may reflect the fact that HFT strategies do not operate identically across market regimes: market-making strategies tend to stabilize prices when liquidity is available and order flow is balanced, whereas speculative or aggressive strategies may amplify price movements during periods of uncertainty.

A related issue concerns the fragility of HFT-provided liquidity. In ordinary conditions, HFT can supply liquidity through passive limit orders and narrow bid–ask spreads. During stressed periods, however, the same firms may cancel their orders or avoid submitting new passive quotes because the risk of adverse selection and inventory losses increases, creating a liquidity gap precisely when liquidity is most needed. A non-peer-reviewed review of the topic (Lin, Rakity, Stan, & Vijayakumar, 2024) describes this problem in relation to market stability, arguing that HFT can improve liquidity and price efficiency in normal conditions but that liquidity may disappear during periods of stress, when rapid order cancellations and the withdrawal of HFT firms can generate

short-term volatility spikes; this discussion, while illustrative of widely held concerns in the practitioner and policy literature, should be read as a secondary, non-academic source rather than as independent empirical evidence.

The evidence from Abreu (2024) adds an important distributional dimension to the analysis of HFT and liquidity. Using proprietary NASDAQ data that identify the trading activity of 26 HFT firms across 120 randomly selected stocks, drawn from selected weeks between 2008 and early 2010 — including the first full week of each quarter in 2008 and 2009, the crisis week of September 2008, and a week in February 2010 — Abreu finds that HFT improves overall market liquidity, but that the improvement occurs unevenly across market participants. When HFT firms supply liquidity to non-HFT traders, effective spreads are wider than in trades between two non-HFT participants, while, among trades involving a non-HFT counterparty, adverse-selection costs are highest when HFT firms demand liquidity from non-HFT traders. The paper therefore suggests that aggregate improvements in liquidity do not necessarily imply that every category of trader benefits equally: HFT firms can exploit their speed advantage to respond to new information and to adversely select slower traders, imposing higher costs on non-HFT participants even while overall market liquidity improves (Abreu, 2024). Because the sample predates the more recent expansion of HFT activity in European and other markets, these findings should be interpreted as evidence from a specific period of the U.S. equity market rather than as a universally applicable estimate of the distributional effects of HFT.

This distinction is particularly important for retail and institutional investors that do not possess comparable low-latency infrastructure. A market may display narrower average spreads and greater quoted depth while some slower participants face higher adverse-selection costs when trading against HFT firms; Abreu's results indicate that the aggregate market-quality effect and the distribution of benefits across market participants should therefore be analyzed separately. HFT may create positive externalities through liquidity provision, but it may also generate private advantages through faster reactions to information and order flow, creating a persistent asymmetry between participants with different technological capacities.

As noted in Section 4.1, Tse's (2024) market-microstructure perspective suggests that market makers become reluctant to trade when adverse-selection risk is high, weakening the process of market clearing. This perspective is consistent with the empirical results of Ben Ammar and

Hellara (2022): during stressed conditions, the withdrawal of passive HFT orders and the increase in aggressive trading reduce the stabilizing role that HFT plays during normal periods.

The Flash Crash of 6 May 2010 provides the most prominent illustration of this possibility: major U.S. equity indices experienced a sharp decline and partial recovery within minutes, demonstrating that a market can remain highly liquid during ordinary periods but become fragile when automated strategies interact under conditions of extreme order imbalance and rapidly falling prices. Clapham et al. (2023) note that HFTs may initially trade in the opposite direction of sudden price movements, but that they can contribute to the acceleration of extreme events when price movements persist or affect several securities simultaneously. Similar concerns appear in the literature on so-called mini flash crashes, where algorithmic feedback, rapid cancellations, and the withdrawal of liquidity providers can amplify short-term price movements (Clapham et al., 2023).

The evidence therefore supports a regime-dependent interpretation of HFT. During stable conditions, HFT can reduce spreads, increase order-book depth, reduce short-term volatility, and improve the ability of investors to transact quickly and at lower cost. During stressed conditions, however, HFT can contribute to higher volatility when algorithms interact rapidly, cancel orders simultaneously, and consume more liquidity than they provide. The same speed that allows HFT firms to revise quotes efficiently can also allow them to withdraw from the market before slower traders can respond. Consequently, HFT should not be described as intrinsically stabilizing or destabilizing; its effects depend on whether the market is functioning under normal conditions or experiencing a sudden imbalance, information shock, or liquidity shortage.

