

This electronic thesis or dissertation has been downloaded from the King's Research Portal at <https://kclpure.kcl.ac.uk/portal/>



Opinion Formation and Herding in Financial Markets

Wang, Chaoran

Awarding institution:
King's College London

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

END USER LICENCE AGREEMENT



Unless another licence is stated on the immediately following page this work is licensed

under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International

licence. <https://creativecommons.org/licenses/by-nc-nd/4.0/>

You are free to copy, distribute and transmit the work

Under the following conditions:

- Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).
- Non Commercial: You may not use this work for commercial purposes.
- No Derivative Works - You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you receive permission from the author. Your fair dealings and other rights are in no way affected by the above.

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Opinion Formation and Herding in Financial Markets



By

Chaoran Wang

The Department of Informatics

King's College London

This thesis is submitted for the degree of

Doctor of Philosophy

April 2023

Abstract

In financial markets, every investor seeks and receives information to decide how they should act (e.g., buy or sell a certain asset). In certain social circles, investors also learn about the decisions of other investors and they might sometimes ignore their own information and take the same decisions as other investors. This phenomenon is known as "herding effect". Many believe that herding can be one of the main causes of crashes and bubbles in financial markets. In this thesis, we adopt empirical methods to explore why investors try to imitate others, the impact of herding on financial markets and whether the trading mechanism used in the market affects herding.

Towards this goal, we connect opinion formation dynamics with herding in financial markets. We model social connections between the traders in different market environment as a graph and adopt a well-established opinion diffusion dynamics. Opinions are translated to trading positions and market prices evolve accordingly. We relate the shape of the graph social network to the equilibria of a game defined as follows. The players are traders that can strategically decide whether to follow the wisdom of the crowd or act upon their own beliefs. Their payoffs are defined as the wealth they accumulate from trading. We adopt Empirical Game-Theoretic Analysis (EGTA) to compute the equilibria of our games.

We first explore the impact of social connections between market participants on herding and market stability in a hypothetical market environment, where orders are always executed at the desired price. We show that the larger the traders' neighbourhood in the social network, the more the traders are willing to imitate others and the less volatile the stock price is. However, when every trader in the market has perfect knowledge of the opinions of all the other traders, the market will still exhibit crashes and bubbles. The definitions of crashes and bubbles in our research are based on changes in stock prices and are inspired by the financial concept of Maximum Drawdown.

The mechanics of trading in an order-driven market environment can influence the behaviour of traders and the idealised setting in our simulated market environment is too simplistic to model real markets. We then investigate opinion formation and herding in order-driven financial markets, which are widely used for many asset classes. We concentrate on Continuous Double Auctions, the principle trading mechanism in this class, and consider two forms of order queuing mechanisms: price-time priority, the de-facto standard, and spread-price/time priority, an alternative recently defined in literature to reduce toxic order flows due to latency arms race. We find that our conclusions are robust and hold in both these realistic market environments; the stronger the social connections between the agents, the more pronounced the herding. Furthermore, our empirical research shows that as the market gives more weight to spread, it becomes more stable thus confirming the findings of related work in our setup.

We conclude our work by enlarging the set of strategies that agents use. We use a meta-game to simplify the actual large game and explore herding of different types of investors in the market with different social connections. The results show that the herding is more pronounced among long-term investors than short-term investors. We see our work as the introduction of a framework that can be used to study more questions about herding in financial markets and other complex systems.

Acknowledgment

I would first like to express my sincerest thanks and appreciation to my supervisor, Professor Carmine Ventre. His guidance, support and encouragement have been invaluable throughout my Ph.D studies. His words were enlightening and refreshing. I am very grateful to him for taking the trouble to recommend papers to me at the beginning of my studies, so that I could find a direction of research that interested me. He has guided me step by step into the academic frontier of finance and computer science. In my studies and in my life, he has always been a source of warmth and hope for me in a foreign country. When I questioned my research, Carmine Ventre was always there to encourage me and offer enthusiastic help. During the difficult times of COVID-19, he remained constantly concerned about my physical and mental health. I would also like to thank my sister Ji Qi, whose care and support over the past five years has helped me to overcome many of my academic and life challenges. They have been with me from my Master's to Doctoral studies. I am honoured and lucky to have met such wonderful supervisor and seniority, who have been my mentors and friends. In addition, thanks to Anton Golub and Ji Qi for sharing the code of their research papers.

At last, I would like to thank my parents. My parents have raised me with boundless love since I was a child. I thank them for supporting me in my choices in life. They have done everything in their power to provide a better environment for me to study and live in. Without their support and encouragement, I would not have been able to study abroad for my Ph.D degree.

Contents

Abstract	iii
Acknowledgement	v
List of Figures	xi
List of Tables	xiii
Abbreviations	xv
1 Introduction	1
1.1 Motivation	3
1.2 Outline of Thesis	4
1.3 Research questions	6
1.4 Thesis contributions	6
2 Literature review	7
2.1 Agent-based model	7
2.2 Opinion Formation	9
2.3 Game theory	10
2.4 Nash equilibrium	11
2.5 Empirical Game-Theoretic Analysis	12
2.6 Opinion Formation in Game Theory	13
2.7 Opinion Formation in finance	15
2.8 Crashes and Bubbles	16

2.9	Maximum Drawdown	18
2.10	Herding effect	19
2.11	Order-driven markets	20
2.12	Continuous double auctions	22
2.13	High-frequency trading	23
3	Opinion formation, herding and market stability	25
3.1	Introduction	25
3.2	An agent-based model (ABM) of Opinion Formation	27
3.2.1	Two sources of information for opinions	28
3.2.2	Opinion formation	28
3.2.3	Trading decision	29
3.2.4	Price clearing	30
3.2.5	Cash and stock	31
3.2.6	Crashes and bubbles	31
3.2.7	Set up parameters	33
3.2.8	Results	34
3.2.9	The initial flash crash	39
3.2.10	Different opinion formation rules	41
3.3	Empirical Game Theory Analysis (EGTA)	43
3.3.1	An example of Nash equilibrium calculation	44
3.3.2	Three Strategies	45
3.3.3	Payoffs	45
3.3.4	Basic experiment setup	46
3.3.5	Choice of the value of α	46
3.3.6	Results of EGTA	46
3.4	Different social connections between agents	48
3.5	Changing risk attitude and trading size	56
3.5.1	About the risk attitude	56
3.5.2	About the trading size	57

3.6	Conclusions	57
4	Opinion formation and herding in order-driven markets	59
4.1	Introduction	59
4.2	The model of limit order book (LOB)	60
4.2.1	Limit order book (LOB)	60
4.2.2	Trading patterns of LOB	61
4.2.3	Agent-based model of LOB	62
4.3	Empirical Game Theory Analysis (EGTA)	64
4.3.1	Three strategies	64
4.3.2	Payoffs	65
4.3.3	Setup of experimental parameters	65
4.3.4	The value of α	66
4.3.5	Experimental results	66
4.3.6	Different social connections	67
4.4	Two types of investors	69
4.4.1	Long-term and short-term investors	69
4.4.2	Experimental results and analysis	70
4.5	Spread/price–time priority	71
4.5.1	Market maker (MM)	71
4.5.2	Order matching mechanism	74
4.5.3	Experimental setup and results	75
4.5.4	Liquidity and stability of the market	77
4.6	Conclusions	80
5	Exploring herding in order-driven markets through a meta-game	81
5.1	Introduction	81
5.2	Meta-game	82
5.3	Experimental setup	83
5.3.1	Meta-strategies	83
5.3.2	Payoffs	83

5.4	Experimental results and analysis	85
5.4.1	Different social connections in the meta-game	86
5.4.2	Long-term and short-term investors	89
5.5	Conclusions	91
6	Conclusions and Future Work	93
6.1	Research Summary	93
6.2	Research Contributions	93
6.3	Academic Publications	95
6.4	Limitations	95
6.5	Future Work	96
	Bibliography	99

List of Figures

3.1	$G = (N, E)$	27
3.2	Trading decision	29
3.3	Definition of crash and bubble	32
3.4	Another situation	33
3.5	Price time series when $\alpha = 0.1$	37
3.6	Price time series when $\alpha = 0.6$	37
3.7	Price time series when $\alpha = 0.65$	38
3.8	Price time series when $\alpha = 0.7$	38
3.9	Price time series when $\alpha = 0.5$	39
3.10	First 100 time steps of $b_i(t)$	40
3.11	Distribution of original set of $b_i(t)$'s	41
3.12	Distribution of new set of $b_i(t)$'s	42
3.13	$\alpha = 0.5$ with new $b_i(t)$	42
3.14	Price time series when $\alpha = 0.021$	43
3.15	Nash Equilibrium	47
3.16	Crashes and Bubbles at the first Nash equilibrium	48
3.17	Crashes and Bubbles at the second Nash equilibrium	49
3.18	Nash equilibria for different graphs	50
3.19	10-a-group and 20-a-group	52
3.20	30-a-group and 40-a-group	53
3.21	50-a-group and 80-a-group	53
3.22	90-a-group and 100-a-group	54

4.1	An agent based model of LOB	63
4.2	The Nash equilibrium of EGTA	66
4.3	The proportion of agents at Nash equilibrium of EGTA	67
4.4	The NE of different social connections	68
4.5	Investors who choose “Imitation” strategy	72
5.1	An example of a heuristic payoff table for 3 meta-strategies and 3 individuals	84
5.2	The Nash equilibrium of meta-game	85
5.3	The proportion of agents at Nash equilibrium	86
5.4	The NE of different social connections in meta-game	87
5.5	Proportion of agents choosing “Neutral” strategy	88
5.6	Investors who choose “Imitation” strategy in meta-game	90

List of Tables

3.1	Results	35
3.2	Duration steps of each Crash and Bubble	35
3.3	The payoff matrix for this game	44
3.4	Number of agents at Nash equilibrium	52
3.5	Crashes and Bubbles at equilibrium in different graph	55
4.1	Example of a limit order book (LOB) for a stock	62
4.2	The proportion of agents at Nash equilibrium	69
4.3	The NE of short-term investors	71
4.4	The table on the left is $\mu = 0.1$ and the right one is $\mu = 0.3$	76
4.5	The table on the left is $\mu = 0.5$ and the right one is $\mu = 0.7$	76
4.6	$\mu = 0.9$	76
4.7	Average volume for different (β, μ)	78
4.8	Average volatility for different (β, μ)	79
5.1	The proportion of agents at Nash equilibrium in meta-game	88
5.2	The NE of short-term investors in meta-game	89

Abbreviations

ABM	agent-based model
MM	market maker
CDA	continuous double auction
EGTA	empirical game-theoretic analysis
HFT	high-frequency trading
LOB	limit order book
NE	Nash equilibrium
ML	machine learning
SVB	Silicon Valley Bank

Chapter 1

Introduction

The efficient market hypothesis (EMH) was introduced in the mid-1960s and is widely accepted as an important financial theory. According to the theory, financial markets are incapable of generating their own internal forces to upset the equilibrium, and large price changes are simply the result of markets reacting to new external information or changing fundamentals. Thus, according to the EMH, there is no room for asset prices to bubble or collapse [1]. However, history and past data show that the EMH has not always been successful in explaining phenomena in financial markets. From the Dutch tulip bubble in 1626, the South Sea corporate bubble in the UK, the French Mississippi corporate bubble in the early 1700s and the Japanese bubble in the 1980s, to the recent US subprime mortgages and the 2008 financial crisis, bubbles and crashes have occurred and continue to occur in financial markets. People increasingly believe that traditional economic models cannot explain the bubbles and crashes of financial markets and related crises. Therefore, agent-based models (ABM) as an alternative method to better understand the complex dynamics of financial markets have received attention [2–6].

Most explanations of crashes look for possible mechanisms or effects that operate on very short time scales (hours, days or weeks at most), see e.g. [7]. But the root cause of the crash could have been months or years before it happened [8]. Crashes happen because the market has entered a phase of instability, near the apex of a bubble, and any small disturbance or process can lead to a crash. Like a feather on one's finger, it is this unstable position that

causes the feather to fall. The cause could be the instability of the finger or the wind. The root cause of the collapse is the instability of the position, not what triggers it. Not a single snowflake is innocent when an avalanche occurs.

There have been many financial crises throughout history, such as the Dutch financial bubble of 1639. It was caused by the speculation of a large number of people on tulips [9]. In 1907, speculation was rife in the US banking sector and the entire financial market was in a state of extreme speculation. In October of that year, the failure of the third largest trust company in the US, the Knickerbocker Trust Company, to acquire shares in the United Copper Company sparked rumours that Knickerbocker was about to go bankrupt. This rumour led to a mad run on the bank's customers and triggered the financial crisis on Wall Street [10]. The US stock market crash of 1929 was caused by a massive sell-off of stocks by investors. The investors didn't actually know what was happening, they just copied each other because they saw other investors dumping stocks [11]. Silicon Valley Bank (SVB) was founded in 1983 in the United States as a subsidiary of Silicon Valley Bank Financial Group. On 10 March 2023, Silicon Valley Bank, the 16th largest bank in the United States, collapsed. It was the second-largest bank failure in US history and the largest since the near collapse of the US financial system in 2008. The Federal Reserve's aggressive interest rate hike policy acted as a trigger that led to the collapse of Silicon Valley Bank. The COVID-19 epidemic is an unprecedented blow to the global economy. The Federal Reserve controls unemployment and inflation by raising interest rates. This has led to a high dollar index and rising yields on US Treasuries. So people were taking cash out of the bank to buy Treasuries, which led to a liquidity crunch in the banks and was the first cause of SVB's bankruptcy [12]. Silicon Valley Bank bought a lot of long-term US Treasuries before the Fed raised interest rates, which caused SVB's cash flow to become even tighter. So Silicon Valley Bank had to sell off its stock and sell the treasury bonds it had previously bought to recoup its capital. Following the announcement, prominent venture capitalists began warning about the situation on social media, some of whom urged SVB customers to withdraw their deposits. People then began to worry whether the money they had deposited with Silicon Valley Bank was still safe. When the panic began to spread, SVB customers withdrew more than 42 billion dollars in a single day in what became known

as the first “Twitter-fueled bank run”. This is the second reason why Silicon Valley Bank went bankrupt [13]. When we look back at the financial crises that have occurred throughout history, it is easy to see that much of the instability in the financial markets was due to social phenomena such as herding.

Financial markets are an important part of the modern economy and their study is essential for investors, policymakers and businesses to understand how they work and their impact on the wider economy [14]. This thesis applies game theory to financial markets and makes a number of innovative attempts and contributions. This chapter presents the motivation and the outline of the thesis.

1.1 Motivation

Financial markets play a key role in supporting business growth and global economic development. By providing access to capital, managing financial risk, allocating resources efficiently, promoting innovation and entrepreneurship, and facilitating international trade, financial markets contribute to increased productivity, competitiveness and prosperity around the world. Financial markets are changing rapidly as technology evolves. Algorithmic trading and high frequency trading are now dominant and these trading methods have not only brought liquidity but also made financial markets unstable [15]. Two of the most notable unstable states of financial markets are market crashes and bubbles [16]. The study of crashes and bubbles in the financial market has therefore become particularly important. Many people believe that the “herding effect” has been a major cause of financial market crashes and mass panics throughout history, see e.g. [17]. The study of herding has therefore become particularly important in the field of finance [18–20]. In financial markets, the herding effect manifests itself in investors following the behaviour or opinions of others rather than making independent decisions based on their own analysis [21]. This can lead to market bubbles or crashes, as prices become disconnected from the fundamentals of the asset [22]. Understanding herding can also help in developing an investment strategy. By studying herding in finance, we can get a better understanding of the dynamics of financial markets and develop strategies to mitigate the negative effects of herd behaviour.

Opinion formation plays a crucial role in financial markets as it influences the behaviour of market participants. Participants in the market form opinions based on a variety of factors such as economic data, company news, political events and market trends [23]. These opinions can have a significant impact on market prices and ultimately determine the direction of the market. Opinions based on incomplete information, bias and sentiment often lead to market inefficiencies and mispricing [24]. As a result, opinion formation by participants can also lead to market volatility and instability. Overall opinion formation is an integral part of financial markets and investors should be aware of its impact on the market.

It is well known that real market environments are extremely complex and unpredictable. Empirical game theory analysis relies on quantitative methods and statistical models, making it a rigorous and objective method for analysing financial markets [25]. By applying empirical game theory analysis to large and complex systems such as financial markets, researchers can better understand the underlying dynamics of financial markets by modelling how different participants interact and respond to each other's decisions [26]. Overall, empirical game theory analysis provides a powerful tool for understanding financial markets and for making investment decisions. By analysing market data and modelling the behaviour of market participants, we can gain insight into the dynamics of markets and identify opportunities to improve market efficiency. This thesis therefore delves into opinion formation and herding in different financial market environments by empirical game theory analysis.

1.2 Outline of Thesis

The first chapter is an introduction to the thesis as a whole, including motivation and the outline of the thesis. Chapter 2 is the literature review. All my contributions are contained in Chapters 3, 4 and 5. Chapter 6 presents the conclusions, limitations of the experiment and future work. The details of each chapter are summarised below:

- ▶ Chapter 2 provides the theoretical basis and sources for this thesis by reviewing the relevant background to the research topic. It first introduces the agent-based model, opinion formation and herding, which are the financial background underlying our research. This is followed by a detailed presentation of the game theory background as a technical premise,

a section that includes game theory, evolutionary game theory and empirical game theory analysis.

- ▶ Chapter 3 defines a new ABM to investigate the extent to which social ties between market participants influence herding and market stability. This is achieved by incorporating classical opinion dynamics in the decision making process of traders and using EGTA to calculate equilibria for appropriately defined strategic games. In these games, traders wish to maximise profits while balancing private and public information about the asset in question. Broadly speaking, we show that the more dense the social graph, the more likely herding are to occur.
- ▶ Chapter 4 applies our agent-based model to order-driven markets and uses empirical game-theoretic analysis to explore opinion formation and herding in financial markets with different social connections. We also investigate the impact of herding on different types of investors. The experimental results show that the herding persists in order-driven markets. In the same market environment, the herding is more pronounced for long-term investors. To check the robustness of our findings to different trading mechanism, we study a new order matching mechanism, called spread/price-time priority mechanism [27], to determine the priority of order execution. The results show that the herding in the market becomes less pronounced as the weight of spread in the order-matching mechanism increases. The spread/price-time priority mechanism does improve the liquidity and stability of the market, but we need to find a specific set of parameter settings by analysing equilibrium market characteristics to make the spread/price-time priority order matching mechanism ensure both low market volatility and a considerable trading volume.
- ▶ Chapter 5 proposes a meta-game model to analyse complex games with meta strategies. The meta-game is concerned with the strategic choices that players make in the actual game. In reality, the strategy set of traders has infinite size and we would like to study herding in this case. So we use a meta-game to explore the impact of herding on different types of investors in the market under different social connections. Three meta-strategies are defined to simplify the game, they are “Imitation”, “Neutral” and “Confident”. The

experimental results show that herding is more pronounced for long-term investors than for short-term investors in order-driven markets. Under the same connectivity conditions, more agents in the meta-game chose 'Neutral' meta-strategy than in the previous game.

- ▶ Chapter 6 gives the conclusions and summarises the contributions and limitations of our work. Our ideas for future work offer some new directions for the study of opinion formation and herding in finance.

1.3 Research questions

The research question of this thesis can be summarized as follows:

- How opinion formation in different market environments and social connections affects herding effect.
- Why investors try to imitate their neighbours.
- What herding effect does to financial markets.
- Does the trading mechanism used in the market affect the herding effect?
- The impact of herding effects on different types of investors.

1.4 Thesis contributions

There are five main contributions to this thesis. First, we adopt a new opinion formation method to explore herding in financial markets. Second, we try to integrate the agent-based model and empirical game-theoretic analysis methods. Then we try different market environments and a new order-matching mechanism. Finally one is the contributions to future research. Our research focused on opinion formation and herding in financial markets. We see our work as the introduction of a framework that can be used to study more questions about herding in financial markets.

The link of code: <https://github.com/Chaoran799/Opinion-Formation-and-Herding-in-Financial-Markets.git>

Chapter 2

Literature review

2.1 Agent-based model

Agent-based models (ABMs) represent a category of computational models designed to replicate the actions and interactions of independent agents and assess their collective impact on a system. This approach draws from various disciplines, including game theory, complex systems, computational sociology, and multi-agent systems. ABMs function at a microscopic level, simultaneously simulating the behaviors and interactions of numerous agents to replicate and predict the emergence of complex phenomena. This computational methodology empowers analysts to construct, analyze, and experiment with synthetic environments composed of interacting agents [28]. ABMs operate on several foundational assumptions, including the idea that agents possess varying degrees of rationality, as determined by either traditional or adaptive behavior rules. These agents exhibit heterogeneity in their characteristics and engage in learning processes [29]. This versatile modeling technique finds applications in analyzing financial markets, particularly those comprised of diverse entities and intricate interactions. ABMs offer a unique lens through which to examine economic phenomena, including financial crises, that are challenging to elucidate using conventional economic tools.

Research in ABMs has evolved significantly since its inception. Recent studies have further advanced the field by incorporating real-world data, sophisticated agent behaviors, and nuanced market structures. For instance, some researchers have extended ABMs to incorpo-

rate sentiment analysis, news data, and social network dynamics, providing a more realistic portrayal of market behavior. These enhancements have offered valuable insights into the role of information and behavioral biases in shaping financial markets. Notable works include the incorporation of agent learning mechanisms, such as reinforcement learning and deep learning, to better mimic the adaptability and evolving strategies of market participants. Previous research endeavors have explored financial markets using ABMs, with one notable early example being the "Artificial Stock Market" model by Palmer, et al. (1994) [30]. In this model, the authors constructed a basic representation of the stock market and demonstrated its capacity to generate price bubbles, crashes, and enduringly high trading volumes. This pioneering work has set the stage for subsequent studies in the realm of ABM-based financial market analysis. Antoine et al. develop an agent-based model of financial markets that contains multiple assets belonging to the fixed income or equity asset classes [31]. The aim is to reproduce the main facts about the emerging dynamics of the yield curve in the fixed-income market. They compare the dynamics generated by different management processes of the risk-free rate with the dynamics of the historical U.S. Treasury market. It is demonstrated that their ABM is able to reproduce the main features of the autocorrelation surface of the volatility of U.S. Treasury yields to maturity over a selected time horizon.

Moreover, recent research on ABMs has focused on exploring the implications of policy interventions and regulations in financial markets. ABMs enable researchers to simulate the effects of regulatory changes, stress tests, and market interventions, providing policymakers with insights into potential outcomes and unintended consequences. In addition to examining market dynamics, ABMs have been employed to study systemic risk in financial systems. These models facilitate the identification and evaluation of interconnectedness and vulnerabilities within financial networks, assisting in the design of more robust and resilient financial systems. Overall, the continuous evolution of ABMs in financial market analysis has expanded the scope and depth of our understanding of market behavior. This evolution has been instrumental in addressing the limitations of traditional economic models and has paved the way for more realistic and comprehensive approaches to studying financial markets and economic phenomena.

2.2 Opinion Formation

Opinion formation describes the dynamics of opinions across a set of interacting factors and is a powerful tool for predicting the evolution and diffusion of opinions. The study of opinion formation in complex networks has developed significantly in recent years, owing to the growing relevance of understanding how opinions and information travel through interconnected societies. One notable avenue of inquiry focuses on the integration of advanced computational techniques and real-world data to create more realistic models of opinion formation. This level of sophistication allows researchers to simulate opinion dynamics under a variety of conditions, contributing to a better understanding of how opinions evolve in different communities. F. Slanina and H. Lavicka describe the Sznajd model of opinion formation and social influence [32]. The Sznajd model is an economic physics model proposed in 2000. The Sznajd model implements a phenomenon called social validation [33]. In brief, the model posits that if two people have the same opinion, their neighbors will begin to agree with them. If the people around them disagree, their neighbors will begin to argue with them. Indeed, many people used ABM to research opinion formation. Others have investigated the links between network structure, viewpoint dynamics, and the ability of adversaries to artificially create differences [34]. These authors address these issues by extending models of viewpoint formation in the social sciences to represent scenarios familiar from recent events in response to them. In this scenario, external actors attempt to destabilize communities through fake news and bots through sophisticated information warfare strategies. They examine the nature of such attacks and consider the best tactics for finding entities that divide adversaries and are responsible for protecting networks from attacks. They then argue that network defenders mitigate such attacks through heterogeneous isolation of nodes. Finally, they generalize these results to two network structures, dynamic opinion processing and decoupling metrics. In summary, research on opinion formation has evolved to encompass a wide range of factors, including advanced computational modeling, the role of external manipulation, and cyber defense strategies. These research efforts have helped provide valuable insights into the complexity and dynamics of opinion formation, allowing us to gain a deeper understanding of the importance

of opinion formation.

2.3 Game theory

Game theory is the study of mathematical models of strategic interactions between rational subjects [35]. It has applications in all areas of social science as well as logic, systems science, and computer science. Game theory was pioneered and only exists as a distinct field by Princeton mathematician John von Neumann [36]. Game theory, as a branch of applied mathematics, provides tools for analysing situations in which parties (called players) make interdependent decisions. This interdependence leads each player to consider the possible decisions or strategies of the other players when formulating a strategy. The solution to a game describes the best decisions of players who may have similar, opposing, or mixed interests, and the possible outcomes of these decisions. Game theory attempts to mathematically and logically determine the actions that players should take to ensure the best outcome for themselves in the game. Games can be classified according to certain distinguishing characteristics, the most obvious of which is the number of players. Thus, a game can be designated as a single-player, two-player, or multi-player (with more than two players) game, each category having its own unique characteristics. In addition, the players can be either individuals or a team.

Nowadays, human beings live in a complex system with different social interactions. Individuals can engage in a variety of activities within society, such as sporting competitions, buying and selling transactions, and scientific research. In these activities, people behave according to their goals and their position. Advanced technology has not only changed our lives but has also brought us new problems. When high frequency trading is possible, it brings not only liquidity to the market but also conflict and risk. Game theory is the formal study of conflict and cooperation [37]. A game consists of players, strategies, and payoffs for each player in a determined environment [38]. Players are essentially agents of the game and strategies are the actions they can take. The payoff is the utility of each player given a certain outcome of the game. There is often strategic interaction between players, which means that payoffs depend not only on the actions taken but also on the actions taken by opponents. In the modern discipline, game theory refers to the study of multiple individuals or teams in a game in a de-

terminated environment, using each other's strategies to implement the corresponding strategies in the game. Throughout the game, the strategies of the other players are important to each player.

The applications of game theory extend far beyond economics to include fields such as political science, biology, and computer science. It has been used to analyze international conflicts, strategic interactions in biology, and the design of algorithms for multi-intelligence systems. In summary, game theory has come a long way since its inception, evolving from a theoretical construct into a practical tool for understanding a wide range of interactions and decision-making processes. Its contributions span multiple fields, making it an important area of study for understanding strategic behavior and rational decision-making in cooperative and competitive environments.

2.4 Nash equilibrium

The concept of Nash equilibrium, named after the mathematician John Nash, has played a key role in the development of game theory and its interdisciplinary applications. John Nash's seminal work on equilibrium points in non-cooperative games was presented in his 1950 doctoral dissertation, "Non-Cooperative Games" [39]. The Nash equilibrium is a central concept in game theory that describes a situation in which no participant has an incentive to unilaterally change his or her strategy, given the strategies chosen by the other participants.

In economics, one of the aims of analysing a game is to find the best strategy to maximise the player's rewards. In general, the optimal strategy depends on the actions chosen by the players. Finding the best strategy for a game is not just a simple question of optimisation. The best strategy should be the one that gives a higher payoff than other strategies, taking into account the possible actions of other opponents. The situation in which each player simultaneously responds optimally to the other's strategy is a Nash equilibrium. The Nash equilibrium represents the action profile of all players in the game and is used to predict the outcome of their decision-making interactions. It models a steady state (i.e. a combination of all players' strategies) in which no player can benefit by unilaterally changing his or her strategy [40].

Nash equilibria have been widely used in economics to analyze markets, oligopolies, auctions, and various competitive environments. The application of Nash equilibria in these contexts provides insights into pricing strategies, market competition, and industry dynamics. In summary, Nash equilibrium represents a fundamental concept in game theory that provides a powerful framework for analyzing strategic interactions. The applications of Nash equilibrium extend far beyond economics, influencing fields as diverse as biology, computer science, and political science.

2.5 Empirical Game-Theoretic Analysis

Empirical game theory analysis is an emerging empirical methodology that bridges the gap between game theory and simulation in practical strategic reasoning [41]. “Game-Theoretic Analysis” usually describes its subject games from their starting point, i.e., a formal model of multiple subject interactions. The equilibrium we obtain in the empirical game is considered likely to be relatively stable in a full game [42]. However, many interesting games go far beyond the bounds of manageable modeling and reasoning. The problem here is not just the complexity of the analysis task or finding a balance. For example, the Trading Agent Competitive Supply Chain Management (TAC/SCM) game [43]. This is a well-defined six-person symmetric game of incomplete information. This game poses a difficult challenge for game-theoretic analysis. Even if a complete strategy is given for all six agents, there is no obvious way to obtain the expected gains unless sampling is done from a random environment using an available game simulator. In Wellman’s article, they break down empirical game theory analysis into three basic steps which are: Parametrize Strategy Space, Estimate Empirical Game and Analyze Empirical Game. They apply this approach to a variety of games, especially market-based scenarios. In some cases, this approach is able to support conclusions in these games that cannot be reached through standard analytical methods. In EGTA, techniques from simulation, search, and statistics are combined with game theoretical concepts. Techniques from search and statistics are combined with concepts from game theory to describe the strategic properties of a domain [44]. It starts with a set of strategies, usually heuristic strategies derived from domain knowledge or experience, and often limited by meaningful features. The

basic steps in EGTA are to simulate a strategy profile and to identify the payoff. Based on the accumulated data, we propose an empirical game model. On this model, we can perform any standard calculation applicable to the game format. Based on these results, we can choose to refine the model by considering more strategies or strategy profiles or by obtaining more strategies or strategy profiles.

Empirical game theoretic analysis has been applied in a variety of fields, including economics, supply chain management, and multi-agent systems. It is particularly valuable in situations where traditional game theory may struggle to capture the complexity of real-world interactions. Empirical game theoretic analysis is a practical and flexible approach to understanding strategic interactions and equilibria in complex real-world scenarios. Combining game theory with simulation, search, and statistical techniques, it provides a powerful framework for modeling and analyzing strategic behavior across a wide range of domains, offering insights not typically available from traditional analytical approaches.

2.6 Opinion Formation in Game Theory

Game theory studies the situation where strategic players can modify the state of a given system without central authorization. Solution concepts, such as Nash equilibrium (NE), have been defined in order to predict the outcome of such situations. It has been pointed out that in a multiplayer game environment, the solution concept should be obtained through a decentralized and reasonably simple process. In many models, a group of people in a social network are studied, each of whom holds his or her own opinion and arrives at a common opinion by repeatedly averaging with neighbors in the network. Since they actually rarely reach a consensus, Morris H. Degroot presents a model that describes how a group of people can reach a consensus on a common subjective probability distribution of parameters by pooling their individual opinions [45]. He explicitly describes the process leading to consensus and identifies the common distribution achieved. Similarly, David Bindel et al. studies a related sociological model in which individuals' intrinsic beliefs counteract the averaging process and generate multiple opinions [46]. They interpret the repeated averaging process as the optimal response dynamics in the underlying game with natural returns and its limits as equilibrium. This al-

allows them to study divergence costs by comparing the difference between equilibrium costs and socially optimal costs. Mengbin Ye et al. attempt to study dynamically changing relative interactions, since interactions may change depending on the issue under discussion. Specifically, they study a matrix of relative interactions that change periodically with the issues [47]. Research suggests that the social power of individuals recognizes a cyclical solution. Thus, Diodato Ferraioli et al. have studied the computation of solution concepts through decentralized dynamics [48]. These algorithms allow participants to take turns to act in order to reduce their own costs, in the hope that the system will "balance out" quickly. They formally analysed the formation of opinions in social networks. Specifically, they studied the optimal response dynamics and showed upper and lower bounds on the convergence of the Nash equilibrium. They also study a noisy version of the best response dynamics, the logit dynamics, and show a range of results for the rate of convergence as the system noise varies. In the framework of strategic game theory, a method is proposed to simulate basic interactions between different individuals by Alessandro Di Mare and Vito Latora [49]. In their model, tolerance thresholds are defined so that individuals with differences of opinion greater than the threshold are unable to interact. They then considered individuals with different propensities to change views and the ability to persuade others. In this way, they obtained the so-called "stubborn individual and speaker" (SO) model. They explored the dynamics of the SO model through numerical simulations. In the paper of Markos Epitropou et al., a continuous opinion formation game with aggregation aspects is studied [50]. In order to find the interaction between global and local influence, they propose a model of an opinion formation game with aggregation. They show that there is a unique equilibrium in the average oriented opinion formation game. They show some results in the context of a general opinion game with negative effects and extend their results to the case where the opinions expressed must come from a restricted area. Kshipra Bhawalkar et al. propose a game model of opinion formation in social networks in which opinions themselves and friendships develop together [51]. In these models, nodes form their opinions by maximizing agreement with their friends, which is influenced by the strength of the relationship, which in turn depends on the differences in opinions with their respective friends. They obtain a function on how many nodes value their own opinions and how much

weight they give to links with friends who more closely agree with them. There is much more work in this area and I only cover the papers more relevant to this project.

2.7 Opinion Formation in finance

Wyart and Bouchaud study a general model of self-referencing behavior in financial markets [52]. In this model, agents use correlations between specific quantities of information and prices estimated through history to create strategies. The effect of these strategies on prices corrects the observed correlations and creates a feedback loop. They argue that these strategies can destabilize the market through efficient behavior. This mechanism can lead to overreaction and excessive volatility. However, their model does not take into account the effects of imitation through social networks. An “information cascade” occurs when a person trusts his neighbors’ information too much and does not consider his own [53]. “Information cascade” is a phenomenon described in behavioral economics and network theory, where many people make the same decision in a continuous manner. In these models, agents know that the information is incomplete. To supplement this information, they look around and ask their neighbors for hints. The agents imitate because they think that other agents have more information. The results show that under the condition that decisions are continuous, one after the other, and irreversible, an agent’s use of other agents’ decisions to make its own will lead to an “information cascade” with probability 1. Some studies have combined and optimized these two lines of research. Georges Hervas and Didier Sornette proposed a simple agent-based model to study how proximate causes of crashes or rallies relate to their underlying mechanisms and vice versa [54]. There are three sources of information in the paper: public information (news), information from their neighbors’ networks, and private information. Agents use the past correlation of these three sources of information to form and adjust their trading strategies. In their experiment, they simulated a market. In this market there are 2,500 agents trading the same stock at given time steps, i.e., the social graph is a grid. They explored the characteristics of crashes and bubbles by observing the change in the price of this stock. Each of their agents has four neighbors. It weights the neighbor information quantified by the coefficient $c_{1i} \in [0, C_1]$. In their conclusion, they argue that the resulting market is approxi-

mately efficient when an agent has little faith in its neighbors. They think that when the agent is more trusting of its neighbor is more prone to crash and bubble. In summary, the field of opinion formation in game theory offers a rich landscape for understanding how individuals' opinions and beliefs are influenced by, and influence, their strategic interactions.

2.8 Crashes and Bubbles

The study of bubbles and crashes in financial markets has been the subject of extensive empirical research. Bubbles, characterized by prolonged price increases, often precede market crashes. Understanding the origins and mechanisms of these phenomena is a central concern in both economics and finance. Blanchard's seminal work in 1979 introduced the concept of rational expectations in the context of bubbles. According to this theory, speculative bubbles eventually lead to market crashes, and these events are consistent with the assumption of rational expectations [55]. This theory laid the groundwork for understanding how rationality can coexist with bubbles. Bubbles can manifest in various forms, making them challenging to detect or reject. The diverse nature of bubbles in financial markets has led to ongoing debates about their existence and characteristics. Researchers have sought to identify rational bubbles in markets, and in many cases, they have found that phenomena like runaway asset prices and market crashes align with rational bubble dynamics [56]. Johansen and colleagues have explored the characteristics of growing bubbles in various markets. They have attributed the ever-increasing rate of price rises to the dynamics of noise traders, who contribute to the escalation of crash risk and, consequently, the occurrence of crashes [57]. An alternative perspective suggests that positive feedback from agents following certain strategies can drive prices up and increase crash risk. This positive feedback loop, rather than noise traders, is seen as a key driver of bubble dynamics [58]. Empirical studies have sought to shed light on the role of speculation in bubble formation and crashes. For instance, Lei et al. conducted experiments in asset markets, where speculation was not possible, to understand the dynamics of bubbles and crashes. Their findings suggested that rational speculation alone does not adequately explain the occurrence of bubbles [59]. Levine and colleagues attempted to explain bubbles in the context of limited rationality. Their focus on the role of bounded

rationality and behavioral factors in bubble dynamics aimed to provide a more comprehensive understanding of these phenomena [60]. Several researchers have utilized agent-based models to explore the existence and characteristics of bubbles and crashes. For example, Lee and Lee introduced the concept of "heterogeneous expectations" and "herding behavior" in their agent-based model. Their findings illustrated how heterogeneous expectations among agents could influence the likelihood of bubbles and crashes in financial markets [61]. Giardina and Bouchaud presented a complex market model that incorporates agents' strategies, wealth, and price dynamics. Their model introduced different mechanisms, leading to oscillatory phases with bubbles and crashes, intermittent phases, and rational market phases [62]. In the context of modern financial markets, Leal et al. developed an agent-based model to investigate the interaction between low-frequency and high-frequency trading. Their research revealed that the presence of high-frequency traders can exacerbate market volatility and contribute to flash crashes. The study emphasizes the role of high-frequency trading in shaping market dynamics [63]. Sornette presented a general theory of financial collapse and stock market instability in 2003 [64]. He discusses the limitations of standard analysis in describing the specificity of crashes and the main causes of crashes, such as imitation and herd behavior among investors. He identifies generic features of crashes. These characteristic patterns have been documented in virtually all crashes in developed and emerging stock markets, currency markets, corporate stocks, and so on. He reviews this discovery in detail and demonstrates how detailed predictions obtained from the model can be used to predict crashes. James et al. designed, implemented and evaluated a model based on hybrid micro-macro agents in which price impacts arise endogenously through the limit order placing activities of algorithmic traders [65]. They characterize the time windows available for regulatory intervention during the propagation of a flash crash crisis, and their results suggest that ex-ante preventive measures may be more effective than ex-post responses.

In conclusion, the empirical exploration of bubbles and crashes in financial markets involves a multidisciplinary approach, drawing from traditional economic theory, behavioral factors, and advanced agent-based modeling. The field continues to evolve, providing valuable insights into the dynamics of these market phenomena and their implications for market

stability and regulation.

2.9 Maximum Drawdown

Maximum Drawdown (MDD) is defined as the maximum loss that occurs from peak to trough over a specified time period. The maximum drawdown at time T of a random process on $[0, T]$ can be defined informally as the largest drop from a peak to a trough [66]. A sharp contraction may even indicate that an already successful trading system is deteriorating due to changes in the market regime. Often, MDD is a very important risk measure [67]. MDD also has been widely used. A drawdown is the cumulative percentage loss due to successive declines in investment prices. It is collected at non-fixed time intervals and its duration is also a random variable. The maximum reduction that occurs over a fixed investment horizon is a flexible measure that provides a different view of the risk and price flows of such investments. There is a paper that proposes a new strategy for modeling maximum loss that separates duration and severity and suggests a flexible distribution from the extreme value theory of modeled loss [68]. The nature of model-based maximum risk extraction is then analyzed, and a new risk measure is proposed. The authors think this approach is also applicable to a variety of different investment strategies. In Hongzhong Zhang and Olympia Hadjiliadis' paper, they study the probabilistic behavior of two quantities that are closely related to market crashes [69]. The first is the drawdown of an asset, and the second is the time interval between the last reset of the maximum value and the impairment before the drawdown. The first is the fall of an investor's current asset from a historical high to a pre-specified level. This is widely used in the financial risk management literature as an indicator of a market crash. The second is the rate at which investors' capital decreases and therefore measures the speed at which a market crash occurs. They call this the speed of market crashes. In the classification of large financial crashes presented in Giulia Rotundo and Mauro Navarra's article, the bursting of speculative bubbles due to endogenous causes is located in the framework of extreme stock market crashes [70]. They further delve into the analysis by examining the drawdown and maximum drawdown of declines in index prices to further describe the rising component of these selected bubbles. An analysis of decreasing duration, which is central to their estimated risk measures, is also

performed. It can be then seen that MDD is widely used to explore crashes and bubbles in the financial market.