This conditional interpretation also has implications for the measurement of market quality. Trading volume alone is insufficient to determine whether a market is liquid, since a high number of trades may occur while passive liquidity is being withdrawn. A more complete analysis should include spreads, order-book depth, order imbalance, adverse-selection costs, realized volatility, and the persistence of liquidity through stressed periods. The three empirical studies considered in this section illustrate precisely this broader approach: Ben Ammar et al. (2020) combine spread and depth measures with adverse-selection components on Euronext; Clapham et al. (2023) examine spreads, order-book depth, imbalance, and two volatility measures around a technical interruption on Xetra; and Ben Ammar and Hellara (2022) distinguish stable conditions from an intraday crash on the same French market. Together, these studies show why liquidity and

volatility should be analyzed jointly rather than as isolated indicators, and why findings obtained on one market, sample period, or trading venue should not automatically be extrapolated to others. The evidence also suggests that AI-driven trading may intensify both sides of this relationship. More advanced algorithms can process larger quantities of data and adjust positions more rapidly, potentially improving liquidity provision and reducing volatility under ordinary conditions. At the same time, if different systems rely on similar data, signals, or risk-management rules, their decisions may become correlated, increasing the probability of simultaneous order cancellations, aggressive selling, and liquidity withdrawal. This risk is particularly relevant when models are trained on historical data that do not adequately represent exceptional market conditions: AI may improve the speed and sophistication of trading without eliminating the underlying problem of model uncertainty and market fragility.

In conclusion, the impact of high-frequency and algorithmic trading on volatility and liquidity is fundamentally conditional. The empirical evidence shows substantial benefits in normal market conditions — lower spreads, greater depth, lower short-term volatility, and reduced transaction costs — but these benefits can weaken or reverse during periods of stress. As shown above, the Brexit-related crash (Ben Ammar & Hellara, 2022) and the Xetra interruption (Clapham et al., 2023) illustrate the same duality from opposite directions: HFT intensity stabilizes prices in calm markets but can amplify volatility under stress, while the sudden absence of low-latency HFT technology widens spreads and reduces depth. Overall, these findings indicate that HFT is an important source of market liquidity, but also that its liquidity provision can be fragile and dependent on market conditions. This dual character creates a central regulatory challenge: preserving the efficiency and cost-reducing benefits of automated trading while ensuring that market infrastructure, liquidity providers, and algorithmic systems remain resilient during episodes of extreme volatility. The discussion so far has examined the effects of AI-driven and high-frequency trading on volatility and liquidity; the next section considers how these technologies influence investor behaviour and the strategic interaction between technologically advanced traders, institutional investors, and retail market participants.

4.3 Changes in Investor Behaviour

The development of algorithmic and high-frequency trading has altered not only the technological organization of financial markets, but also the behaviour of the participants operating within them. Investors now trade in an environment characterized by continuous electronic information, very short execution times, automated order submission, and increasing differences in technological capacity. These changes affect how investors obtain information, formulate trading decisions, manage orders, and interact with liquidity providers. Investor behaviour should therefore be analysed not only at the level of individual preferences, but also in relation to the market infrastructure and strategic environment created by automated trading.

One of the most important changes concerns the growing importance of speed in the trading process. High-frequency traders can use sophisticated systems to identify news, process trading signals, and execute orders within milliseconds. Abreu (2024) describes news and directional trading as one of the principal HFT strategies documented in the literature: algorithms can search financial news for relevant information, identify the securities affected, and submit orders before slower market participants have fully interpreted the information. Other HFT firms analyse order-flow patterns and attempt to anticipate the trades of larger investors. When an institutional investor gradually builds a position, algorithms may identify the sequence of orders and trade ahead of the expected price movement (Abreu, 2024). This possibility has encouraged institutional investors to divide large orders into smaller transactions, use hidden orders, or trade through dark pools in order to reduce the information revealed to high-frequency strategies. The increasing importance of speed has consequently changed the relationship between informed, institutional, retail, and technologically advanced traders.

In traditional markets, investors could rely more heavily on research, fundamental analysis, and direct observation of market conditions. In contemporary electronic markets, however, information may be incorporated into prices before a slower investor can act upon it. Clapham, Haferkorn, and Zimmermann (2023) note that HFT can accelerate the incorporation of private information into prices, but may also reduce the time available to informed traders to recover the costs of producing that information. This creates a potential long-term trade-off: prices may become more efficient in the short run, while the incentives to undertake costly research may weaken over time.