In conclusion, Maximum Drawdown (MDD) is a widely used risk measure in financial markets for assessing the potential downside risk of investment portfolios. Its applications extend to understanding market crashes, the speed at which they occur, and the dynamics of speculative bubbles. Researchers have explored different modeling approaches and have used MDD as a critical tool for risk management and analysis in the financial industry.

2.10 Herding effect

Herd behavior is the behavior of individuals in a group who act collectively without a centralised command. Herd behavior occurs not only in flocks of birds, schools of fish, etc., but also in humans. Voting, demonstrations, general strikes, sporting events, decision-making, and opinion formation are all forms of human-centered herd behavior [71]. Often the same thing happens in financial markets as in our everyday lives. We imitate other people's decisions rather than analysing the available information and making our own decisions. This decision mimicry can lead to collective hysteria and investments can be affected by these panic situations. In financial markets, every investor receives information about how they should act. They are also aware of other people's decisions, even though they do not know the information they receive. Using this information, each investor can make his or her own decisions. But blind faith in other people's decisions can sometimes lead investors to ignore the information they themselves receive and instead make the same decisions as their counterparts. This effect on the economy is known as the "Herding effect". Many believe that this phenomenon has been a major cause of stock market crashes and mass panics throughout history [17, 72].

There is a herding effect in financial markets when a group of investors ignore their own information and only listen to the decisions of other investors. Early studies on herding in finance were rooted in the field of behavioral economics. Researchers like De Long et al. (1990) observed the existence of herding behavior in stock markets, highlighting the disconnect between market prices and fundamental values [73]. There are many scenarios in which herding effects can occur in financial markets. Thomas et al. examine herding behavior in

global markets by applying daily data for 18 countries from May 25, 1988, through April 24, 2009 [74]. They find evidence of herding in advanced stock markets (except the US) and in Asian markets. For example, when investors realise that the government is unable to pay its debts, they convert the government's currency into real assets (such as gold) or foreign currency. This is when a herding effect tends to occur in the currency markets. It tends to cause moderate inflation in the short term. When consumers realise that inflation has increased for the goods they need, they start to stock up. This accelerates the rate of inflation even faster and eventually leads to a collapse of the currency. A similar thing can happen in the stock market [75]. Some studies have shown that herding effects can influence the decisions of normal market participants [76,77]. The emergence of collective decision-making in the market is not necessarily harmful [72]. Indeed, when all agents tend to align their expectations with those of one or a few leaders, the herd effect may greatly reduce market efficiency. However, market dynamics become efficient when each agent considers multiple opinions and thus follows the wisdom of the crowd.

In conclusion, the herding effect represents a fascinating and complex aspect of financial decision-making and market behavior. It is rooted in behavioral biases and has significant implications for market dynamics, asset prices, and investor decision-making. Research continues to explore the causes and consequences of herding behavior and its role in shaping financial markets.

2.11 Order-driven markets

The order-driven market is a financial market. In an order-driven market, trades are executed using trading rules. Orders from both buyers and sellers are displayed in the order book, and trades are only executed when the buy and sell prices match. In an order-driven market, dealers are not involved, instead, traders buy and sell directly from each other. An order-driven market consists of a continuous flow of buy and sell orders from market participants. There is no designated liquidity provider and the two basic types of orders are market orders and limit orders. Market orders are executed at the best possible price in the market. However, there is no guarantee of their execution or price. A limit order sets a maximum purchase limit

or a minimum sale limit. Orders can only be executed below, not above, the buy limit and not below the sell limit. The greatest advantage of an order-driven market is transparency, as the entire order book is displayed to investors who wish to access this information. The order-driven trading system ranks buy and sell orders according to price, matching the highest ranked order (if possible) with the minimum order size. If there are any shares left to buy or sell in a given order, the trading system matches that order with the next highest ranked sell or buy order.

Since 2000, order-driven markets have been widely used in major financial institutions around the world. At the same time, economic physics has become an established area of research for physicists [78,79]. But in a multi-agent framework, the physicist's approach to the microstructure of order-driven markets has always lacked an appropriate utility function [80]. As a result, there is a preference in economics for simple statistical methods to study the problem. Bak, Paczuski and Shubik constructed a simple model of the stock market and argued that large changes in stock prices may be due to herdings, where agents imitate each other's behavior [81]. Different scaling behavior is obtained when agents imitate each other and react to recent market fluctuations. They investigate the interaction between "rational" traders, who act by analysing basic information about stocks, and "noise traders", who make decisions by studying market dynamics and the behavior of other traders only. When the relative number of rational traders is small, bubbles typically occur, where market prices are outside the range justified by fundamental market analysis. When the number of rational traders is high, market prices are generally locked within their defined price range. Maslov introduced and examined a simple model of a limit order-driven market [82]. Traders in this model can either trade the stock at the market price or place a limit order. The choice between these two options is completely random. Numerical simulations of the model show that, despite these minimalist rules, the model generates price patterns with realistic characteristics. This line of research is known as zero intelligence [83].

2.12 Continuous double auctions

A continuous double auction (CDA) is a dynamic market environment in which bidders repeatedly exchange bids to buy and sell units of a commodity in order to maximize trade surpluses [84]. CDA is one of the most common types of auctions and underpins most financial and commodity exchanges. Double means that both buyers and sellers bid, while continuous means that once bidding occurs, the auction will be held for a period of time. The bids are then cleared. After a new sell (buy) bid is received, the auction checks to see if it matches the best outstanding buy (sell) bid. If so, each bidder's transaction price is usually determined by the previous bid, and a new offer is posted. A typical offer consists of a bid and offers pair, where BID is the highest outstanding purchase price and ASK is the lowest outstanding sale price. The standard CDA allows bidding for multiple units of a good, which are usually considered to be divisible quantities. Thus, a bid effectively expresses a willingness to trade any part of a specified quantity at a given price. Unmatched or partially matched bids are listed in the order book and usually remain unpaid until matched, substituted, withdrawn, or canceled at the end of the auction. There are many variations of this overall scheme, with different rules for bidding, clearing and quoting.

The standard model in economics emphasises the role of intelligent agents who maximise utility. However, in some cases, constraints imposed by the market system govern the behavior of strategic agents. Farmer, Patelli and Zovko tested a simple model using data from the London Stock Exchange [83]. In this model, a minimum number of intelligent agents place random orders to trade. The model deals with orders, prices and statistical mechanisms in the context of a continuous double auction and generates simple laws relating order arrival rates to statistical properties of the market. It demonstrates the existence of simple laws linking prices to order flow. And it shows that in some cases the strategic behavior of agents may be governed by other considerations. Order-driven markets are based on the principle of continuous double auctions. It explains the inflow of demand and supply by constantly adjusting the prices of assets. Orders placed by traders are organised and matched in the limit order book (LOB) which ultimately defines the microstructure of the market. Bartolozzi introduced a new multi-

agent model that aims to reproduce the dynamics of the dual auction market at the micro time scale by faithfully modeling the matching mechanism in the limit order book [85]. The agents' actions are related to a random variable, namely market sentiment. He defined this as a mixture of public and private information. Although the model makes only a few basic assumptions about the agents' trading strategies, it is able to reproduce empirical features of the high-frequency dynamics of the market microstructure. Generally, a trade is successful if the highest price offered by the buyer matches the lowest price offered by the seller. In addition, the second indicator of matching orders is the timestamp of each order when the bid/ask prices of the orders are the same. CDA is commonly used by the financial markets for high frequency trading.

2.13 High-frequency trading

High frequency trading refers to trading at very high frequencies, but there is still no single definition. The US Securities and Exchange Commission (SEC) and the US Commodity Futures Trading Commission (CFTC) define a high frequency trader as a proprietary trading firm that uses high speed systems to monitor market data and submit large volumes of orders. The European Securities and Regulation Commission (CESR) considers high frequency trading to be a form of automated trading, using sophisticated computer and IT systems to execute trades at millisecond speeds and hold positions briefly during the day, profiting from trading financial instruments between different trading platforms at ultra-high speeds. High-frequency trading refers to computerised trading that seeks to profit from extremely short-lived market movements that people cannot take advantage of. In recent years, high frequency trading has seen rapid growth. In the United States, high-frequency trading accounts for over 70% of the total equity trading volume. The same trend is evident in Europe, where the number of contracts traded at high frequencies has experienced exponential growth [86].

Despite the minimal profits from a single trade, high-frequency traders make huge profits by using computers to place and withdraw orders up to tens of thousands of times per second, either by pre-empting trades or by programming the logic to be fixed and completing operations as soon as market information enters. High-frequency trading can free up the mar-

ket's hidden liquidity and increase its liquidity. High-frequency traders can play the role of formal or informal market makers placing simultaneous limit orders to buyers and sellers on limit sheets, which provides liquidity to market participants trading at the same time. High-frequency trading results in lower bid/ask spreads in the market. The spread is the difference in price between buying and selling an investment product, which is the cost of trading to the investor. However, for regulators, high-frequency trading is difficult to control effectively and regulation is more difficult.

Due to the boom in high-frequency trading in recent years, the study of it has become a major area of interest for financial practitioners. Brogaard, Hendershott and Riordan examined the role of high frequency traders in price discovery and price efficiency [87]. In general, HFTs promote price efficiency by trading in the direction of permanent price changes and in the opposite direction to transient pricing errors, both at average prices and on the highest volatility days. They argue that the direction of trading in high-frequency trading is related to public information, such as macro news announcements, market-wide price movements and limited order books. High-frequency trading is also widely used in the foreign exchange market. Golub, Dupuis and Olsen outlined key insights from the academic literature on the fundamental mechanisms of the foreign exchange market and the impact of high-frequency traders in the foreign exchange market and discuss actual market events where there are short-term price disruptions [27]. They focus on the behavior of the relevant high-frequency traders. They also propose a new approach to achieving price stability. They suggest that a limit order thin queuing system rewards market participants by offering competitive two-way prices. The results of the simulations suggest that this may well enhance market stability.

Chapter 3

Opinion formation, herding and market stability

3.1 Introduction

Our research in this chapter attempts to study crashes and bubbles by using an agent-based model that reflects traders' (agents') opinions and herding behavior. The agents are located on a social graph and trade only one stock in the simulated market. A single stock assumption limits the liquidity of the markets we model. While liquidity is probably a catalyst in the development of crashes and bubbles, it may not be the nucleating and triggering factor [54]. We aim to find the crashes and bubbles by observing the changes in the price of this stock under different levels of trust of agents [88]. Specifically, we try to explore crashes and bubbles by connecting social graphs, opinion formation, and markets. We set a parameter similar to the degree of stubbornness of agents. A higher level of agent stubbornness means that the agent is more likely to believe in itself than its social connections. By varying the value of this parameter, we observe how the stock price changes for agents with different levels of stubbornness. We use Maximum Drawdown to define crashes and bubbles [66]. Our results allow to relate crashes and bubbles and their duration to different levels of stubbornness in cliques social graph.

We use ABM to connect two complex dynamics in financial markets. The first is concerned

with how traders' decisions are influenced by their professional connections. We adopt a model of opinion dynamics, introduced in [89] building upon the seminal work of DeGroot [45], and studied extensively by the community working on incentives in a multi-agent system, see, e.g., [46, 48]. In this model, each agent (modeling a trader) needs to balance its own personal belief about the future movement of a financial asset with the opinions of the other traders in the market. The interplay between steering (acting upon personal belief) and herding (acting upon the wisdom of the crowd) is in fact the second dynamic we want to study. To capture this, we link the opinion dynamics with trading decisions, which naturally define market price dynamics. In this setup, we study the emergence of herding in financial markets, as a function of the structure of the social network of traders.

We will use empirical game-theoretic analysis to validate our results and explore some new relationships between stubbornness and crashes and bubbles. We compute the equilibria of the game, where agents decide at each step whether to follow or steer, for a variety of social graphs. Once we reach a stubbornness equilibrium, we observe how the stock price changes by using our ABM. We further assess how herding (or lack thereof) affects markets by studying crashes and bubbles. Our results overall confirm that at equilibrium, the more traders are connected in the social network the more likely it is that herding will occur. However, herding will not cause substantial price movements in the market unless the social graph is a clique. In this case, in fact, we find that a minority of agents will maintain a high level of stubbornness at equilibrium, which will lead to a few bumps in the price time series.

An idealised market environment is simulated in our experiments. In such a market, all participants have access to perfect information and are rational in their decision making, seeking to maximise their individual utility. Although real markets are often more complex and vulnerable to imperfections and external factors that can deviate from the idealised market in various ways. But idealised markets can be a useful benchmark for analysing real markets. Understanding the differences between ideal and real markets can help promote market efficiency. In Chapter 4 and 5 we will simulate a more realistic market.


3.2 An agent-based model (ABM) of Opinion Formation

We try to make some adjustments based on a paper that we mention in Section 2.7 [54]. In our model, public news is included in private information. Moreover, in our experiment, we use a fully connected clique. More neighbors of the agent can make the agent's opinion formation more realistic and allow a more complete study of herd behavior. We also used a new agent-based model of opinion formation inspired by [48].

We set up a fixed number of N agents who are buying and selling a stock. As mentioned in [54], this can be thought of as a market portfolio, trading in an organized market coordinated by market makers. At each time step, the agent may trade or hold. We focus on understanding the detailed issues behind bubbles and crashes.

The information that underlies an agent's decision is limited to two sources (neighbors and themselves) and the returns that have been realized. Our agent will use all available information to maximize profits.

Figure 3.1: $G = (N, E)$



In our experiments, the decisions made by agents at each time step depend on their opinion at the current time step. We denote the opinion for agent i at time step t , as $s_i(t) \in [0, 1]$. The graph $G = (V, E)$, $V = N$, is undirected with $E = \{(i, j) | i, j \in N\}$. Figure 1 represents a clique where $E = V \times V$. N_i as the neighborhood of i in G , i.e., $N_i = \{j \in V | (i, j) \in E\}$. For cliques, we have $N_i = V \setminus \{i\}$. We assume that agents are aware of each other's decisions at the last time step.

3.2.1 Two sources of information for opinions

At each time step, the agents examine and weigh their available information (including public news) to form their own private opinion on what the future price variation will be. Based on their private opinion, they decide whether to trade or remain inactive. The information they use has two different origins.

Private information. The first source of information is private information denoted by $b_i(t)$, which is information obtained by agent i at time t in its own unique ways, but which is not public. It reflects each agent's particular subjective view of the stock's future performance. Intuitively, $b_i(t)$ can be seen as a belief of agent i about the stock; the bigger $b_i(t)$ the more the agent believes that the stock will increase in value (and then need to be bought). The private information $b_i(t)$ is different for each agent and there is no correlation between each agent's private information; moreover each $b_i(t)$ changes with the time step t .

Neighbor information. The second source of information is provided by the past decisions of other agents at time t , $s_j(t)$, for $j \in N$ and $t \leq T$, we let T denote the trading horizon. With limited access to information and limited bounded rationality, some argue that imitating others is optimal [90]. In our model, motivated by work in game theory, agents gather information on the opinions of their "neighbors" in their social network and incorporate it as an ingredient into their trading decisions.

3.2.2 Opinion formation

The cost of the opinion $s_i(t)$ for agent i at time step t is denoted by $C(s_i(t))$ [48]. It represents the degree of match between the opinion of the agent i and the information it receives. The smaller the gap between its opinion and the two sources of information, the lower the cost to

the agent.


$$C(s_i(t)) = (1 - \alpha)[s_i(t) - b_i(t)]^2 + \alpha \frac{\sum_{[(i,j) \in E, i \neq j]}^{N-1} [s_i(t) - s_j(t-1)]^2}{|N_i|} \quad (3.1)$$

As we can see in (3.1), $[s_i(t) - b_i(t)]$ means the difference between the opinion of agent i at time t , $(s_i(t))$ and the private information it got $(b_i(t))$; α denotes that how much agents trust their neighbors, $\alpha \in [0, 1]$. A bigger α represents more trusting other agents' opinions, and vice versa. In this experiment, we limited each agent to only two sources of information (private information and neighbor information) in the market. So $(1 - \alpha)$ means the agent's level of trust in itself; this is what we call stubbornness. The value of $s_i(t)$ at the minimum of $C(s_i(t))$ is the optimal opinion chosen by our agent i at time t . In this sense, traders are rational as they minimize their costs.

3.2.3 Trading decision

Buy or sell. we now relate the opinions to the trading decisions in terms of decision and size. The signal for the trading is determined by the opinion $s_i(t)$ for agent i at time t . We set a buying threshold th_B and a selling threshold th_S , both in $[0, 1]$, with $th_B > th_S$. They can be seen as the agent's risk tolerance in each experiment. When $s_i(t) \geq th_B$, the agent i will buy the stock. When $s_i(t) \leq th_S$, the agent i will sell the stock. When $th_S < s_i(t) < th_B$, the agent i will hold, as shown in Figure 2.

Figure 3.2: Trading decision



Trading size. The assets of agent i consist of the amount of cash it holds, $cash_i(t)$, and the number of stocks traded in the market, $stock_i(t)$; $price(t)$ is the price of the stock of time t .

When an agent decides to buy, it uses a fixed fraction g of its cash to buy stocks. When an agent decides to sell the stock, it sells a fixed fraction g of its stock. Therefore, its action $A_i(t)$ and the direction $B_i(t)$ of its decision is made by:

If $s_i(t) \geq th_B$

$$A_i(t) = g \cdot \frac{cash_i(t)}{price(t)} \quad (3.2)$$

$$B_i(t) = +1(\text{buying}) \quad (3.3)$$

If $s_i(t) \leq th_S$

$$A_i(t) = g \cdot stock_i(t) \quad (3.4)$$

$$B_i(t) = -1(\text{selling}) \quad (3.5)$$

3.2.4 Price clearing

When all the agents have completed the transaction, a new price is created. In the following formula (3.6), $r(t)$ is the average of the stock return; γ represents the relative impact of the excess demand upon the price. We assume for simplicity the existence of a market maker, who accepts all transactions from agents and has an unlimited amount of cash and stock. The assumption makes sense for markets with many high frequency traders, providing liquidity to the market. The price is determined by:

$$r(t) = \frac{1}{\gamma \cdot N} \sum_{i=1}^N A_i(t) \cdot B_i(t) \quad (3.6)$$

$$\log[price(t+1)] = \log[price(t)] + r(t) \quad (3.7)$$

In (3.6), we assume that the linear market impact function is a rough approximation of a time scale that is significantly larger than the trade-by-trade time scale on which the nonlinear impact function is observed [91]. This is similar to [54].

3.2.5 Cash and stock

We assume a frictionless market that has no transaction costs. When the return and new price of the stock are determined by (3.6) and (3.7), the amount of cash and stock held by agent i will be determined by the following formulas:

$$cash_i(t) = cash_i(t-1) - A_i(t) \cdot price(t) \cdot B_i(t) \quad (3.8)$$

$$stock_i(t) = stock_i(t-1) + A_i(t) \cdot B_i(t) \quad (3.9)$$

3.2.6 Crashes and bubbles

To the best of our knowledge, there are no accepted quantitative definitions of these concepts. The definitions of crashes and bubbles in our research are based on changes in stock prices inspired by the financial concept of Maximum Drawdown. We let P_{max} and P_{min} denote maximum and minimum stock prices in a suitable time interval, they are initially set at $price(0)$. We let P_t denote a shorthand for the current price at time t , $price(t)$. We assume that there are two thresholds δ_{in} and δ_{out} . These two thresholds allow to study different market conditions. When the stock price trends downwards, δ_{in} is the percentage decline in stock price, used to determine the starting point of the crash and δ_{out} is the growth after the stock price reaches its lowest point, used to determine the end point of the crash; δ_{in} and δ_{out} are similarly used to determine where bubbles start and end when stock prices tend to rise. The point-in-time at which the crashes and bubbles begin is named t_{in} , and the point at the end of the crashes and bubbles is named t_{out} .

We let t_{max} denote the time step in which the price was P_{max} and t_{min} be the time step when the price was P_{min} . Our notion of crash is given in Algorithm 1. The duration of this crash is from t_{in} to t_{out} . When $\delta_1 > \delta_{in}$ and $\delta_2 > \delta_{out}$, there is a crash, see Figure 3.3. Bubbles are defined similarly in Algorithm 2. When $\delta_3 > \delta_{in}$ and $\delta_4 > \delta_{out}$, there is a bubble, see Figure 3.3.

In the other case, as shown in Figure 3.4, although $\delta_1 > \delta_{in}$ and $\delta_3 > \delta_{in}$, $\delta_2 < \delta_{out}$ and $\delta_4 <$

Algorithm 1 Definition of crash

Let $\delta_1 = \frac{P_{max} - P_{min}}{P_{max}}$, $\delta_2 = \frac{P_t - P_{min}}{P_{min}}$

```

while  $\delta_1 > \delta_{in}$  do
  |  $t_{in} \leftarrow t_{max}$ 
  | if  $\delta_2 > \delta_{out}$  then
  | |  $t_{out} \leftarrow t_{min}$ 
  | end

```

end

Crash in $[t_{in}, t_{out}]$

Set $P_{max} \leftarrow P_{min}$

Algorithm 2 Definition of bubble

Let $\delta_3 = \frac{P_{max} - P_{min}}{P_{min}}$, $\delta_4 = \frac{P_{max} - P_t}{P_{max}}$

```

while  $\delta_3 > \delta_{in}$  do
  |  $t_{in} \leftarrow t_{min}$ 
  | if  $\delta_4 > \delta_{out}$  then
  | |  $t_{out} \leftarrow t_{max}$ 
  | end

```

end

Bubble in $[t_{in}, t_{out}]$

Set $P_{min} \leftarrow P_{max}$

Figure 3.3: Definition of crash and bubble

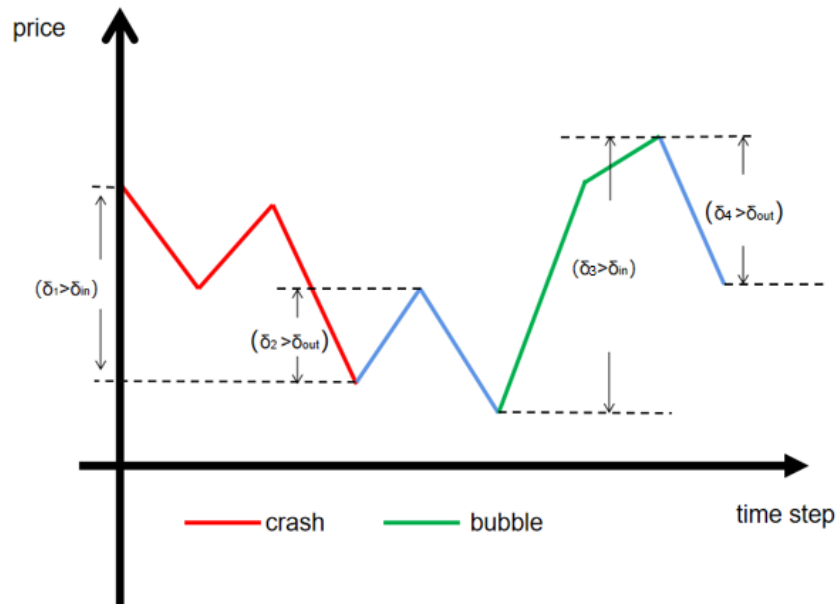



Figure 3.4: Another situation



δ_{out} , there are no crashes or bubbles. In our experiment, δ_1 , δ_2 , δ_3 and δ_4 are the proportion of changes in stock prices.

3.2.7 Set up parameters

Our agents will invest a certain proportion of their cash in the only stock traded in the market. Thus, when most agents buy the stock with their money, the local maximum value of the stock's price is reached. Until there are few agents to buy the stock, the price of the stock will peak, after which its price will fall.

Basic parameters. Our experiment is based on [54], in order to observe the difference and connection between the two sets of experimental results. The graph $G = (V, E)$ is a clique. We let T denote the trading horizon and set $T = 10,000$ in our experiments, this is the same as Georges's experiment. In our simulations, we fix the number of agents in the system to $N = 100$, $\gamma = 0.25$, the fraction of their cash or stocks that investors trade per action to $g = 2\%$. This means that each agent has 99 neighbors, which is enough to support us through the later experiments. The initial amount of cash and stocks held by each agent to $cash_i(0) = 1$ and

$stock_i(0) = 1$; $price(0) = 1$. For the value of the private information, $b_i(t)$, we use random values from 0 to 1 because the private information changes with every time step and each agent has different personal information. However, this does not mean that the agents change their minds at every time step. The agents just get different information about the stock at every time step.

Trading threshold. Because the agents in our experiment have only three trading strategies (buy, sell and hold), we set $th_B = \frac{2}{3}$, $th_S = \frac{1}{3}$. That means when the opinion for agent i at time step t , $s_i(t) > \frac{2}{3}$, the agent i will use its cash at the rate of g to buy the stock. When $s_i(t) < \frac{1}{3}$, agent i will sell g of its stock. When the value of $s_i(t)$ is between $\frac{1}{3}$ and $\frac{2}{3}$, the agent will not trade, it will keep its cash and portfolio unchanged.

δ_{in} and δ_{out} . As mentioned above what we try to find is the duration of crashes and bubbles that occur in the market. In fact, crashes can occur with small changes in stock prices [92]. As the value of δ_{in} and δ_{out} become smaller, the probability of crashes and bubbles will become higher. If we try to change the value of crashes and bubbles, we will get results under different market environments. In our experiments, we try to explore the impact of agents with different levels of stubbornness on market stability. In order to better observe the results of the experiment, we do not need large changes in stock prices to define crashes and bubbles. As from Section 3.2.6, δ_{in} and δ_{out} determine the market crash and bubble's in point and out point. In our experiment we set $\delta_{in} = 5\%$, $\delta_{out} = 2.5\%$ [65, 93].

3.2.8 Results

We did this experiment by picking different values of α . For each different α , we run 100 simulations and then we take the average of the prices at each time step. To make the experiment fairer and more convincing, we used the same set of data of $b_i(t)$ for each different value of α .

In Table 3.1 and 3.2, we can clearly see the number and timing of crashes and bubbles at different values of α . When $\alpha = 0$, the stubbornness of agent i reaches its maximum. That means agent i only trust itself, $s_i(t)$ depends entirely on its private information ($b_i(t)$). When $\alpha > 0.7$, few agents trade in the market. The price of the stock remains almost unchanged.

Table 3.1: Results

Crashes and bubbles in different α		
Value of α	Number of crashes	Number of bubbles
0.1	0	0
0.2	2	0
0.3	4	1
0.4	5	1
0.5	5	1
0.6	8	7
0.65	8	9
0.7	1	0

Table 3.2: Duration steps of each Crash and Bubble

Value of α	Duration of crashes	Duration of bubbles
0.1	none	none
0.2	3127, 3894	none
0.3	2, 2481, 4512, 2045	90
0.4	3, 1913, 1907, 2915, 1817	169
0.5	4, 1972, 2072, 2708, 1490	129
0.6	5, 855, 679, 538, 1126, 906, 1253, 535	149, 96, 75, 383, 548, 164, 500
0.65	7, 256, 372, 344, 327, 620, 297, 129	712, 390, 670, 730, 212, 298, 592, 427, 403
0.7	10	none

We find that four states emerge, depending on how much the agent is influenced on average by its neighbors (controlled by α). With small α systems (low herding regime), agents are sometimes influenced by the news and sometimes by their neighbors. But since by default they have little or no trust in their neighbors, they do not over-imitate them. Agents' decisions are largely determined by their private information. As shown in Figure 3.5, the stock price is relatively stable and tends to fall.

By increasing α above a certain threshold, the system enters a second state in which agents on average pay more attention to the information of their neighbors than their private information. Because agents are always trying to better match their decisions with the information they receive, it makes sense for agents to follow the majority and "surf" during bubbles or crashes. This process is characterized by the fact that agents' strategies are locally adjusted to market sentiment and therefore fluctuate greatly from the underlying price. As shown in Figure 3.6, the red part denotes crashes and the green part denotes bubbles, it is the same for the following figures. We can see there are 5 crashes and 1 bubble occurring under this regime, but stock prices still tend to decline.

As we continue to increase the value of α , the system will enter a third state. In this regime, the agent pays more attention to information about its neighbors. As a result, the underlying price of the stock is still highly volatile and there are many crashes and bubbles. However, as shown in Figure 3.7, the stock price is rising.

It is not until the value of α increases to another threshold that the system enters the fourth state. In this state, agents are reluctant to trust private information about themselves. Instead, they are more willing to trust their neighbors' information. As noted above, the agents' decisions at the beginning of the experiment depended heavily on their private information. As the time steps increases, few agents are inclined to trade, and they are more likely to hold their stocks. As a result, as shown in Figure 3.8, stock prices fluctuate little and trend upward.

In general, our model is like an efficient market, except that during bubbles and crashes these huge price changes occur. As the value of α increases, the stubbornness of agents reduces and the likelihood of crashes and bubbles increases. This is similar to the conclusion of [54]. The difference is that when the value of α reaches a certain threshold, there are very






Figure 3.5: Price time series when $\alpha = 0.1$ Figure 3.6: Price time series when $\alpha = 0.6$ 

Figure 3.7: Price time series when $\alpha = 0.65$ Figure 3.8: Price time series when $\alpha = 0.7$ 

few agents in the market who will trade and the stock prices will not change much. In the experiment, we have another interesting discovery. The price of stocks in the market will always tend to decline until the value of α is bigger than a certain threshold. At this point, the price of the stock will tend to increase.

3.2.9 The initial flash crash

We can see there is a flash crash in the first few time steps, as is shown in Figure 3.9. One of our conjectures is that the large difference from the base price at the beginning of the experiment stems from private information. A similar set of private information (including news) appears randomly in order to push prices in one direction and begin the process of synchronization of agents. This is realistic because economic news tends to be consistent. As a result, traders are either confident (a string of good news) or overly pessimistic (a string of bad news) about the future of the economy and they will push stock prices in one direction. Once the market reacts to a sharp drop in stock prices and they trust their neighbors' information, they will mimic their neighbors' decisions and produce the same behavior as in our model.



Through our experiments, we find that the reason for the flash crash is the price of the stock is actually calculated at the first time step. In Figure 9, our initial price is 1 ($price(0) = 1$). When $t = 0$, the opinion of agent i is only decided by $b_i(t)$ as shown by formula (1). We use the same set of $b_i(t)$. In this set of data, there are a lot of $b_i(t)$ that are smaller than th_S in the first few time steps. That means that a lot of agents will choose to sell the stock at the beginning of the simulation. This leads to a flash crash at the beginning. As agents in the market continue to trade, the market quickly returns to normal. This also reflects the self-regulating function of our simulated markets.

As shown in Figure 3.10, in our experiment, the average of the first five time steps of $b_i(t)$ is 0.31. That is smaller than th_S . The average of each time step of $b_i(t)$ is 0.47 as shown in Figure 3.11. So we try to use a set of private information where the average of $b_i(t)$ for the first few time steps is less one-sided as shown in Figure 3.12. The average of the first five time steps of $b_i(t)$ is 0.5. The average of each time step of $b_i(t)$ is also 0.5.

Figure 3.10: First 100 time steps of $b_i(t)$

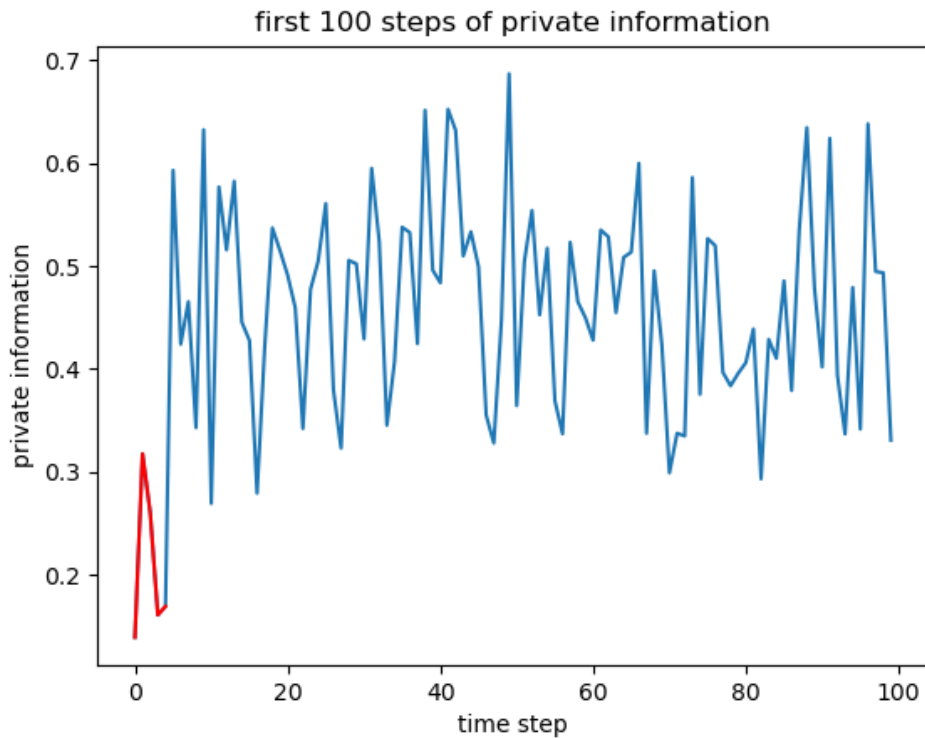



Figure 3.10 shows the average of the first 100 time steps of private information, $b_i(t)$. The red part is the average of the first 5 time steps of $b_i(t)$.





We rerun the experiment with a new set of private information at $\alpha = 0.5$, see Figure 3.12. The result of the experiment is as expected. The flash crash at the beginning of the experiment disappeared, see Figure 3.13:

Figure 3.13 shows the price change of the stock when $\alpha = 0.5$ by using the new set of $b_i(t)$. There are 4 times crashes and no bubble. As we can see in Figure 3.11 and Figure 3.12, in both the original and new set of private information, the value of $b_i(t)$ is more clustered in $[0, 0.3]$. This could be the reason for the overall downward trend in stock prices. It is only when the value of α reaches a certain threshold that the herding of agents comes into play leading to an upward trend in stock prices even though everyone is pessimistic.

3.2.10 Different opinion formation rules


In (3.1), the denominator is $|N_i|$. We think that this denominator is going to cause opinions that are not too linked to our graph since we take the average of the difference between agent i and its neighbors. So we try to change our opinion formation to :

Figure 3.12: Distribution of new set of $b_i(t)$'sFigure 3.13: $\alpha = 0.5$ with new $b_i(t)$ 

$$C(s_i(t)) = (1 - \alpha)[s_i(t) - b_i(t)]^2 + \alpha \sum_{[(i,j) \in E, i \neq j]}^N [s_i(t) - s_j(t-1)]^2 \quad (3.10)$$

The results show that agents in the market only trade when the value of α is very small. As we discussed earlier, the clique is too strong for (3.10) and herding occurs for very small values of α . This is not an effective opinion formation rule. As shown in Figure 3.14, when $\alpha = 0.021$ the stock price fluctuations reach their maximum. At this point, we get 13 crashes and 13 bubbles. The market never moves significantly when α is more than 0.27.

Figure 3.14: Price time series when $\alpha = 0.021$



3.3 Empirical Game Theory Analysis (EGTA)

In the real market, traders cannot all have the same level of stubbornness in every trade as we set out above. Every trader should have their own choice. For simplicity in this chapter, we will consider a small game to see the extent to which incentives change our result. In Chapter

5, we look at a more general setup. In our game, we start by identifying three strategies and find a Nash equilibrium in the game by calculating the average payoff of the agents under each strategy. Crashes and bubbles occurring when the market is in equilibrium are then observed based on the number of agents choosing each strategy.

3.3.1 An example of Nash equilibrium calculation

Calculating a Nash equilibrium using the Empirical Game-Theoretic Analysis (EGTA) method is a data-driven process that involves several steps. Let's go through a step-by-step example using a simplified scenario:

Step 1: Define the Game. Let's consider a version of the classic "Prisoner's Dilemma" game. There are two players, Player A and Player B. They each have two strategies: "Cooperate" (C) or "Defect" (D).

Step 2: Parameterized strategic space. Define the strategy space for Player A and Player B. In our example, they have two strategies each: "Cooperate" (C) and "Defect" (D).

Step 3: Generate Data. We need to collect real-world data or data from simulations that reflect how Player A and Player B make decisions in the game. This can be based on observation or experimentation.

Step 4: Estimate Empirical Game. Use the collected data to estimate empirical payoffs for each strategy profile. The empirical payoffs are based on the frequencies of observed outcomes. The payoff matrix for this game is shown in Table 3.3

Table 3.3: The payoff matrix for this game

Player B Player A	Cooperate (C)	Defect (D)
Cooperate (C)	(3 , 3)	(0 , 5)
Defect (D)	(5 , 0)	(1 , 1)

For Player A choosing "Cooperate" and Player B choosing "Cooperate", we assume that in three simulations of the strategy profile (C, C) we get payoffs (2, 2) (3, 3) and (4, 4). Then

the empirical payoff will be the average A: 3, B: 3 (3, 3). Similarly, for Player A choosing "Cooperate" and Player B choosing "Defect", the empirical payoffs are A: 0, B: 5 (0, 5).

Step 5: Use a solver to compute the equilibrium. Analyze the empirical game to identify Nash equilibrium. In our example, you can calculate the following:

For Player A: The best response to Player B choosing "Cooperate" is "Defect" (D). The best response to Player B choosing "Defect" is also "Defect" (D).

For Player B: The best response to Player A choosing "Cooperate" is "Defect" (D). The best response to Player A choosing "Defect" is "Defect" (D).

Given the analysis, the strategy profile (D, D) is a Nash equilibrium since neither player has an incentive to unilaterally change their strategy.

3.3.2 Three Strategies

In our experiments, each agent has three strategies to choose from. They are "Imitation", "Neutral" and "No imitation". As mentioned above, a higher level of agent stubbornness means that the agent is more likely to believe in itself. This means that we can define these three strategies in terms of the value of α . For example, we use $\alpha = 0$ to represent the "No imitation" strategy. Then we choose a higher value of α to represent the "Imitation" strategy. The value of α for the "Neutral" strategy can be chosen within the interval of the α values of the two strategies mentioned above.

3.3.3 Payoffs

As we mentioned above, the amount of cash and stock held by agent i will be determined by 3.8 and 3.9. So the total assets of agent i at time step t will be determined by the following formula:

$$Z_i(t) = \text{cash}_i(t) + \text{stock}_i(t) \cdot \text{price}(t) \quad (3.11)$$

Then the total payoff of agent i after t time steps will be determined by:

$$\begin{aligned}
R_i(t) &= (Z_i(1) - Z_i(0)) + (Z_i(2) - Z_i(1)) + \cdots + (Z_i(t) - Z_i(t-1)) \\
&= Z_i(t) - Z_i(0)
\end{aligned} \tag{3.12}$$

3.3.4 Basic experiment setup

Our experiments were carried out on the basis of all the above settings. We went 10,000 time steps per experiment which means that $T = 10,000$. In our experiments, each agent can choose one of the three strategies per experiment. At the end of each experiment, according to 3.11 and 3.12, we record the number of agents who chose each strategy and their average total payoff, respectively. By varying the number of agents for each strategy, we obtained each of the possible scenarios and their average total return. We then use EGTA to process these data analytically and find the Nash equilibrium. We will then observe the crashes and bubbles that occur in markets at these Nash equilibrium points.