The speed advantage also affects the behaviour of liquidity providers. Abreu (2024) distinguishes between HFT firms that demand liquidity and those that supply it. HFT market makers post limit orders on both sides of the electronic order book, allowing other traders to execute transactions without waiting for a new counterparty. However, HFT liquidity providers can revise or cancel their orders rapidly when new information arrives. This ability reduces the risk of being adversely selected by informed traders, but it also makes limit orders submitted by slower participants more exposed to stale-price risk. In this context, non-HFT investors may face higher adverse-selection costs because they cannot cancel or modify their orders as quickly as HFT firms (Abreu, 2024).

As discussed in Section 4.2, Abreu's (2024) NASDAQ evidence shows that HFT firms exploit their speed and information-processing advantage at the expense of non-HFT participants, an effect reported to be especially pronounced for smaller-capitalization stocks.

The resulting environment may influence how institutional investors execute large orders. Abreu (2024) reports that institutional traders attempt to avoid revealing their intentions by splitting trades into smaller units, using hidden orders, or accessing dark pools. These practices represent behavioural adaptations to the strategic presence of HFT. Instead of simply submitting a large order to the visible limit order book, an institutional investor may choose to conceal part of the order or distribute execution across time and venues. Such strategies can reduce the information available to high-frequency traders, but they can also increase the complexity and fragmentation of the trading process.

Algorithmic trading also affects retail investors, although the nature of this effect differs from the institutional response. Retail traders generally lack the low-latency infrastructure, direct market connectivity, and analytical capacity available to HFT firms. They may therefore benefit from narrower aggregate spreads and greater market liquidity, while remaining exposed to adverse selection when trading against faster participants. Abreu (2024) shows that improvements in overall liquidity may occur at the expense of non-HFT traders, indicating that aggregate market-quality measures do not necessarily describe the experience of every investor. A market can become more liquid on average while slower traders face less favourable execution conditions in specific transactions.

The growth of electronic and automated trading may also alter the way investors interpret market signals. Algorithmic systems can react to short-term order-flow changes, price discrepancies, and news signals that may have little relationship with a company's long-term fundamentals. Investors

who observe rapid price movements may respond by placing additional market orders, cancelling limit orders, or following the direction of the market. These reactions can reinforce short-term price movements, especially when multiple participants respond to the same information or to the same algorithmically generated signal. A practitioner review of the topic links this type of interaction to the possibility of feedback loops, flash crashes, and sudden changes in liquidity, particularly when automated systems operate in highly interconnected markets (Lin, Rakity, Stan, & Vijayakumar, 2024).

The effects of automation on investor behaviour are particularly evident during periods of market stress, confirming the state-dependent pattern already discussed for the Brexit-related crash in Section 4.2 (Ben Ammar & Hellara, 2022): strategies that provide liquidity during normal conditions may instead withdraw liquidity or increase market pressure during periods of exceptional uncertainty. The withdrawal of liquidity can also influence the behaviour of slower investors. When visible orders disappear and spreads widen, investors may delay execution, submit more aggressive orders, or exit positions at less favourable prices. Automated stop-loss orders may further amplify this process by converting price declines into additional sell orders. Consistent with the empirical evidence on the Flash Crash documented by Kirilenko, Kyle, Samadi, and Tuzun (2017), a student-authored commentary on algorithmic trading (Park, 2025) describes how the withdrawal of HFT liquidity during the 2010 Flash Crash is considered to have contributed to downward momentum, while automated stop-loss orders and further algorithmic selling intensified short-term price movements.

These mechanisms show that investor behaviour is not independent of market design: the structure and speed of the trading system can shape the reactions that investors make to changing prices. The development of AI-driven trading adds another layer to this process. Machine-learning systems can analyse large quantities of market data, identify patterns that are difficult for human investors to observe, and adjust trading parameters as market conditions change. This may improve the speed and consistency of investment decisions, but it may also increase the correlation between trading strategies if different institutions use similar data, models, or risk-management procedures. When several systems generate similar signals, their actions may become concentrated on the same side of the market. Such synchronized behaviour can amplify price movements and reduce the diversity of views that normally contributes to continuous market functioning.

At the same time, automated trading does not eliminate the influence of human behaviour. Human traders and institutions remain responsible for choosing model structures, selecting data, defining risk limits, and determining when algorithms should be activated or suspended. Human decisions may therefore influence the market indirectly through the design and governance of automated systems. Moreover, when algorithms perform poorly or behave unexpectedly, traders may intervene by cancelling orders, changing parameters, or withdrawing from the market. Investor behaviour in an AI-driven environment should therefore be understood as a combination of direct human decisions and decisions delegated to computational systems.