3.3.5 Choice of the value of α


Based on the results of Table 3.1, it is easy to see that when the value of α exceeds 0.7, the price of the stock only fluctuates very little and only very few crashes or bubbles occur. After several tests, we have determined that the price of a stock will no longer change when $\alpha = 0.75$. So we set $\alpha = 0.75$ as one of the strategies, called “Imitation”. Similarly, the results of Table 3.1 show that when the value of α exceeds 0.5, there are significantly more crashes and bubbles in the market and the fluctuation of price becomes greater. Therefore we set $\alpha = 0.5$ as the second strategy, called “Neutral”. The last strategy is “No imitation” and it is represented by $\alpha = 0$.

3.3.6 Results of EGTA

To make our results more robust, we run 100 simulations for each different profile and get the average payoff of the agents choosing each strategy. In each of these 100 simulations, we used completely different sets of private information ($b_i(t)$), and for each profile, we used the same 100 sets of $b_i(t)$. Figure 3.15 shows the Nash equilibrium that we obtained. The triangle represents the strategy space in two dimensions. The points in the triangle represent the unit

simplex strategy. The vertices represent the pure strategies. The open circle in the triangle represents a Nash equilibrium. Each line represents a trajectory that starts from an Original strategy and then applies the function repeatedly until it reaches an equilibrium. By looking at the dynamic trajectory of the agents' choice of strategy in the figure we can find two Nash equilibrium points. At the first point we have 92 agents choosing "Imitation" strategy, 0 agents choosing "Neutral" strategy, and 8 agents choosing "No imitation" strategy. And at the second point we have 93 agents choosing "Imitation" strategy, 7 agents choose "Neutral" strategy and 0 agent choose "No imitation" strategy.

Figure 3.15: Nash Equilibrium




With these two equilibrium points, we obtain the results shown in Figure 3.16 and 3.17.

There are 3 crashes and 4 bubbles at the first Nash equilibrium, 1 crashes, and 3 bubbles at the second Nash equilibrium.

3.4 Different social connections between agents


In the above experiment, we set up 100 agents that are connected to each other. This means that each agent has 99 neighbors. To round out our experiment, we will try to use different social graphs. We can change the way agents contact each other so that the information they get about their neighbors changes. We obtain the Nash equilibrium points as shown in Figure 3.18:

Figure 3.16: Crashes and Bubbles at the first Nash equilibrium



In Figure 3.18, ‘ x -a-group’ represents a graph where each group has x agents and each agent has $x - 1$ neighbors. The graph is a ring where agent i has an edge with the nearest $x - 1$ agents. For example, the “10-a-group” as it appears in Figure 3.18 represents that each


Figure 3.17: Crashes and Bubbles at the second Nash equilibrium



agent knows the opinion of 9 other agents. In other words, each agent has 9 neighbors, and so on. As shown in Figure 3.18, each social connection has multiple Nash equilibriums. To better study the extent to which social connections between market participants affect market stability, we will choose the point with the lowest price volatility for each social connection as the Nash equilibrium point for this social connection. In Chapter 4 and 5 we do not need to choose the Nash equilibrium because the Nash equilibrium of different social connections is unique. We use the average standard deviation of prices to express price volatility (we run 100 times for each Nash equilibrium). As shown in Figure 3.15, there are two Nash equilibriums at this social connection. The price volatility of these two Nash equilibrium points are 11% and 13% respectively, and we choose the point where the price volatility is 11% as the Nash equilibrium point of this social connection. In the same way, the Nash equilibrium of each social connection is shown in Table 3.4.

As shown in Table 3.4, as agents get more of their neighbors' opinions more people choose the "Imitation" strategy. This is quite understandable, traders are easily influenced in their

Figure 3.18: Nash equilibria for different graphs



judgment by outside information in real life. This means that in the real market, the more information traders have about others, the more they will want to imitate them. In Table 3.4 we can see that at “90-a-group” no agent chose the “No imitation” strategy, which just proves that agents are also becoming more rational in our experiments. At “100-a-group” we find 8 agents choosing the “No imitation” strategy. That means when each agent is given almost the same information, there will always be a few agents who believe that imitation does not yield higher returns. So they choose to trust their own private information.

The data in Table 3.5 shows that the more neighbors agents have the fewer crashes and bubbles there will be in the market. In particular, there is no crash or bubble at “90-a-group”. This result indicates that when most people in the market are willing to imitate others and no one believes their own private information, the market will reach a relatively stable state. At “100-a-group”, as we mentioned above, each agent is aware of all the other agents’ information. In our experiments, agents who choose the “No imitation” strategy means they only trust their own private information, $b_i(t)$. The value of $b_i(t)$ is random from 0 to 1. In other words, at “100-a-group” when a few agents who choose the “No imitation” strategy make some random decisions other agents are trying to imitate their decisions. This is why there were 3 crashes and 4 bubbles in the market at “100-a-group”.


When we relate Table 3.4 to 3.5, it is easy to see that as agents are given more information, more agents choose the “Imitation” strategy. Crashes and bubbles that occur in the market also decrease as agents are given more information. This means that as agents are given more information, the more willing they are to imitate others and the less volatile the price of the stock. In general, the less neighbors’ information the agents have in our experiments the more unstable the market becomes. But when every agent in the market knows all the other agents’ information, there will still be crashes and bubbles in the market. The market reaches its most stable state at “90-a-group”.

Figure 3.19 shows the percentage of agents who chose each strategy at ‘10-a-group’ and ‘20-a-group’. Based on the information in the graph we can see that the largest number of agents chose the “No imitation” strategy, with very few agents choosing the other two strategies at ‘10-a-group’. And at ‘20-a-group’, the number of agents choosing the “Imitation”

Table 3.4: Number of agents at Nash equilibrium


	Agents of Imitation	Agents of Neutral	Agents of No imitation
10-a-group	4	4	92
20-a-group	36	5	59
30-a-group	92	4	4
40-a-group	92	0	8
50-a-group	94	0	6
60-a-group	94	0	6
70-a-group	93	0	7
80-a-group	94	0	6
90-a-group	98	2	0
100-a-group	92	0	8

Figure 3.19: 10-a-group and 20-a-group




strategy increases significantly.

Figure 3.20: 30-a-group and 40-a-group



At '30-a-group', as shown in Figure 3.20, the agents choosing the "No imitation" strategy become very few, while those choosing the "Imitation" strategy become the most numerous. Very few agents choose the "Neutral" strategy at '40-a-group'.

Figure 3.21: 50-a-group and 80-a-group



As we can see in Figure 3.21, the proportion of agents choosing each strategy remains almost unchanged as our social graph gradually changes from '50-a-group' to '80-a-group'. The majority of agents still choose the "Imitation" strategy and we hardly see any agents choosing the "Neutral" strategy.




Figure 3.22 shows the agent's choice of strategy at '90-a-group' and '100-a-group'. In Figure 3.22 we can see that at '90-a-group' almost all agents have opted for the "Imitation" strategy and no agent chose the "No imitation" strategy. At '100-a-group' we can find several agents chose the "No imitation" strategy. That means when each agent is given almost the same information, there will always be a few agents who believe that imitation does not yield higher returns. So they choose to trust their own private information.

According to the information given in Figures 3.19 to Figure 3.22, we can find that as agents get more information from neighbors the more people choose the "Imitation" strategy. This is quite understandable, traders are easily influenced in their judgment by outside information in real life. This means that in the real market, the more information traders have about others, the more they will want to imitate them.

Based on the data in Table 3.4, we record the number of crashes and bubbles that have occurred in the market in Table 3.5:

Based on the data at Nash equilibrium of different social graphs, we record the number of crashes and bubbles that have occurred in the market in Table 3.5. The data in Table 3.5 shows that the more neighbor information agents get the fewer crashes and bubbles there will be in the market. In particular, there is no crash or bubble at '90-a-group'. This result indicates that when most people in the market are willing to imitate others and no one believes their own private information, the market will reach a relatively stable state. At '100-a-group',

Table 3.5: Crashes and Bubbles at equilibrium in different graph

Different graph	Number of crashes	Number of bubbles
10-a-group	13	9
20-a-group	6	3
30-a-group	2	3
40-a-group	1	5
50-a-group	2	3
60-a-group	1	3
70-a-group	0	3
80-a-group	0	2
90-a-group	0	0
100-a-group	3	4

as we mentioned above, each agent is aware of all the other agents' information. In our experiments, agents who choose the "No imitation" strategy means they only trust their own private information, $b_i(t)$. The value of $b_i(t)$ is random from 0 to 1. In other words, at '100-a-group' when a few agents who choose the "No imitation" strategy make some random decisions other agents are trying to imitate their decisions. This is why there were 3 crashes and 4 bubbles that occurred in the market at '100-a-group'.

As for the singularity of the '90-a-group', as we stated above (Figure 3.22) when each agent is essentially endowed with progressively more information from '30-a-group' to '80-a-group' as shown in Figure 3.19 to Figure 3.21, there will always be a few agents who believe that imitation will not yield higher returns and choose to trust their own private information. However, in the case of '90-a-group', the opinions of other traders are not ignored; we can speculate that this is because the herding is maximized for this setting. Basically, all agents in the market choose the "Imitation" strategy, but few prefer to keep a positive weight for their belief. This changes in the '100-a-group', where every agent is given the information of the other 99 agents. In this full information setup, some agents remain stubborn.

Crashes and bubbles that occur in the market also decrease as agents are given more information. This means that as agents are given more information, the more willing they are to imitate others and the less volatile the price of the stock. In general, the less neighbors' information the agents have in our experiments the more unstable the market becomes. But when every agent in the market knows all the other agents' information, there will still be crashes and bubbles in the market. The market reaches its most stable state at '90-a-group'.

3.5 Changing risk attitude and trading size

3.5.1 About the risk attitude

As from Section 3.2.3, when $th_S < s_i(t) < th_B$, the agent i will hold. We can think of $(th_B - th_S)$ as the 'risk aversion' of the agent. When the value of $(th_B - th_S)$ is bigger, it means the agents are more likely to want to avoid the risk and vice versa. We set $th_B = \frac{2}{3}$ and $th_S = \frac{1}{3}$ in the above experiment. To explore what effect different risk attitudes have on the results of our experiments, we try to move the threshold of trading to $th_B = \frac{3}{5}$, $th_S = \frac{2}{5}$, thus capturing risk seeking traders. This also allows to check the robustness of our ABM; the riskier the traders the more volatile the market should be.

No matter the social graphs we get the same result at equilibrium that there are 4% agents choosing the "Imitation" strategy, 4% agents choosing the "Neutral" strategy, and 92% agents choosing the "No imitation" strategy. This means that when the trading risk is high enough, most agents will not want to imitate their neighbors no matter how much neighbors' information they are given. The number of crashes and bubbles that occur in the market at any given social graph remains virtually unchanged. This is because when the vast majority of agents choose the "No imitation" strategy, their decisions depend on their private information, $b_i(t)$. We used the same $b_i(t)$'s for each different social graph in our experiments. This resulted in the number of crashes and bubbles being similar in each of the social graphs. We got 8 crashes and 4 bubbles in this experiment, which is higher than the data in Table 3.5. This confirms that high risk makes markets unstable.

Then we change the threshold of trading to $th_B = \frac{4}{5}$, $th_S = \frac{1}{5}$. This means that trading in the market has become less risky. The results of the experiment were very similar to the

results in Section 3.4. They have the same pattern which is at ‘10-a-group’ and ‘20-a-group’, the majority of agents in the market choose the “No imitation” strategy and from ‘30-a-group’ to ‘80-a-group’ The choice of agents hardly varies much. In this state, approximately 94% of agents choose the “Imitation” strategy and 6% agents choose the “No imitation” strategy. The result of ‘90-a-group’ and ‘100-a-group’ is similar to Figure 3.22. At ‘90-a-group’ almost all agents have opted for the “Imitation” strategy and no agent chose the “No imitation” strategy and several agents chose the “No imitation” strategy at ‘100-a-group’. From these results, we can see that, unlike in the high-risk state, the pattern of agents’ choice decisions in the low-risk market is very similar to that in Section 3.4, especially the singularity of the ‘90-a-group’.

3.5.2 About the trading size

In our experiment, the fraction of cash or stocks that investors trade per action g as mentioned in Section 3.2.3 is the same for all agents following [54] and $g = 2\%$. When we change g to 1%, crashes and bubbles in the market have become less frequent. When we increase g to 5%, the number of crashes and bubbles will increase with g . Then we randomly choose g from 1% to 5% for each agent of different social graphs and get the Nash equilibrium. The results are nearly the same as the Figure 3.18. That means the trading size g only changes the stability of the market and cannot affect the Nash equilibrium of different social graphs.

3.6 Conclusions

We have defined a new agent-based model to study the extent to which social connections between market participants affect herding and market stability. This is achieved by incorporating classical opinion dynamics in the decision making process of the traders and using EGTA to compute the equilibria of a suitably defined strategic game. In these games, traders want to maximise their profits whilst balancing private and public information about the asset. Loosely speaking, we show that herding is the more likely the denser the social graph. There is however a surprising discontinuity point in what herding actually means for market stability. Essentially, the herd responds quickly to the actions taken by the stubborn traders.

We see our work as the introduction of a framework that can be used to study more ques-

tions about herding in financial markets. There are in fact still some gaps between our simulated market environment and the real market. This is a limitation we can overcome with an order-driven market or different market trading mechanisms. In Chapter 4, we will explore opinion formation and herding in order-driven markets with different social connections. We also study a new order-matching mechanism to determine the priority of order execution.

Chapter 4

Opinion formation and herding in order-driven markets

4.1 Introduction

Order-driven markets are used in many different financial markets, including stock exchanges, futures markets, and some foreign exchange markets. One of the advantages of order-driven markets is that they provide transparency and equal access to market information for all market participants. To refine our experiments, we apply our agent-based model to an order-driven market and use empirical game theory analysis to explore opinion formation and herding in financial markets with different social connections. We also investigate the impact of herding on different types of investors. The experimental results show that herding persists in order-driven markets. In the same market environment, herding is more pronounced for long-term investors.

For our experiment above, it is worth mentioning one important limitation, which is the absence of market makers. Market Makers (MMs) play an important role in financial markets, facilitating trading in order-driven markets. A market maker is a liquidity provider who can be a company or an individual. They quote both bid and ask prices for tradable assets held in inventory, hoping to make a profit from the bid/ask spread or pass-through [94]. The primary goal of market making is to ensure that financial markets operate smoothly and efficiently.

To overcome the limitation of not having market makers in our experiments, we added a new order matching mechanism, spread/price-time priority mechanism [27], which determines by involving three indicators (spread, price, and time) order execution priorities. By analysing the experimental results we find that as the weight of spreads in the order matching mechanism increases, the herding of the market becomes less pronounced. The spread/price-time priority mechanism can indeed improve market liquidity and stability. However, in order to make it a more efficient mechanism, the weighting of spreads in the order matching mechanism also needs to be adjusted.

4.2 The model of limit order book (LOB)

4.2.1 Limit order book (LOB)

As a result of the explosion of technological innovation over the last two decades, most financial transactions are now electronic and high in frequency. Limit Order Books (LOBs) are primarily used to facilitate electronic trading [95]. A LOB can be thought of as a file that stores a list of all market participants' orders for a particular asset; it contains a time record of each order, including the positions bought or sold and their corresponding quantities. The LOB ranks the orders on both sides of the market and regularly matches and executes the top-ranked buy and sell orders according to market rules.

A limit order book is a record of outstanding limit orders kept in the market. A limit order is an order to buy or sell a security at a specific price or better. A buy limit order is an order to buy at a pre-determined price or lower, while a sell limit order is an order to sell a security at a pre-determined price or higher. When a limit order for a security is entered, it is kept on record. As buy and sell limit orders are placed for securities, these orders are recorded in the order book. When the market moves to a pre-specified price, the market will execute the order at or better than the specified limit price. The market must ensure that the highest priority orders in the limit order book are executed before other orders in the book and before orders of equal or higher price are held or submitted by other traders on the floor (such as floor brokers and market makers).

LOBs are a simple way to encapsulate the complex dynamics of finance, where micro

interactions between traders lead to complex macro phenomena [96]. The dynamics of LOBs can be used to understand market price discovery and efficiency [97], help improve regulation and ensure the orderly functioning of capital markets. The study of the microstructure of LOBs is therefore also of interest to researchers and regulators. Overall, the LOB provides traders with valuable information on market dynamics, liquidity and sentiment.

4.2.2 Trading patterns of LOB

The limit order book is a real-time list of all limit orders for each participant over a period of time [98]. It consists of two main parts, buyers and sellers. Buyers provide orders with bids and order sizes to buy and sellers provide orders with sell prices and order sizes to sell. The highest bid refers to **Best Bid**, while the lowest ask refers to **Best Ask**. The **Spread** of the market is the difference between the highest price the buyer is willing to pay (the bid) and the lowest price the seller is willing to accept (the ask) (**Spread= Best Ask - Best Bid**). Traders can observe the bid/ask spread in the LOB to determine the liquidity of the market and the price at which they can enter or exit a trade. Participants can trade in the order book by using two types of orders - **Limit order** and **Market order**. A limit order is where a participant offers a specific price and quantity that he is willing to buy or sell at any time. A market order acts automatically by setting a specific number of shares; its advantage is that it is fast. It follows that market orders demand market liquidity, while limit orders provide liquidity to the market. Canceling an order is another action that participants can take to close their position in the market.

Table 4.1 shows an illustration of a limit order book for a stock. The next step is to illustrate the ranking of order execution and queuing through Table 4.1. If a limit order comes into the market and a participant wants to sell the stock of size 100 at 47.5£. This order will soon match the best bid price on the limit order book. However, if the sell price of this limit order is changed to 48.5, it will not be executed until the best bid price is at or above 48.5. On the other hand, suppose a buy market order of size 350 is entered into the market for this stock. This market order will immediately be filled by an order from Trader A (300 of 50£) and Trader B (50 of 51£). The limit order book therefore changes over time as a result of the interaction

Table 4.1: Example of a limit order book (LOB) for a stock

Trader	Ask/Bid	Price(£)	Size	
	Bid	42	500	
	Bid	45	350	
	Bid	47	400	
	Bid	48	200	Best Bid
A	Ask	50	300	Best Ask
B	Ask	51	100	
C	Ask	54	450	
D	Ask	56	550	


of the market and limit orders. At the same time, it provides an insight into a market where all participants are seeking the best price for an asset to gain maximum benefit, and where they interact with each other in the market. The orders in the trading example above are all at different prices. In fact, orders are matched following **price-time priority**. When there are two limit orders at the same price, the order that enters the market first will be executed first.

4.2.3 Agent-based model of LOB

We will explore opinion formation and herding in order-driven markets on the basis of Bartolozzi's proposed multi-agent model [85] by using EGTA. Figure 4.1 is a schematic depiction of an agent-based model of limit order book.

As shown in Figure 4.1, agents refer to participants who aim to gain benefits from trading in the market. Agent i gets its own opinion $s_i(t)$ at each time step t by the information they receive (private information $b_i(t)$ and neighbor information $s_j(t)$) and by their own stubbornness α (as in Section 3.2). When $s_i(t) \geq th_B$, the agent i will submit a bid order. When $s_i(t) \leq th_S$,

Figure 4.1: An agent based model of LOB



the agent i will submit an ask order. When $th_S < s_i(t) < th_B$, the agent i will not submit new orders.

At each time step agents submit orders to the market, the order generation involves three points: order type, order price and order size. In our experiments, we will address these issues using the same approach as [85]. Our decision to submit a limit order or a market order is related to the “*aggressiveness*” of the agent. If an order needs to be filled quickly, an agent may accept the cost of paying the difference up front and send a market order. The alternative is a less aggressive limit order, which will be executed at a pre-determined price. Orders that are further away from the opposite optimum position will take more time to complete. The price at which an order is submitted, and its “*aggressiveness*”, is derived from a log-normal distribution whose q -quantile is centered on the best bid or best offer. The “*direction*” of the distribution is also chosen according to the direction of the order. The size of an order that a trader is willing to buy or sell may be related to several factors: the specific execution strategy, the available liquidity in the market and the volatility at that point in time, among others. In a real trading environment, single large orders are often split into smaller orders to

minimise the costs associated with their impact. While this is an important issue in practice, to keep things simple we are only sending orders in a single block in this experiment. The number of contracts an agent is willing to buy or sell (g) is consistent with the randomness of the process, which is drawn from a log-normal distribution, rounded to whole numbers. The value of g must be greater than or equal to 1 and less than a quarter of the total volume contained on the appropriate side of the LOB.

At each time step, the agent in the LOB with the outstanding order will decide whether to delete it or not depending on whether it times out or not. That is, if the order is not executed within the T_{max} time step, it is automatically deleted. This time step is relatively long if compared to the average duration of the transaction.

Finally, orders entering the LOB at each time step will be matched according to **price-time priority**. Each agent will be given new neighbor information and will repeat the above steps at the next time step to make a new decision based on its own opinion.

4.3 Empirical Game Theory Analysis (EGTA)

In this section, we will use Empirical Game Theory Analysis (EGTA) to empirically solve the game and analyse the equilibrium state of the market described above. In our game, we first identify three strategies and find the Nash equilibrium in the game by calculating the average payoff of the agents under each strategy.

4.3.1 Three strategies

In our experiments, each agent has three strategies to choose from. They are “Imitation”, “Neutral” and “No imitation”. As mentioned above, a higher level of agent stubbornness means that the agent is more likely to believe in itself. This means that we can define these three strategies in terms of the value of α . For example, we use $\alpha = 0$ to represent the “No imitation” strategy. Then we choose a higher value of α to represent the “Imitation” strategy. The value of α for the “Neutral” strategy can be chosen within the interval of the α values of the two strategies mentioned above. These three strategies are set up in the same way as in Section 3.3. In our experiments, each agent can choose one of the three strategies per

experiment. At the end of each experiment, we record the number of agents who chose each strategy and their total payoff, respectively. By varying the number of agents for each strategy, we obtained each of the possible scenarios and their average total return. We then process these data analytically and find the Nash equilibrium by EGTA.

4.3.2 Payoffs

As defined in Section 3.3, the total assets of agent i at time step t will be the sum of the value of the stocks ($stock_i(t)$) it holds and the cash it currently owns ($cash_i(t)$) at time t (see (3.11) for details). Then the total payoff of agent i after t time steps will be determined by:

$$R_i(t) = Z_i(t) - Z_i(0) \quad (4.1)$$

4.3.3 Setup of experimental parameters

In our experiments, we set up 1000 agents, i.e. $N = 1000$. Initially, we would have liked to set up 10,000 agents as in previous work, but the EGTA part was computationally out of reach. Numerical tests show that when the number of agents reaches a certain value, the results are independent of the number of agents [85]. Our experiment will run for 1000 time steps. In previous experiments, we found that $T = 1000$ was sufficient to ensure that the simulation reached a steady state. We set $Tmax = 100$ which determines whether an order is time-out. This means that when the order exceeds 100 time steps and is still not matched, the order will be canceled. We set $th_B = \frac{2}{3}$, $th_S = \frac{1}{3}$. That means when the opinion for agent i at time step t , $s_i(t) > \frac{2}{3}$, the agent i will submit a bid order. When $s_i(t) < \frac{1}{3}$, agent i will submit a ask order. When the value of $s_i(t)$ is between $\frac{1}{3}$ and $\frac{2}{3}$, the agent will not offer any orders at time t . P_{Ba} and P_{Bb} are the best ask and best bid, respectively. The price $price(t)$ of the stock at each time step is the mid-price of the current time step of the limit order book ($price(t) = (P_{Ba} + P_{Bb})/2$). And the distribution, q-quantile, $q = 0.5$ in our case the same as [85]. And δ_{in} and δ_{out} determine the market crash and bubble's in point and out point. In our experiment, we set $\delta_{in} = 5\%$, $\delta_{out} = 2.5\%$ [65, 93]. Our experiment will be repeated 100 times and results are average over these runs.

4.3.4 The value of α

For the choice of the value of α , we use the approach of Section 3.3. We set $\alpha = 0.75$ as the “Imitation” strategy and set $\alpha = 0.5$ as the “Neutral” strategy. The last strategy is “No imitation” and it is represented by $\alpha = 0$. This strategy setup allows us to better analyse and compare opinion formation and herdings in the market environment we set in Section 3 and the order-driven market environment.

4.3.5 Experimental results

After 100 simulations we obtained the results shown in Figure 4.2 and the proportion of agents at Nash equilibrium is shown in Figure 4.3. We have 94% agents who choose “Imitation” strategy, 6% agents choosing “Neutral” strategy, and 0% agents choosing the “No imitation” strategy. This result is similar to what we got in Section 3.3.6. The only difference is that in the order-driven market environment, no agent chooses “No imitation” strategy. This means that in the CDA market, we have designed, no agent to trust only its own private information. In contrast to the previous results, most agents in an order-driven market environment are willing to trust and imitate the decisions of their neighbors.

Figure 4.2: The Nash equilibrium of EGTA





Figure 4.3: The proportion of agents at Nash equilibrium of EGTA

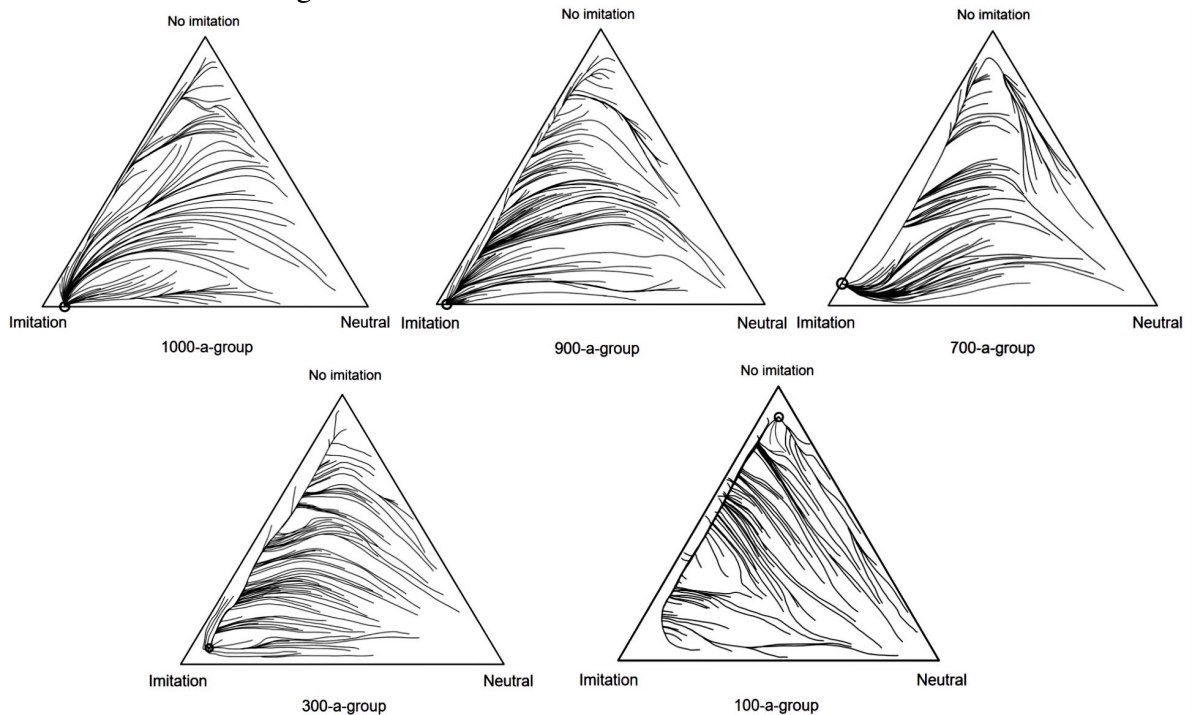


4.3.6 Different social connections

Having obtained these results, we will continue to explore whether the different social connections between agents will have an impact on the results of our experiments. The Nash equilibrium of different social connections shown in Figure 4.4 and Table 4.2 shows the proportion of agents at the Nash equilibrium of different social connections. As in Section 3.4, ‘ x -a-group’ represents a graph where each group has x agents and each agent has $x - 1$ neighbors.

The first column in Table 4.2 shows the social ties between agents, and the second to fourth columns show the proportion of agents choosing different strategies at the Nash equilibrium point, respectively. From the information given in this figure, we can see that most agents chose the “No imitation” strategy and a small number of agents chose “Imitation” strategy and “Neutral” strategy at ‘100-a-group’ and ‘200-a-group’. Starting with the ‘300-a-group’, most agents prefer the “Imitation” strategy. No agent chose “Neutral” strategy from ‘400-a-group’ to ‘800-a-group’ and only a small number of agents chose “No imitation” strategy. The choice of agent has changed at ‘900-a-group’, there are 97% agents choosing the “Imitation” strategy and 3% agents choosing the “Neutral” strategy, and no agent chooses the “No imitation”

Figure 4.4: The NE of different social connections



strategy anymore.

Based on the results obtained, we can analyse that most agents prefer to trust their own private information when they are given relatively little neighbor information (like '100-a-group' and '200-a-group'). They believe that the information they receive about their neighbors is not sufficient to change their decisions. When an agent is given enough information about its neighbors (from '300-a-group' to '800-a-group'), most agents will choose "Imitation" strategy, they are willing to imitate the trading decisions of their neighbors. But there will still be a small number of 'stubborn' agents in the market who still believe only in their own private information. At this point, herding is already present in the market. When agents have access to the information of almost all other agents in the market ('900-a-group'), there will be no agents in the market who believe only in their own private information. Almost all agents will choose "Imitation" strategy due to herding. But when every agent in the market knows all the other agents' information ('1000-a-group'), the proportion of agents choosing the imitation strategy decreases. This may be because for each agent in the market, they are given almost the same information. A small proportion of them do not believe that choosing an imitation

strategy will lead to higher returns.

Table 4.2: The proportion of agents at Nash equilibrium

	Imitation	Neutral	No imitation
1000-a-group	94%	6%	0%
900-a-group	97%	3%	0%
800-a-group	93%	0%	7%
700-a-group	94%	0%	6%
600-a-group	94%	0%	6%
500-a-group	93%	0%	7%
400-a-group	91%	0%	9%
300-a-group	91%	4%	5%
200-a-group	5%	6%	89%
100-a-group	5%	5%	90%

4.4 Two types of investors

4.4.1 Long-term and short-term investors

There are two types of investors in the market: long-term investors and short-term investors. They are both essential presences in the market. Long-term investors and short-term investors are two types of investors with different investment objectives and strategies. Short-term investors aim to make quick profits in a relatively short period of time. Long-term investors, on the other hand, take a more patient approach to building wealth over a longer period of time.

Long-term investing usually refers to putting money into an investment vehicle and holding it for an extended period of time with a view to achieving a higher return. This investment strategy requires patience and a long-term perspective. Because markets are volatile, short-term fluctuations may cause investors to panic. However, if held for the long term, it can

often yield better returns. In contrast, short-term investing is more short-sighted and usually involves putting money into the market with the aim of achieving high returns in the short term. This investment strategy is usually more flexible and potentially more profitable but is also more risky due to higher market volatility. In short, both long-term and short-term investments have their advantages and disadvantages. Long-term investments may generate higher returns, but they require time and patience. Short-term investments, on the other hand, may be more profitable but involve a higher level of risk. In this section, we build on our previous experiments to explore whether herdings have different implications for long-term and short-term investors.

4.4.2 Experimental results and analysis

In simple terms, the difference between long-term and short-term investors in our experiments is the difference in the number of time steps agents take to trade in the market. In the above experiment, we run 1000 time steps per experiment. We consider the agents in Section 4.3.6 to be long-term investors. We then changed the number of running time steps to 100 for each experiment and simulated it 100 times. The results obtained are shown in Table 4.3.

From the information in Table 4.3 we can see that no agent chose “Imitation” strategy at ‘100-a-group’. There are 96% agents who chose “No imitation” strategy and 4% agents choosing “Neutral” strategy at this social connection. The proportion of agents choosing “Imitation” strategy increases with the amount of information available to agents in the market and reaches a peak at ‘900-a-group’ (90%). On the other hand, the number of agents who are willing to trust only their own private information is becoming less and less. This proves that herding has an impact on short-term investors as well. The reason for the specificity in the ‘1000-a-group’ should be the same as for Section 4.3.6 is that when agents in the market are given almost identical information, there will be a small fraction of agents who will no longer be willing to imitate their neighbors’ decisions.

To get a more intuitive view of the impact of the herding on long-term and short-term investors in the market, we compare the proportion of agents choosing the “Imitation” strategy in Table 4.2 and Table 4.3. From Figure 4.5 we can see that, given the same social connections,

Table 4.3: The NE of short-term investors

	Imitation	Neutral	No imitation
1000-a-group	86%	6%	8%
900-a-group	90%	5%	5%
800-a-group	88%	5%	7%
700-a-group	77%	12%	11%
600-a-group	65%	17%	18%
500-a-group	53%	26%	21%
400-a-group	21%	22%	57%
300-a-group	12%	14%	73%
200-a-group	4%	6%	90%
100-a-group	0%	4%	96%

long-term investors always have more agents choosing “Imitation” strategy than short-term investors. This is because as the time step increases, more and more agents are willing to accept and imitate the behavior of their neighbors. This means that the herding is more pronounced for long-term investors in the same market environment.


4.5 Spread/price–time priority

4.5.1 Market maker (MM)

For our experiment above, it is worth mentioning one important limitation, which is the absence of a market maker. Market makers (MM) play an important role in financial markets, facilitating trading in order-driven markets. A market maker is a liquidity provider who can be a company or an individual. They quote both bid and ask prices for tradable assets held in inventory, hoping to make a profit from the bid/ask spread or pass-through [94]. The primary goal of market making is to ensure that financial markets operate smoothly and efficiently.

Market makers can be broadly divided into two types, the first being official market mak-

Figure 4.5: Investors who choose “Imitation” strategy



ers. They are often designed to act as market makers and dealers, and they are contractually obliged to provide reasonable spreads and maintain order and fairness in the market. At the London Stock Exchange, official market makers are available for many securities. A number of LSE member firms have taken on the obligation to always set a two-way price for each stock in which they make a market. Their price is the one displayed on the Stock Exchange Automated Quotation (SEAQ) system. The second type are unofficial market makers, who are not obliged to always set two-way prices, but who also do not have the advantage of everyone having to deal with them. In a limit order book (LOB), investors can submit orders for both sides and quotes can be stored in the LOB until the order is traded or canceled [99]. They also provide liquidity to the market without necessarily having to comply with contractual restrictions. Thus, in order-driven markets, electronic trading systems allow participants to become market makers and seek profits from their high-frequency trading strategies through two-sided quotes. In our research, we consider market makers to be unofficial market makers who trade in the market without obligations or contractual restrictions.

Previous research has shown that the impact of market makers is expected to be positive, promoting market liquidity, price discovery and market stability [100–102]. Market makers generally adapt to temporary order imbalances by carrying some inventory, they not only make a profit but also stabilise prices [103]. Because of this, market makers may accumulate large inventories over a period of time. While they can effectively hedge their risk with future contracts, this inventory risk can affect the liquidity of the market [104]. Overall, the impact of market makers on market liquidity and stability has been positive.

Due to the benefits that the study of market maker has brought to the market, a large research literature on optimal market making has emerged around the world [105–108]. It is worth noting that the study of MM can contribute to the optimisation and design of trading mechanisms. We can maintain market stability by incentivising market makers [27].

In high-frequency trading, participants can become market makers on a voluntary basis. They maximise their profits by managing inventories and taking advantage of the spreads. The widely used **price-time** ranking system will stimulate high-frequency traders to always submit orders with very low latency and close to the best ask/best bid price. This is because, in a price time ranking system, the best bid/ask price received first will be executed first in the LOB. As this matching mechanism only takes into account the order price, this can lead to high-frequency traders offering a large number of one-sided unbalanced orders, resulting in large price swings, as was the case in the flash crash of 2010. If one thinks about this problem from the perspective of an order-matching mechanism, the matching mechanism should not only encourage traders to enter the market but also avoid the possibility of such imbalances. The paper [27] proposes a new type of order matching mechanism called **spread/price-time** priority, which determines three indicators (spread, price and time) to determine the order execution priorities, thus overcoming these uncoordinated incentives. The article suggests that a queuing system for limited order books rewards market participants by offering competitive prices in both directions. The model simulations they present suggest that this new order-matching mechanism is likely to enhance market stability. In order to get rid of the limitation of the absence of market makers in the market and to explore the impact of herding on the market environment with different trading matching mechanisms, we add **spread/price-time**

priority to our previous experiments and attempt to explore the impact of this new order matching mechanism on the liquidity and stability of the order-driven market mentioned above.

4.5.2 Order matching mechanism

In the previous experiment, orders in the market are matched according to price-time priority. Golub et al. developed a new matching system based on price-time priority spread/price-time priority mechanism [27]. They make the spread another indicator for the ranking of executed orders. In essence, this new matching mechanism provides a new way to reward participants who submit two-sided limit orders rather than one-sided orders.

Spread/price-time priority has two ranking attributes, price and spread. In this setting, the top-ranked orders are executed with the highest priority. The ranking is a weighted average ranking of the spread ranking and the price ranking based on the parameters $\beta \in [0, 1]$, as in Formula (4.2).

$$Rank(\beta) = \beta \cdot spread + (1 - \beta) \cdot price \quad (4.2)$$

In our experiments, the price ranking *price* is the distance from the best ask/bid. We use P_{ask} and P_{bid} to indicate the price of the ask/bid order provided by the agent, P_{Ba} and P_{Bb} to determine the price of the **Best Ask** and **Best Bid**. The spread ranking *spread* is generally the difference between the prices of two-side orders offered by the agent. For agents offering only one-side orders, the spread on their order is the highest spread between the two sides of the limit order in the LOB, which we denote by $spread_{Max}$. The *price* and *spread* are defined as follow:

$$price = \begin{cases} 0, & \text{new best buy/sell price} \\ P_{ask} - P_{Ba}, & \text{sell limit order} \\ P_{bid} - P_{Bb}, & \text{buy limit order} \end{cases}$$

$$spread = \begin{cases} spread_{Max}, & \text{one-side limit order} \\ P_{ask} - P_{bid}, & \text{two-side limit order} \end{cases}$$

Orders with smaller spreads will result in a more balanced market. Conversely, if the

spread is large, the volatility of the market will increase. Therefore, the smaller the spread, the higher the priority of the order. Limit orders on two sides have a higher spread priority than limit orders on one side. In summary, spread/price-time is an optimised version of price-time priority, where β is used to adjust the importance of the spread relative to the primary ranking of the price. A larger value of β means that the spread has a greater impact on the matching priority of an order, and vice versa. In simple terms, β is how much importance the market’s order matching mechanism places on spreads. In addition, it provides more liquidity to the market by encouraging participants to submit two-side limit orders. The **spread/price-time** priority matching mechanism also reduces the market’s spreads, making the market more stable [27].

4.5.3 Experimental setup and results

As noted above, we replace the price-time priority order matching mechanism from the experiments in Section 4.3.5 with the current spread/price-time priority. We use μ to represent the probability of an agent submitting a two-side order, $\mu \in [0, 1]$. The probability of an agent submitting a one-sided order is $(1 - \mu)$. That is, in our agent-based model of LOB there are two new degrees of freedom: the parameter $\beta \in [0, 1]$ determines the primary ranking (β), and the parameter $\mu \in [0, 1]$ determines the probability of submitting a two-sided limit order. We will test the following combinations of values for these two parameters: $\beta \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$ and $\mu \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. In total, we obtained 55 games with different pairs of β and μ , each set of parameter pairs we denote by (β, μ) . For each (β, μ) we perform 100 simulations and run each experiment for 1000 time steps, as we have done in Section 4.3.5. Similarly, every agent in the market has three strategies to choose from: “Imitation” strategy ($\alpha = 0.75$), “Neutral” strategy ($\alpha = 0.5$) and “No imitation” strategy ($\alpha = 0$).