The distribution of information also affects the relationship between retail and institutional investors. HFT firms may process news and order-flow information in milliseconds, while retail investors often rely on publicly available information and slower execution channels. This difference does not necessarily imply that retail investors are always disadvantaged, since they may benefit from the liquidity supplied by HFT firms. However, it does imply that they face a different strategic environment. Retail traders may submit market orders at prices that have already moved, respond to short-term volatility, or place limit orders that become stale before they can be updated. The result is that differences in speed and information processing can influence not only execution costs but also the timing and direction of investment decisions.

The consequences for investor behaviour are therefore ambivalent. Automated trading can increase market access, reduce execution frictions, and allow investors to trade quickly across multiple venues. It can also improve the ability of institutional traders to manage large orders and can provide continuous liquidity during normal periods. At the same time, the same systems may encourage short-termism, increase the importance of technological speed, and expose slower participants to adverse selection. During periods of stress, algorithmic withdrawals and correlated trading responses may further limit the opportunities available to investors who rely on traditional execution methods. Overall, AI-driven and high-frequency trading have not removed investor behaviour from financial markets; they have changed the conditions under which that behaviour occurs. Institutional investors adapt by fragmenting orders and concealing trading intentions, liquidity providers respond to adverse-selection risk by revising or cancelling quotes, and retail investors operate in an environment where speed, automation, and algorithmic feedback shape price movements. These changes reinforce the need to analyse investors not as isolated decisionmakers, but as participants embedded in a technological and strategic market structure.

Having examined how automated trading modifies investor behaviour and the distribution of advantages among market participants, the following section considers the ethical and regulatory issues raised by these developments.

4.4 Ethical and Regulatory Aspects

The expansion of algorithmic and high-frequency trading has created important ethical and regulatory questions concerning market fairness, transparency, manipulation, systemic risk, and the resilience of financial infrastructure. HFT can improve liquidity, accelerate price discovery, and reduce arbitrage opportunities, but its benefits are not distributed equally among participants. The technological advantages associated with low latency, co-location, and sophisticated data processing may enable some firms to react to information before slower investors, creating a persistent asymmetry in execution and information processing. The central regulatory challenge is therefore not simply whether HFT should be allowed, but how its benefits can be preserved while limiting the risks associated with speed, automation, and concentrated technological capacity.

The ethical dimension of HFT is closely related to the unequal access to speed and market information. High-frequency traders invest heavily in computing infrastructure, specialized connections, and co-location services to reduce the time required to receive information and submit orders. Pagnotta and Philippon (2018) describe this process as competition on speed, in which trading venues differentiate themselves by offering faster execution and charge higher fees to participants that value immediate transactions. Although such competition can improve market integration and price discovery, it may also widen the gap between firms with advanced technological resources and slower institutional or retail investors. This raises questions about whether a market can be considered fair when participants formally operate under the same rules but possess substantially different capacities to process and act upon information.

The problem is particularly relevant when HFT strategies anticipate the order flow of other investors, as in the front-running dynamic already discussed previously (Abreu, 2024). These strategies are not necessarily illegal, but they illustrate the ethical tension between legitimate technological advantages and the possibility that slower traders bear higher costs as a result of adverse selection: HFT improves overall market liquidity while imposing higher effective spreads and adverse-selection costs on non-HFT participants, indicating that improvements in aggregate market quality do not necessarily imply equal treatment of all investors.

A second ethical concern relates to the use of order submissions and cancellations to create misleading impressions of market demand or supply. A practitioner review of the topic discusses practices such as flickering quotes, quote stuffing, layering, and spoofing, in which algorithms place and rapidly cancel orders in ways that may influence the perceived depth or direction of the market (Lin, Rakity, Stan, & Vijayakumar, 2024). A high order-to-trade ratio is not automatically evidence of manipulation, since HFT market-making strategies legitimately require frequent quote updates and cancellations. Nevertheless, the distinction between liquidity provision and manipulative behaviour can be difficult to establish when orders are submitted without a genuine intention to execute. This difficulty is intensified when the decision process is performed by an automated system and regulators must determine whether a particular sequence of orders reflects a legitimate strategy or an attempt to mislead other participants.

The 2010 Flash Crash illustrates the potential systemic consequences of these issues. An automated sell programme involving 75,000 E-mini S&P 500 futures contracts is generally considered to have initiated intense selling pressure (Kirilenko, Kyle, Samadi, & Tuzun, 2017; Zhu, 2025); HFT firms initially absorbed part of the order flow, but subsequently liquidated their inventories, contributing to a downward price spiral and a severe liquidity imbalance. The event demonstrates that HFT was not necessarily the original cause of the disturbance, but that interactions among automated strategies could amplify the shock once it had begun. This episode is widely interpreted in the literature as an example of the recurring vulnerabilities of algorithmic finance, in which feedback loops, stop-loss orders, forced sales, and the so-called "hot-potato" effect created by rapid trading among HFT firms (Zhu, 2025) can transform a localized shock into a systemic event. These mechanisms show why market stability cannot be evaluated only by examining the behaviour of individual algorithms: systemic risk may emerge from their interaction.

The Flash Crash also revealed the importance of liquidity obligations. Traditional market makers were generally expected to continue providing quotes and managing inventories during periods of stress, whereas HFT firms may withdraw when adverse-selection and inventory risks become too high; this asymmetry is thought to have contributed to the withdrawal of liquidity during the Flash Crash and to have intensified downward momentum (Kirilenko, Kyle, Samadi, & Tuzun, 2017). These concerns about the regulatory implications of HFT-related liquidity and volatility risk are echoed more broadly in the recent literature (Zhu, 2025). Regulation must therefore consider not

only whether HFT improves liquidity during ordinary conditions, but also whether sufficient safeguards exist to prevent a sudden withdrawal of liquidity when it is most needed.

A further challenge concerns the regulation of fragmented markets. Competition between venues encourages exchanges to invest in lower latency and attract speed-sensitive participants, but it can also disperse trading activity across multiple platforms. Pagnotta and Philippon (2018) show that fast venues can improve price integration by allowing HFTs to identify and eliminate price differences across markets. At the same time, fragmented trading may reduce the amount of liquidity available at any single venue and make supervision more difficult. Regulators must therefore balance the benefits of competition and price integration against the risk that fragmented infrastructures create information gaps, inconsistent safeguards, or channels through which instability can spread rapidly across markets.

The regulatory response has developed through a combination of authorization requirements, algorithmic controls, transparency obligations, market safeguards, and enforcement against manipulation.

Jurisdiction / Framework	Main instrument	Primary objective
European Union - MiFID II	Authorization requirement, record-keeping, algorithmic testing obligations	Ensure firms engaging in algorithmic trading do not contribute to disorderly market conditions
Germany - High-Frequency Trading Act	Licensing requirements, algorithmic circuit breakers, co-location restrictions	Limit risks specifically associated with low-latency trading infrastructure
United States - SEC Rule 15c3-5	Broker risk-management controls for market access	Prevent erroneous or manipulative orders from reaching the market
Exchange level - Circuit breakers (ex-ante / ex-post)	Trading halts triggered by price limits	Interrupt feedback loops during sudden price movements
Exchange level - Minimum resting times	Minimum period before an order can be cancelled or modified	Reduce misleading, short-lived liquidity signals

Table 4.3 – Selected regulatory instruments addressing algorithmic and high-frequency trading

Table 4.3 summarizes the main regulatory instruments that have been introduced specifically in response to the growth of algorithmic and high-frequency trading.

Under the European MiFID II framework, firms engaging in algorithmic trading must be authorized, preserve digital records, and ensure that their algorithms do not contribute to disorderly market conditions. Germany's High-Frequency Trading Act introduced additional requirements, including licensing provisions, algorithmic circuit breakers, and restrictions affecting co-location services. In the United States, Rule 15c3-5 requires brokers to maintain risk-management controls for market access, while enforcement actions may target practices such as spoofing. These measures reflect a shift from regulating only financial institutions toward regulating the technological systems through which orders are generated, routed, and executed.

Circuit breakers represent one of the main tools available to limit the effects of sudden price movements. Ex-ante circuit breakers establish price limits that, when triggered, temporarily halt trading across a market; ex-post mechanisms instead interrupt trading in individual securities after a specified price movement. Simulation-based research on agent-based markets suggests that ex-ante mechanisms can prevent the continuation of simulated flash crashes, whereas ex-post mechanisms may lengthen their duration (Jacob Leal & Napoletano, 2016). Although the effectiveness of circuit breakers depends on their design, they can provide time for information to be assessed, reduce the speed of feedback loops, and prevent a temporary imbalance from developing into a broader market disruption.

Minimum resting times provide another possible regulatory mechanism. Under such a rule, an order cannot be cancelled or modified until a specified minimum period has elapsed. The aim is to reduce misleading short-lived signals and limit the ability of algorithms to create the appearance of liquidity that disappears before other participants can respond. Minimum resting times may also reduce the technological gap between HFT firms and slower investors by preventing the immediate cancellation of orders. However, the measure may also reduce legitimate market-making flexibility and increase execution costs if liquidity providers are forced to maintain positions in rapidly changing conditions. Its design would therefore need to distinguish between manipulative order activity and legitimate risk management.