The results obtained are as follows, Table 4.4, 4.5 and 4.6 show the Nash equilibrium we obtain for different pairs of (β, μ) . The first column of the table shows the value of β , and the second to fourth columns show the proportion of agents choosing “Imitation” strategy, “Neutral” strategy and “No imitation” strategy at Nash equilibrium.

Table 4.4: The table on the left is $\mu = 0.1$ and the right one is $\mu = 0.3$

β	Imitation	Neutral	No Imitation
0	94%	6%	0%
0.1	91%	8%	1%
0.2	86%	10%	4%
0.3	81%	12%	7%
0.4	70%	17%	13%
0.5	60%	20%	20%
0.6	40%	18%	42%
0.7	37%	14%	49%
0.8	20%	9%	71%
0.9	7%	6%	87%
1	0%	0%	100%

β	Imitation	Neutral	No Imitation
0	94%	6%	0%
0.1	90%	10%	0%
0.2	85%	12%	3%
0.3	81%	14%	5%
0.4	72%	19%	9%
0.5	56%	22%	22%
0.6	35%	20%	45%
0.7	37%	13%	50%
0.8	19%	9%	72%
0.9	6%	5%	89%
1	0%	0%	100%

Table 4.5: The table on the left is $\mu = 0.5$ and the right one is $\mu = 0.7$

β	Imitation	Neutral	No Imitation
0	94%	6%	0%
0.1	90%	10%	0%
0.2	86%	11%	3%
0.3	80%	14%	6%
0.4	72%	18%	10%
0.5	55%	23%	22%
0.6	41%	18%	41%
0.7	34%	13%	53%
0.8	19%	7%	74%
0.9	4%	5%	91%
1	0%	0%	100%

β	Imitation	Neutral	No Imitation
0	94%	6%	0%
0.1	89%	11%	0%
0.2	80%	15%	5%
0.3	71%	18%	11%
0.4	62%	22%	16%
0.5	52%	25%	23%
0.6	45%	20%	35%
0.7	32%	14%	54%
0.8	18%	8%	74%
0.9	5%	5%	90%
1	0%	0%	100%

Table 4.6: $\mu = 0.9$

β	Imitation	Neutral	No Imitation
0	94%	6%	0%
0.1	85%	9%	6%
0.2	76%	15%	9%
0.3	67%	20%	13%
0.4	57%	24%	19%
0.5	46%	29%	25%
0.6	34%	28%	38%
0.7	20%	23%	57%
0.8	11%	16%	73%
0.9	2%	4%	94%
1	0%	0%	100%

From the information in the figures above we can see that as the weight of spread (β) gets larger, fewer agents choose the “Imitation” strategy and more agents choose the “No imitation” strategy. This proves that herding on the market is becoming less pronounced. As the market gives less and less weight to price ($1 - \beta$), agents do not get higher returns by choosing the “Imitation” strategy. When $\beta = 1$, all agents choose the “No imitation” strategy. In our experiments, the agents’ choice of the “No imitation” strategy means that the agents are only influenced by their own private information. This shows that when the priority of order matching in the market is influenced only by the spread, agents will only make decisions based on their own private information and will not choose the imitation strategy.

4.5.4 Liquidity and stability of the market

To better observe the impact of the new order-matching mechanism on the liquidity and stability of our simulated order-driven market. We will also analyse the trading volume and price volatility of the market in each experiment. The trading volume we obtained is the average of the cumulative trading volume per 1000 time steps over 100 simulated experiments. To account for different frequency traders, the price volatility is calculated on smaller and larger scales. In other words, the price volatility σ in our experiment is the average of the price volatility at 20, 40, 100, 200, 300, 400, 500, and 1000 time steps [27]:

$$\sigma = \frac{\sigma_{20} + \sigma_{40} + \sigma_{100} + \sigma_{200} + \sigma_{300} + \sigma_{400} + \sigma_{500} + \sigma_{1000}}{8},$$


where σ_x is the standard deviation of price over windows of x steps. The results of the experiment are shown in Table 4.7.

Based on the information in Table 4.7 it is easy to see that either parameter β or parameter μ can influence the average volume of the market. When we fix the weight of the market order matching mechanism on the spread (the value of β), we compare the information in the table horizontally. We can see that there is a clear trend towards an increase in the average volume of the market as the probability of an agent submitting a two-sided order (μ) increases. When we look at the information in the table vertically we can see that the average market volume increases as the value of β becomes bigger. However, this increasing trend is not evident when

$\beta > 0.5$. When $\beta \leq 0.5$, the average market volume increases significantly as the value of β becomes bigger. We have the following explanation for this: a higher value of μ represents a higher probability for each trader to become a market maker and submit two-sided orders. As a result, the number of orders in the market will increase, which leads to more executions. The **spread/price-time** priority mechanism will result in a significant increase in the average trading volume of the market and this growth trend will slow down as spread dominates more in the ranking ($\beta > 0.5$).

Table 4.7: Average volume for different (β, μ)

β	$\mu=0.1$	$\mu=0.3$	$\mu=0.5$	$\mu=0.7$	$\mu=0.9$
0	39976	50963	60479	71835	80456
0.1	45355	54316	66841	75247	84932
0.2	49753	60874	70523	80634	89255
0.3	55422	66028	75714	86744	95472
0.4	61257	71149	80571	91481	107933
0.5	71448	80633	91483	101479	115647
0.6	85564	95421	105822	115482	125429
0.7	84508	98364	104965	114833	130966
0.8	87011	97465	107837	117895	128431
0.9	85497	100522	115478	120178	124867
1	87631	95629	108974	114786	128763



Volume increases

The average volatility for different (β, μ) are shown in Table 4.8. The information in this figure shows that increasing the value of μ reduces volatility, which means the greater the probability of an agent submitting a two-sided order in the market, the lower the average volatility of the market. When the weight of the price in the order matching mechanism is higher than the weight of the spread, the volatility of the market becomes higher ($\beta < 0.5$). It

is clear that the volatility of the market is highest when the contribution of the spread is zero. When we fix the probability of an agent submitting a two-sided order, the lowest volatility of the market always occurs at $\beta = 0.5$. When $\beta > 0.5$, there is a relatively weak increase in the volatility of the market. The following explanation can be given for these phenomena. The higher the value of μ , the higher the probability that an agent will offer a two-sided order. This implies an increase in the number of orders in the market and a sufficient number of orders contributes to market stability. Price-time priority ($\beta = 0$) gives priority to the execution of orders close to the mid-price, which may lead to a large number of unilateral orders in the market causing large price fluctuations. The order matching mechanism with spread solves this problem, as it encourages agents to submit a two-sided order. As shown in Table 4.8, the optimal weighting of the spread is $\beta = 0.5$. However, when the spread is fully dominant, the priority of order trading is determined entirely by the spread. This will result in a bilateral order being offered at a price that is significantly different from the market's mid-price, but it will still be executed and the volatility of the market will increase as a result.

Table 4.8: Average volatility for different (β, μ)

β	$\mu=0.1$	$\mu=0.3$	$\mu=0.5$	$\mu=0.7$	$\mu=0.9$
0	227563	206394	178671	150967	119752
0.1	212458	195471	170584	146982	110347
0.2	185739	176697	157961	127736	104728
0.3	142076	135941	117430	87234	67201
0.4	75279	65237	53479	54912	40537
0.5	3146	3079	2984	3108	2943
0.6	3484	3395	3314	3357	3297
0.7	3759	3743	3671	3619	3681
0.8	4294	4176	4093	4129	4081
0.9	4637	4692	4561	4621	4543
1	5179	5140	5086	5127	4982


Volatility increases

In general, with the increasing weight of spread in the order-matching mechanism, the herding in the market is becoming less pronounced. Increasing the probability of market agents submitting two-sided orders can improve the liquidity and stability of the market. Increasing the weight of spreads in the order-matching mechanism can also reduce market volatility and increase the volume of orders traded in the market. However, these positive effects on the market slowly disappear as the contribution of spreads reaches a threshold. According to the experimental results above, the spread/price time priority mechanism encourages traders to submit two-sided orders. But in order to make it a more effective mechanism, it is crucial to fine-tune the weighting of spreads in the order matching mechanism.

4.6 Conclusions

In this chapter, we leverage agent-based modeling to delve into order-driven markets. Through empirical game-theoretic analysis, we explore opinion formation and herding effects in financial markets with different social connections. Our exploration also delves into the influence of the herd effect on various investor types. The experimental outcomes unveil a persistent presence of the herd effect in order-driven markets. Furthermore, it becomes evident that this herd effect is notably more pronounced among long-term investors operating within the same market environment. To gauge the robustness of our findings across varying trading mechanisms, we meticulously examine a novel order-matching system known as the spread/price-time priority mechanism. These results reveal a noteworthy observation: as the weight assigned to the spread in the order-matching mechanism increases, the herding effect within the market tends to diminish. While the spread/price-time prioritization mechanism shows promise in enhancing market liquidity and stability, it is evident that adjustments to the weightings of spreads within the mechanism may be necessary for optimal performance and efficiency.

Chapter 5

Exploring herding in order-driven markets through a meta-game

5.1 Introduction

In Chapter 4 we examine opinion formation and herding in order-driven markets through an agent-based model (ABM) and Empirical Game-Theoretic Analysis (EGTA). In our experiment, we assume that the number of agents is constant, and strategies are fixed. By assuming a constant number of agents, we can remove the complexity of agents entering or exiting the system, thus making the model easier to manage and analyze. A fixed policy means that agent behavior remains consistent throughout the simulation, which can lead to stable and predictable results. With constant agents and fixed strategies, we can focus more on the interaction between agents and the market environment. However, assuming a constant number of agents and a fixed policy may not accurately represent dynamic real-world systems. If the strategies are fixed, we may miss the opportunity to explore the influence of agents by other strategies. In the previous experiments, each agent had three strategies to choose from. They are “Imitation”, “Neutral” and “No imitation”. We defined these three strategies in terms of the value of $\alpha \in [0, 1]$ (α means how much agents trust their neighbors, see Section 3 for more details). In fact, the value of α can be any value within the range of values it takes. In other words, in the games we design, the agents can choose from a wide range of strategies,

such as $\alpha = 0.1, 0.2, 0.3$, etc. To refine our experiments, we enrich the set of strategies that agents can choose from. Correspondingly, this game will become very complex. It will have a large number of possible entries, which will make it impossible for us to handle. Therefore, we will use a meta-game to simplify this game. The results of the experiment show that in order-driven markets herding is more pronounced for long-term investors than for short-term investors. Given the same connections, more agents in the meta-game chose the “Neutral” strategy than in the previous game.

5.2 Meta-game

A normal-form game is a matrix presentation of a game. It involves players, possible actions, and the rewards of choosing the appropriate strategy. As a game becomes complex, with more players and strategies, the entries in a normal-form game increase exponentially. In contrast to simple games, complex systems can be analysed at a high level by grouping the strategies available to the players. The set of strategies obtained from these groupings is called meta-strategy [109]. Rather than exploring all possible strategies, a meta-strategy is a classification based on all possible strategies in the game. A new meta-strategy will be defined as the set of strategies that participate in a meta-game.

In addition, in game theory, games can be divided into symmetric and asymmetric games. They refer to the degree to which the participants in a game and their strategies are similar or different. In a symmetric game, all players have the same set of strategies and payoffs, and the game is identical from each player’s perspective [110]. In other words, the players have the same roles and there is no inherent advantage or disadvantage to any player. Examples of symmetric games include the classic Prisoner’s Dilemma and Rock-Paper-Scissors games. In contrast, in an asymmetric game, players have different sets of strategies and payoffs, and each player may have a different perspective. In other words, players play different roles, and different players may have inherent advantages or disadvantages. In our experiments, the focus is on symmetric games.

In our experiments, the choice of strategies available to agents in the market is too rich to model. So we try to define a set of three simple meta-strategies to reduce the complexity

of the game. By analysing the experimental results of the meta-game, we can get a more comprehensive understanding of the impact of herding in order-driven markets at different social connections between agents.

5.3 Experimental setup

Based on the experiments in Section 4.3, we have added a new meta-strategy setting and updated the calculation of the payoff for agents who choose meta-strategies.

5.3.1 Meta-strategies

As we mentioned above, the strategies available to the agent are determined by the degree of trust the agent has in its neighbors (the value of $\alpha \in [0, 1]$). We, therefore, tried to reduce the complexity of the game by defining all the strategies available to the agent as three meta-strategies: “Imitation”, “Neutral” and “Confident”. The first meta-strategy is the “Imitation” strategy, An agent choosing this meta-strategy is always willing to trust its neighbor information and imitate its neighbors’ decisions. So we define this meta-strategy by $\alpha \in (\frac{2}{3}, 1]$. We use $\alpha \in [\frac{1}{3}, \frac{2}{3}]$ to define the “Neutral” strategy. The agent who chooses this meta-strategy remains neutral in its decision whether or not to imitate its neighbor. The last meta-strategy is “Confident” strategy. We replaced the previous “No imitation” strategy with this meta-strategy. Agents who choose this meta-strategy are more likely to trust their own private information. We define this meta-strategy by $\alpha \in [0, \frac{1}{3})$.

5.3.2 Payoffs

We need to analyse the evolution of the meta-strategy through the relative expected payoffs between the strategies and observe the market conditions in the process. A heuristic payoff table is a tool used in decision-making that provides a simplified representation of the possible outcomes of a decision, along with the probabilities of those outcomes and their corresponding payoffs. It is often used when it is difficult or impossible to determine the exact probabilities and payoffs of different decisions. By using a heuristic payoff table, we can compare the expected values of different decisions and choose the one that offers the highest expected payoff [111]. Figure 5.1 indicates an example for generating a heuristic payoff table for 3

meta-strategies (s_1, s_2 and s_3) and 3 individuals. N represents a matrix which contains all the discrete distributions and U is the corresponding payoff matrix. $P = (N, U)$ represents the heuristic payoff table which captures information on payoffs for all possible distributions among a finite population. In our experiments, we record the average payoff of each agent in a heuristic payoff table according to formula (3.11) and formula (3.12).

Figure 5.1: An example of a heuristic payoff table for 3 meta-strategies and 3 individuals

First Step:

Enumerate a matrix N of all possible distributions over 3 meta-strategies for 3 individuals.

$$N = \begin{matrix} & s_1 & s_2 & s_3 \\ \begin{bmatrix} 0 & 0 & 3 \\ 0 & 1 & 2 \\ 0 & 2 & 1 \\ 0 & 3 & 0 \\ 1 & 0 & 2 \\ 1 & 1 & 1 \\ 1 & 2 & 0 \\ 2 & 0 & 1 \\ 2 & 1 & 0 \\ 3 & 0 & 0 \end{bmatrix} & & & \end{matrix}$$

10 simulations

Second Step:

Each distribution over strategies can be simulated, finally returning a matrix U of corresponding payoffs of N .

$$U = \begin{matrix} & U(s_1) & U(s_2) & U(s_3) \\ \begin{bmatrix} 0 & 0 & 0.8 \\ 0 & 0.4 & 0.5 \\ 0 & 0.6 & 0.6 \\ \dots & \dots & \dots \\ \dots & \dots & \dots \\ \dots & \dots & \dots \\ \dots & \dots & \dots \\ \dots & \dots & \dots \\ \dots & \dots & \dots \\ \dots & \dots & \dots \\ 0.6 & 0 & 0 \end{bmatrix} & & & \end{matrix}$$

Third Step: Heuristic Payoff Matrix: $P = (N, U)$

N			U		
0	0	3	0	0	0.8
⋮	⋮	⋮	⋮	⋮	⋮
3	0	0	0.6	0	0

Below we show a simplified meta-game payoff table for a 1000 player meta-game and three meta-strategies with $(1001 + 1) \cdot \frac{1001}{2} = 501,501$ entries. Each row on the left, representing one combination of players, sums three positive integers to 1000. The three integer columns represent the number of agents who choose “Imitation”, “Neutral” and “Confident” strategies respectively. U_I , U_N and U_C respectively correspond to the payoffs for choosing the


three strategies. Each row in the payoff table P can be obtained by one simulation, and an entire meta payoff table will be 501,501 simulations.

$$P = \left(\begin{array}{ccc|ccc} 0 & 0 & 1000 & U_I(0) & U_N(0) & U_C(1000) \\ 10 & 0 & 990 & U_I(10) & U_N(0) & U_C(990) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1000 & 0 & 0 & U_I(1000) & U_N(0) & U_C(0) \end{array} \right)$$

5.4 Experimental results and analysis

As in Section 4.3, our experiments simulate 100 times, with a time step of 1000 for each experiment. The Nash equilibrium is shown below.


Figure 5.2: The Nash equilibrium of meta-game



The proportion of agents choosing the three meta-strategies at this Nash equilibrium is shown in Figure 5.3. The proportion of agents choosing the “Imitation” strategy is 90%, the proportion of agents choosing the “Neutral” strategy is 10% and no agents chose the “Confident” strategy. We can see that herding is still present in the market, but that more agents choose “Neutral” strategy and fewer agents choose “Imitation” strategy compared to the results in Section 4.3.5. This means that when we have a less stringent definition of the strategies

that agents can choose, more will choose the neutral strategy.

Figure 5.3: The proportion of agents at Nash equilibrium

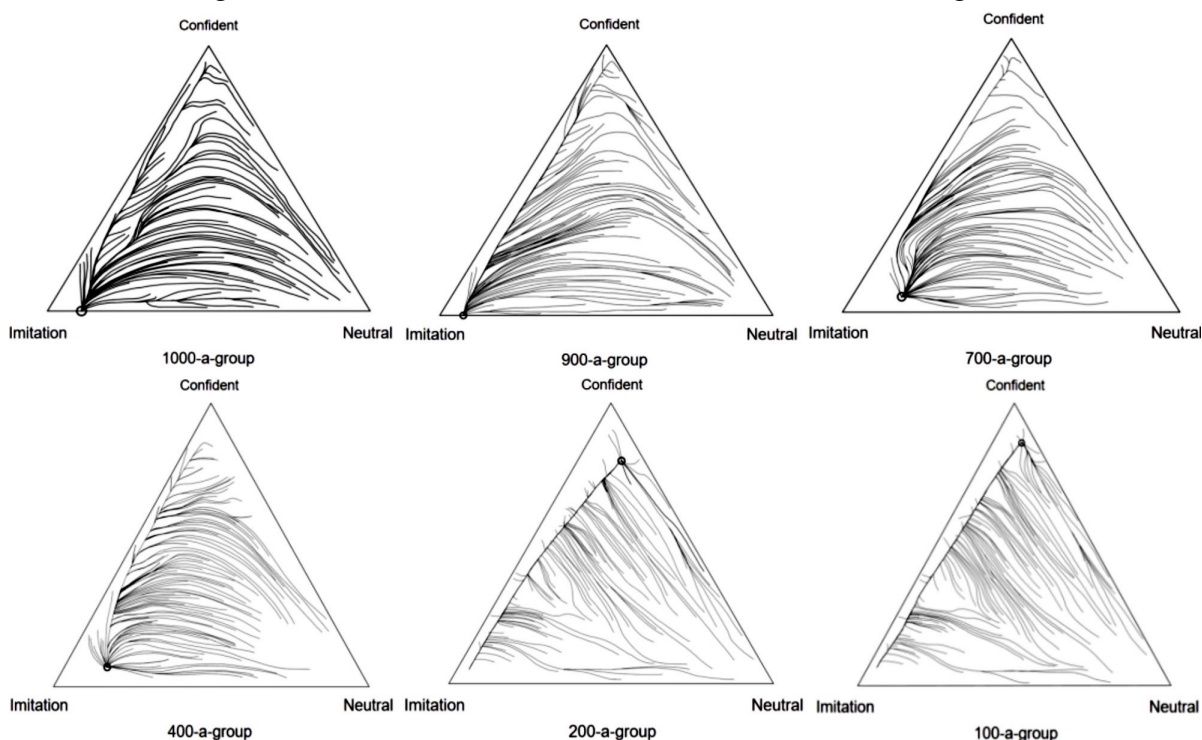


5.4.1 Different social connections in the meta-game

Having obtained these results, we attempted to explore the effects of different social connections between agents on the strategy choices of agents at Nash equilibrium in the meta-game. There are 10 different types of social connections between agents as the same in Section 4.3.6. We obtain the following results in Figure 5.4.