Regulators may also consider speed limits, batch auctions, or transaction taxes. Such measures could reduce the competitive advantage created by extremely low latency, limit excessive order cancellation, and address the technological arms race among venues and HFT firms. Pagnotta and Philippon (2018) similarly argue that the welfare benefits of continuously reducing latency may be limited relative to the high costs of the technological investments required. A frequent batch auction, for example, would replace continuous priority based exclusively on speed with periodic order matching, reducing the value of winning by a few microseconds. A transaction tax could discourage extremely short-term strategies, although it might also reduce legitimate liquidity provision and increase the cost of trading for other participants.

The design of regulation must also address the difficulty of identifying and attributing algorithmic misconduct. HFT is not defined by a single strategy, and the same technical behaviour — such as repeated order submission and cancellation — may occur in both legitimate market-making and manipulative activity. The requirement to prove intent can create particular difficulties when decisions are generated by automated systems: an algorithm does not possess intention in the same way as a human trader, but it operates according to parameters selected by a firm or programmer. Effective supervision may therefore require a combination of behavioural thresholds, audit trails, algorithmic testing, and institutional accountability, rather than relying exclusively on proving the subjective intention of an individual trader.

Operational resilience is another essential regulatory issue. Clapham, Haferkorn, and Zimmermann (2023) show that the technical interruption of HFT connections on the Xetra platform significantly reduced liquidity and increased trading costs, even though the market did not collapse. The event demonstrates that financial-market stability depends not only on the algorithms used by trading firms, but also on the reliability of exchange gateways, data feeds, communication systems, and order-routing infrastructures. Operators and regulators therefore need procedures for managing technical failures, testing algorithmic systems, monitoring unusual order activity, and ensuring that trading can continue in an orderly manner when a particular technological component becomes unavailable.

The ethical and regulatory challenges become even more complex when AI and machine learning are incorporated into HFT systems. AI-driven models may react to patterns that are difficult for their users or regulators to interpret, and their simultaneous use by multiple institutions may generate correlated behaviour. As already noted (Tse, 2024), market makers may become reluctant

to trade when they cannot adequately hedge their positions; this risk is amplified when the internal decision-making process of an AI model is opaque, making it difficult to determine why liquidity was withdrawn, why an order was submitted, or whether a particular sequence of trades resulted from a technical error, a risk-management rule, or a manipulative strategy.

These issues suggest that regulation should combine innovation-friendly principles with safeguards directed at market integrity and resilience. A complete prohibition of HFT would remove potential benefits related to liquidity, price discovery, and price integration across venues. Conversely, an entirely unregulated environment could allow speed advantages, correlated algorithms, and fragile liquidity to generate costs for slower investors and systemic risks for the wider market. A proportionate regulatory framework should therefore focus on transparency, pre-trade risk controls, algorithmic testing, surveillance, circuit breakers, operational resilience, and clear responsibility for the conduct of automated systems.

Overall, HFT and AI-driven trading create a regulatory dilemma rather than a simple case for either intervention or non-intervention. These technologies can narrow spreads, improve the speed of price discovery, and support continuous trading, but they can also increase adverse selection, amplify feedback loops, encourage manipulative practices, and widen the technological gap between market participants. The evidence reviewed in this chapter suggests that the central objective of regulation should be to preserve the economic functions of automated trading while limiting its destabilizing effects. This requires rules capable of adapting to technological change, because static regulation may quickly become ineffective as trading systems, market structures, and AI-based strategies continue to evolve. Building on this discussion of the ethical and regulatory issues associated with AI-driven and high-frequency trading.

CONCLUSIONS

This thesis set out to understand how artificial intelligence is reshaping online trading and, more broadly, the way financial markets themselves function. The journey moved from a basic definition of financial markets and the historical rise of electronic trading, through the development of AI techniques in finance, to the specific applications of AI-driven trading and what they mean for market efficiency, volatility, liquidity, investor behaviour, and regulation.

The first chapter argued that financial markets cannot be reduced to a fixed set of institutions: they are better understood as an evolving architecture built to perform enduring functions — channeling savings, sharing information, allocating risk — whose institutional form keeps adapting to technology, transaction costs, and regulation. The shift from floor-based trading to electronic and online systems, together with the behavioural changes that accompanied it, made this ongoing adaptation visible in practice and laid the technological and institutional groundwork on which AI-driven trading now operates.

The second chapter situated artificial intelligence within that groundwork, tracing its trajectory from early expert systems to today's machine-learning and natural-language-processing tools, and showing how these have found their way into prediction, optimization, and language-related tasks across financial services. Robo-advisory platforms and automated data-analysis tools stood out as clear illustrations of a broader pattern: artificial intelligence tends to automate specific, well-defined tasks rather than replace financial decision-making altogether.

Building on this foundation, the third chapter turned to AI-driven trading itself, examining algorithmic execution, machine-learning-based market analysis, and sentiment-based prediction, with a real-world illustration drawn from Cube Investimenti's Dual Sentiment Index. Across these applications, artificial intelligence proved genuinely useful for its speed, consistency, and ability to process large and varied datasets. It nonetheless carries real limitations, tied to data quality, model opacity, and the historical assumptions embedded in most learning algorithms — limitations that become most visible in unusual episodes such as the 2010 Flash Crash.

The fourth chapter carried these trade-offs to the level of the financial system as a whole, yielding a conditional picture rather than a uniform one. Under ordinary conditions, algorithmic and high-frequency trading tend to narrow spreads, deepen order books, and speed up how quickly information is incorporated into prices; under stress, the same technologies can instead amplify

volatility, withdraw liquidity precisely when it is most needed, and concentrate advantages in the hands of whoever is fastest and best equipped. Investor behaviour has not disappeared under automation but has simply been redirected: institutional investors split up and conceal their orders, liquidity providers reprice the risk of trading against better-informed counterparties, and retail investors operate within a market increasingly shaped by speed and algorithmic feedback loops. This leads to the chapter's central conclusion regarding regulation: the aim should be neither to prohibit nor to unconditionally endorse algorithmic and AI-driven trading, but to preserve what makes it efficient while embedding the transparency, testing, and resilience required to contain its capacity to destabilize markets.

Drawing these four chapters together, artificial intelligence emerges as a force that has not altered the fundamental purpose of financial markets, but has changed, often substantially, the technological and institutional means through which that purpose is fulfilled. AI-driven trading is neither a straightforward improvement over traditional approaches nor an inherently destabilizing force: its effects depend on prevailing market conditions, on the quality of the data used to train the underlying models, and on the robustness of the surrounding regulatory and operational infrastructure.

This work is, of course, subject to certain limitations, which merit explicit acknowledgement. The empirical evidence discussed in Chapter 4 draws primarily on a limited number of natural experiments and event studies — most notably the Xetra interruption and the Brexit-related crash — concerning specific markets, instruments, and time periods, and their generalization to other trading venues or market regimes should be approached with due caution. Certain mechanisms discussed in relation to investor behaviour and manipulation-related practices were illustrated through more practitioner-oriented material, employed not as primary evidence but as a means of broadening the range of mechanisms under consideration, and should be read accordingly. Similarly, the Dual Sentiment Index discussed in Chapter 3, while a valuable real-world illustration, rests on a relatively limited historical sample.

Looking ahead, this analysis could be extended by examining a wider range of trading venues and asset classes, by drawing on further natural experiments as AI-driven trading becomes more widespread, and by looking more directly at how regulatory frameworks such as MiFID II evolve in response to the growing use of machine-learning techniques in trading. As AI-driven tools continue to spread across both institutional and retail trading, understanding how their benefits and

risks are distributed among market participants will remain a central question — both for those studying financial markets and for those tasked with regulating them.

BIBLIOGRAPHY

Abreu, D. (2024). High frequency traders and liquidity. SSRN Working Paper.

Bahoo, S., Cucculelli, M., Goga, X., & Mondolo, J. (2024). Artificial intelligence in finance: A comprehensive review through bibliometric and content analysis. SN Business & Economics, 4(2), Article 23.

Barber, B. M., & Odean, T. (2001). The Internet and the investor. Journal of Economic Perspectives, 15(1), 41–54.

Barber, B. M., & Odean, T. (2002). Online investors: Do the slow die first? The Review of Financial Studies, 15(2), 455–488.

Belanche, D., Casaló, L. V., & Flavián, C. (2019). Artificial intelligence in FinTech: Understanding robo-advisors adoption among customers. Industrial Management & Data Systems, 119(7), 1411–1430.

Ben Ammar, I., Hellara, S., & Ghadhab, I. (2020). High-frequency trading and stock liquidity: An intraday analysis. Research in International Business and Finance, 53, Article 101235.

Ben Ammar, I., & Hellara, S. (2022). High-frequency trading, stock volatility, and intraday crashes. The Quarterly Review of Economics and Finance, 84, 337–344.

Cao, L. (2021). AI in finance: Challenges, techniques, and opportunities. ACM Computing Surveys, 55(3), Article 64.

Cavalcante, R. C., Brasileiro, R. C., Souza, V. L. F., Nobrega, J. P., & Oliveira, A. L. I. (2016). Computational intelligence and financial markets: A survey and future directions. Expert Systems with Applications, 55, 194–211.

Chami, R., Fullenkamp, C., & Sharma, S. (2009). A framework for financial market development (IMF Working Paper No. 09/156). International Monetary Fund.

Chung, H., Sheu, H.-J., & Hsu, S. (2010). Trading platform, market volatility and pricing efficiency in the floor-traded and E-mini index futures markets. International Review of Economics & Finance, 19(4), 742–754.

Clapham, B., Haferkorn, M., & Zimmermann, K. (2023). The impact of high-frequency trading on modern securities markets: An analysis based on a technical interruption. Business & Information Systems Engineering, 65(1), 7–24.

Committee on the Global Financial System. (2001). The implications of electronic trading in financial markets (CGFS Papers No. 16). Bank for International Settlements.

Cube Investimenti. (2026). Dual Sentiment Index (DSI): Internal proprietary documentation. Unpublished internal material, made available to the author during a curricular internship.

Davidson, P. (2000). Is a plumber or a new financial architect needed to end global international liquidity problems? World Development, 28(6), 1117–1131.

Goldstein, I., Spatt, C. S., & Ye, M. (2021). Big data in finance. The Review of Financial Studies, 34(7), 3213–3225.

Goldstein, I., Spatt, C. S., & Ye, M. (2025). The next chapter of big data in finance. The Review of Financial Studies, 38(3), 605–622.

Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. Journal of Behavioral and Experimental Finance, 32, Article 100577.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223–2273.

Hendershott, T., Jones, C. M., & Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? The Journal of Finance, 66(1), 1–33.

Hendershott, T., & Riordan, R. (2013). Algorithmic trading and the market for liquidity. Journal of Financial and Quantitative Analysis, 48(4), 1001–1024.

Jacob Leal, S., & Napoletano, M. (2016). Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading (OFCE Working Paper No. 2016-12).

Joby, N. P., & Vandana, V. C. (2024). The impact of artificial intelligence on stock market trading: Trends and implications. ShodhKosh: Journal of Visual and Performing Arts, 5(1), 1694–1700.

Khan, K., Shaikh, S., Sherwani, F. K., & Khan, M. U. (2024). Unlocking the application of AI for portfolio maximization investment decision. In Conference proceedings of the Centre for Business & Economic Research: ICGEEE-2024 (pp. 122–130). Centre for Business & Economic Research.

Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. The Journal of Finance, 72(3), 967–998.

Levine, R. (2005). Finance and growth: Theory and evidence. In P. Aghion & S. N. Durlauf (Eds.), Handbook of economic growth (Vol. 1A, pp. 865–934). Elsevier.

Lin, P.-H., Rakity, I., Stan, C. G., & Vijayakumar, S. (2024). Assessing the impact of high-frequency trading on market efficiency and stability. OxJournal.

Liu, J., & Fu, S. (2024). Financial big data management and intelligence based on computer intelligent algorithm. Scientific Reports, 14, Article 59244.

Merton, R. C., & Bodie, Z. (2004). The design of financial systems: Towards a synthesis of function and structure (NBER Working Paper No. 10620). National Bureau of Economic Research.

Milana, C., & Ashta, A. (2021). Artificial intelligence techniques in finance and financial markets: A survey of the literature. Strategic Change, 30(3), 189–209.

Mudinas, A., Zhang, D., & Levene, M. (2019). Market trend prediction using sentiment analysis: Lessons learned and paths forward.

Pagnotta, E. S., & Philippon, T. (2018). Competing on speed. Econometrica, 86(3), 1067–1115.

Park, J. (2025, April 4). Algorithmic trading and market volatility: Impact of high-frequency trading. Michigan Journal of Economics.

Rohan, A., Hossen, M. D., Pranto, M. N., Hossain, B., Yoshi, A. M., & Islam, R. (2026). Artificial intelligence in financial market prediction: Advancements in machine learning for stock price forecasting. Frontiers in Artificial Intelligence, 8, Article 1696423.

Stoll, H. R. (2006). Electronic trading in stock markets. Journal of Economic Perspectives, 20(1), 153–174.

Tse, W.-M. (2024). High-frequency trading, asset pricing, and market microstructure. SSRN Electronic Journal.

World Bank. (2002). World development report 2002: Building institutions for markets. Oxford University Press.

Zhu, F. (2025). High-frequency trading: Impact on market liquidity, volatility, and regulatory challenges. FIEM 2025 Conference Proceedings, 163–168.