The proportion of agents at Nash equilibrium of different social connections is shown in Table 5.1. From the results given in Table 5.1, we can see that at ‘900-a-group’ there are 93% agents who choose “Imitation” strategy and 7% agents choose “Neutral” strategy, no agent chooses “Confident” strategy. Starting with the ‘800-a-group’, most agents prefer the “Imitation” strategy and a small number of agents choose “Neutral” strategy and “Confident” strategy. Most agents choose “Confident” strategy at ‘100-a-group’ and ‘200-a-group’. In general, the same results as those obtained in Section 4.3.6, as agents gain more neighbor

Figure 5.4: The NE of different social connections in meta-game




information, the proportion of agents choosing “Imitation” strategy becomes larger. Most agents prefer to trust their own private information when they are given relatively little neighbor information. However, when every agent in the market knows information about all other agents (‘1000-a-group’), the proportion of agents choosing the imitation strategy decreases compared to ‘900-a-group’. This is because when the agents in the market are socially related to the ‘1000-a-group’, for each agent they receive almost the same information. This leads to a small number of agents who do not believe that choosing an imitation strategy will lead to higher returns. When we compare the results in Table 5.1 with Table 4.2, it is easy to see that the proportion of agents choosing the “Neutral” strategy has increased significantly as shown in Figure 5.5, while the number of agents choosing the other two strategies has decreased. This phenomenon can be explained by the fact that agents in our meta-game have more strategies to choose from compared to previous games, which leads to fewer extreme agents choosing “Imitation” strategy and “Confident” strategy and more agents choosing “Neutral” strategy.

Table 5.1: The proportion of agents at Nash equilibrium in meta-game

	Imitation	Neutral	Confident
1000-a-group	90%	10%	0%
900-a-group	93%	7%	0%
800-a-group	87%	8%	5%
700-a-group	86%	10%	4%
600-a-group	84%	9%	7%
500-a-group	86%	8%	6%
400-a-group	83%	11%	6%
300-a-group	81%	10%	9%
200-a-group	6%	12%	82%
100-a-group	4%	8%	88%

Figure 5.5: Proportion of agents choosing “Neutral” strategy



5.4.2 Long-term and short-term investors

Next we will try to explore whether the herding has different implications for long-term and short-term investors in the meta-game. In our experiments, the difference between long-term and short-term investors is the difference in the number of time steps agents take to trade in the market. In the above experiment, we run 1000 time steps per experiment. We consider the agents in Section 5.4.1 to be long-term investors. We then changed the number of running time steps to 100 for each experiment and simulated it 100 times. We consider the agents in this experiment to be the short-term investors. The results obtained are shown in Table 5.2.


Table 5.2: The NE of short-term investors in meta-game

	Imitation	Neutral	Confident
1000-a-group	83%	10%	7%
900-a-group	86%	8%	6%
800-a-group	81%	9%	10%
700-a-group	72%	12%	16%
600-a-group	59%	21%	20%
500-a-group	43%	29%	28%
400-a-group	19%	30%	51%
300-a-group	11%	22%	67%
200-a-group	2%	14%	84%
100-a-group	0%	8%	92%

From the information in Table 5.2 we can see that no agent chose the “Imitation” strategy at ‘100-a-group’. There are 92% agents who chose the “Confident” strategy and 8% agents chose “Neutral” strategy at this social connection. The proportion of agents choosing “Imitation”

strategy increases with the amount of information available to agents in the market and reaches a peak at ‘900-a-group’ (86%). On the contrary, the number of agents who are willing to trust only their own private information is becoming less and less. This proves that herding has an impact on short-term investors in meta-game as well. When we compare the results in Table 4.3 and 5.2 we can see that more short-term investors choose the neutral strategy and fewer agents choose the other two strategies in the meta-game with the same connection compared to the results in Section 4.4. We have compared the proportion of long-term and short-term investors who choose imitation strategies based on the information in Table 5.1 and 5.2. As shown in Figure 5.6, it is easy to see that there are always more long-term investors opting for “Imitation” strategy than short-term investors with the same connection. This is because as the time step increases, more and more agents are willing to accept and imitate the behavior of their neighbors and Fewer investors are willing to trust their own private information.

Figure 5.6: Investors who choose “Imitation” strategy in meta-game



In general, a greater proportion of long-term investors have always opted for “Imitation” strategy than short-term investors, both in the previous game and in the meta-game. This implies that herdings are more pronounced for long-term investors in order-driven markets than

for short-term investors. With the same connection, more agents choose “Neutral” strategy in the meta-game than in the previous game. This is because, when the agents in the game have more strategies to choose from, there will be fewer relatively extreme agents who are willing to trust only themselves or who are more inclined to imitate the decisions of others.

5.5 Conclusions

In this chapter, we propose a meta-game model to analyse complex games with meta-strategies. The meta-game primarily focuses on the strategic decisions made by players within the actual game. Given that, in practice, traders have an infinitely large strategy set, our aim is to investigate herding behavior in this context. To achieve this, we employ the meta-game framework to examine how herding influences various investor types in the market, considering different social connections. Within this framework, we define three distinct meta-strategies: “Imitation,” “Neutral,” and “Confident.” These meta-strategies serve to streamline the game’s complexity. The empirical findings indicate that, in order-driven markets, herding tendencies are more pronounced among long-term investors than their short-term counterparts. Interestingly, under equivalent connectivity conditions, a higher proportion of agents within the meta-game opt for the ‘neutral’ meta-strategy compared to the previous game. This suggests a shift in strategic decision-making patterns when considering the meta-game perspective.

Chapter 6

Conclusions and Future Work

The final chapter provides a summary of my research work, outlining important research contributions. In addition, this chapter also discusses future work.

6.1 Research Summary

My PhD research focused on opinion formation and herding in financial markets. We see our work as the introduction of a framework that can be used to study more questions about herding in financial markets. Opinion formation and herding are important in finance because they can influence market behavior, investment decisions, and the overall health of the financial system. Research into them can contribute to the creation of a stable and well-functioning market. The main research objective of this thesis is to explore the opinion formation and herding of participants in different market environments and different social connections. My work focuses on an empirical approach to the above-mentioned problems in finance. Most previous research has focused on agent-based models, as the complexity of financial scenarios poses many difficulties for theoretical analysis. We have adopted a simulation-based approach (EGTA) to generate empirical games, thus making possible the analysis of some of the results in ABM.

6.2 Research Contributions

The significant contributions of my research are briefly outlined as follows:

1. **The impact of social connections between market participants on herdings and market stability (Chapter 3)**

- We have defined a novel ABM to study the extent to which social connections between market participants affect herding and market stability.
- We incorporate classical opinion dynamics in decision making process of traders and use EGTA to calculate the equilibrium of an appropriately defined strategic game. In these games, traders want to maximise their profits whilst balancing private and public information about the asset.
- In the EGTA experiment, we mainly explored the Nash equilibrium of the game in different contact between agents. We also study the relationship between the stability of the market and the strategies chosen by agents in different social graphs.
- Loosely speaking, our results show that herding is the more likely the denser the social graph. Essentially, the herd responds quickly to the actions taken by the stubborn traders.

2. **Opinion Formation and Herding in order-driven market (Chapter 4)**

- We bridge some of the gaps between our simulated market environment and the real market through order-driven markets and different market trading mechanisms.
- We extend the known model to an order-driven market environment. In this market environment, agents strategically choose the level of trust they place in their neighbors based on different social connections in order to make a profit on their transactions.
- We have conducted separate experiments with long-term and short-term investors and compared the results.
- To overcome the limitation of not having a market maker in the market, we have introduced a new order-matching mechanism called spread/price-time priority into the order-driven market. The spread/price-time priority is determined by involving

three indicators (spread, price and time) order execution priorities, thus overcoming these uncoordinated incentives.

- The results of the experiment show that herdings still appear in order-driven markets. Long-term investors are more susceptible to herding than short-term investors. The spread/price time priority mechanism helps to improve market stability and liquidity. However, in order to make it a more efficient mechanism, the weighting of spreads in the mechanism needs to be fine-tuned.

3. Exploring the herding in order-driven markets through a meta-game (Chapter 5)

- To refine our experiments, we enriched the set of strategies available to agents and used a meta-game to simplify the game and explore the impact of herdings in order-driven markets on both short-term and long-term investors.
- We define a set of three simple meta-strategies to reduce the complexity of the game and record the average payoff per agent in a heuristic payoff table.
- Overall, in a similar conclusion to the previous one, long-term investors in the meta-game are more vulnerable to herdings than short-term investors. With the same connection, more agents choose “Neutral” strategy in the meta-game.

6.3 Academic Publications

Some of the work from my Ph.D. research has been published or submitted.

- Chaoran Wang, Ji Qi, Carmine Ventre. ‘Opinion Formation and Herding in Financial Markets’, International Conference on Autonomous Agents and Multiagent Systems, Workshop on Learning with Strategic Agents. 2022.

6.4 Limitations

Agent-Based Modeling (ABM) is a powerful and versatile approach. ABM is very effective in modeling complex systems where individual agents interact with each other and their environment. It provides a bottom-up approach to understanding emergent phenomena. ABM

allows the modeling of agents with different characteristics, behaviors, and decision-making processes, which allows for highly realistic simulations. ABM can account for heterogeneity among agents, which is often critical for capturing the diversity of behaviors and characteristics in real-world systems. While ABM has many advantages, it also has some limitations. Constructing detailed ABMs typically requires large amounts of data and significant computational resources, which can be challenging to obtain and manage. Running ABM simulations can be computationally intensive, so it is necessary to consider the time and resources required for large-scale or long-duration simulations. If the ABM model is complex, validating or calibrating it against real-world data may become difficult, especially if there is limited empirical data available.

The network topology used to model social connections is relatively simple - all agents have the same number of connections, i.e. the underlying network structure is homogeneous. Whether agents would lie or share misinformation was not considered in our experiment and will be part of my future research. The topology of real-world financial markets is very different, e.g., like scale-free networks, nodes may have different numbers of connections. The most important limitation of our game model is the reliance on the number of strategies, as generating payoff tables requires sufficient computational power. Due to computational power limitations, we used three strategy/meta-strategy games in the above study. An increase in the number of strategies in the game may lead to more discoveries, but also to an exponential increase in computing power.

6.5 Future Work

Based on current research on opinion formation and herdings in the financial field, the following three main directions merit further exploration:

- **Refining the agent-based model.** In Chapter 4 and Chapter 5, additional parameters can be changed or added to our agent-based model to restore a realistic financial market environment. Firstly, the model can explore the performance of traders with different ‘risk aversion’ in an order-driven market environment by varying the trading threshold (th_B and th_S). Secondly, in reality, market maker transaction fees and market maker

rebates play an important role in high frequency trading markets and this should also be reflected in the model. Thirdly, the inclusion of public information in the model, such as news, could be considered and it should be used as a new factor influencing the decisions made by participants. Finally, we can calibrate our ABM with some real-world data. In addition, we can also consider different concepts of crashes and bubbles where the length is limited in time, such as flash crashes. And try to limit the rate of crashes and bubbles by certain parameters.

- **In-depth exploration of spread/price-time priority.** In Chapter 4 we have shown that the spread/price-time priority mechanism helps to improve market stability and liquidity, but we need to fine-tune the spread weights in this mechanism. To do this, we can find a specific set of parameter settings by analysing equilibrium market characteristics for different values of β and μ . Under this parameter setting, the spread/price-time priority order matching mechanism ensures both low market volatility and a considerable trading volume.
- **Other analyses and scenarios about the meta-game.** In Chapter 5, we propose a meta-game model to analyse complex games with meta strategies. In the future, we can reverse-engineer strategies based on desired system behavior. We can also increase heterogeneity through social connections between agents by learning from data. This is a widely used technique to deal with large games [112]. In future experiments, we can simulate a more realistic scenario by mixing long-term and short-term investors in the meta-game.
- **Extending symmetric games to asymmetric games.** In our study above, we focus on symmetric games. That is, games in which participants in our experiments have the same set of options and each option has the same payoff. This means that each player can use the same strategy and receive the same reward for each outcome. In a real market environment, however, each player has different options and motivations. This means that our games require more strategic thinking and decision-making skills. Therefore, we can try to extend symmetric games to asymmetric games. Players in

an asymmetric game can have different sets of options and different payoffs for each option. This means that each player has a unique set of strategies available to them and receives a different reward for each outcome.

In summary, in this thesis, we have studied opinion formation and herding in different market environments and social connections through ABM and EGTA. Essentially, in our experiments, the flock reacts quickly to actions taken by recalcitrant traders. We show that the denser the social graph, the more likely the herding is. The experimental results demonstrate that adding spread/price-time priority mechanism to the order-driven market can improve liquidity and stability of the market, but that the spread weights in the order matching mechanism need to be fine-tuned. In addition, an increase in the spread weighting in the order-matching mechanism would make the herding in the market less pronounced.

Our research provides new perspectives for understanding financial markets and developing strategies. Applied to real data and markets, this research helps bridge the gap between theory and practice, enabling informed decision-making, risk management and optimization of investment outcomes in real-world financial environments.

Bibliography

- [1] W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler, “Asset pricing under endogenous expectations in an artificial stock market,” *The economy as an evolving complex system II*, vol. 27, 1996.
- [2] J. D. Farmer and D. Foley, “The economy needs agent-based modelling,” *Nature*, vol. 460, no. 7256, pp. 685–686, 2009.
- [3] L. Wang, K. Ahn, C. Kim, and C. Ha, “Agent-based models in financial market studies,” in *Journal of Physics: Conference Series*, vol. 1039, no. 1. IOP Publishing, 2018, p. 012022.
- [4] A. Mandes and P. Winker, “Complexity and model comparison in agent based modeling of financial markets,” *Journal of Economic Interaction and Coordination*, vol. 12, pp. 469–506, 2017.
- [5] Y. Yoshimura, H. Okuda, and Y. Chen, “A mathematical formulation of order cancellation for the agent-based modelling of financial markets,” *Physica A: Statistical Mechanics and its Applications*, vol. 538, p. 122507, 2020.
- [6] R. L. Axtell and J. D. Farmer, “Agent-based modeling in economics and finance: past, present, and future,” *Journal of Economic Literature*, 2022.
- [7] D. Sornette, “Why stock markets crash, princeton univ,” 2003.
- [8] A. Johansen, D. Sornette *et al.*, “Shocks, crashes and bubbles in financial markets,” *Brussels Economic Review*, vol. 53, no. 2, pp. 201–253, 2010.

- [9] A. M. Van der Veen, “The dutch tulip mania: the social foundations of a financial bubble,” *Department of Government College of William & Mary*, 2012.
- [10] R. F. Bruner and S. D. Carr, “Lessons from the financial crisis of 1907,” *Journal of Applied Corporate Finance*, vol. 19, no. 4, pp. 115–124, 2007.
- [11] J. B. Foster and F. Magdoff, *The great financial crisis: Causes and consequences*. NYU Press, 2009.
- [12] H. Geman, “From lehman to silicon valley bank and beyond: Why are mistakes repeated in the us banking system?” Policy Center for the New South, Tech. Rep., 2023.
- [13] D. Ciuriak, “The silicon valley bank failure: Historical perspectives and knock-on risks,” *Available at SSRN*, 2023.
- [14] K. Pilbeam, *Finance & Financial Markets*. Macmillan International Higher Education, 2018.
- [15] D. Busch, “Mifid ii: Regulating high frequency trading, other forms of algorithmic trading and direct electronic market access,” *Law and Financial Markets Review*, vol. 10, no. 2, pp. 72–82, 2016.
- [16] H. L. Vogel, *Financial market bubbles and crashes: features, causes, and effects*. Springer, 2018.
- [17] S. Bikhchandani and S. Sharma, “Herd behavior in financial markets,” *IMF Staff papers*, vol. 47, no. 3, pp. 279–310, 2000.
- [18] A. Park and H. Sabourian, “Herding and contrarian behavior in financial markets,” *Econometrica*, vol. 79, no. 4, pp. 973–1026, 2011.
- [19] S. Spyrou, “Herding in financial markets: a review of the literature,” *Review of Behavioral Finance*, vol. 5, no. 2, pp. 175–194, 2013.
- [20] S. Barde, “Direct comparison of agent-based models of herding in financial markets,” *Journal of Economic Dynamics and Control*, vol. 73, pp. 329–353, 2016.

- [21] M. Ouarda, A. El Bouri, and O. Bernard, “Herding behavior under markets condition: Empirical evidence on the european financial markets,” *International Journal of Economics and Financial Issues*, vol. 3, no. 1, pp. 214–228, 2013.
- [22] C. Avery and P. Zemsky, “Multidimensional uncertainty and herd behavior in financial markets,” *American economic review*, pp. 724–748, 1998.
- [23] M. F. Granha, A. L. Vilela, C. Wang, K. P. Nelson, and H. E. Stanley, “Opinion dynamics in financial markets via random networks,” *Proceedings of the National Academy of Sciences*, vol. 119, no. 49, p. e2201573119, 2022.
- [24] V. M. Eguiluz and M. G. Zimmermann, “Transmission of information and herd behavior: an application to financial markets,” *Physical review letters*, vol. 85, no. 26, p. 5659, 2000.
- [25] E. Brinkman, “Understanding financial market behavior through empirical game-theoretic analysis,” Ph.D. dissertation, 2018.
- [26] L. J. Schvartzman and M. P. Wellman, “Stronger cda strategies through empirical game-theoretic analysis and reinforcement learning,” in *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2009, pp. 249–256.
- [27] A. Golub, A. Dupuis, and R. Olson, “High-frequency trading in fx markets,” *High-frequency Trading: New Realities for Traders, Markets and Regulators; Easley, D., Prado, MMLD, O’Hara, M., Eds*, pp. 21–44, 2013.
- [28] M. Buchanan, “Meltdown modelling: could agent-based computer models prevent another financial crisis?” *Nature*, vol. 460, no. 7256, pp. 680–683, 2009.
- [29] M. Yoon and K. Lee, “Agent-based and “history-friendly” models for explaining industrial evolution,” *Evolutionary and Institutional Economics Review*, vol. 6, no. 1, pp. 45–70, 2009.

- [30] R. G. Palmer, W. B. Arthur, J. H. Holland, B. LeBaron, and P. Tayler, “Artificial economic life: a simple model of a stockmarket,” *Physica D: Nonlinear Phenomena*, vol. 75, no. 1-3, pp. 264–274, 1994.
- [31] A. Kopp, R. Westphal, and D. Sornette, “Agent-based model generating stylized facts of fixed income markets,” *Journal of Economic Interaction and Coordination*, vol. 17, no. 4, pp. 947–992, 2022.
- [32] F. Slanina and H. Lavicka, “Analytical results for the sznajd model of opinion formation,” *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 35, no. 2, pp. 279–288, 2003.
- [33] K. Sznajd-Weron and J. Sznajd, “Opinion evolution in closed community,” *International Journal of Modern Physics C*, vol. 11, no. 06, pp. 1157–1165, 2000.
- [34] J. Gaitonde, J. Kleinberg, and E. Tardos, “Adversarial perturbations of opinion dynamics in networks,” *arXiv preprint arXiv:2003.07010*, 2020.
- [35] R. Myerson, “Game theory: Analysis of conflict harvard univ,” *Press, Cambridge*, 1991.
- [36] J. von Neumann, “10.1007/bf01448847,” *Math. Ann*, vol. 100, pp. 295–320, 1928.
- [37] B. Von Stengel, “Game theory basics,” *Lecture Notes, Department of Mathematics, London School of Economics, Houghton St, London WC2A 2AE, United Kingdom*, 2008.
- [38] D. Bloembergen, K. Tuyls, D. Hennes, and M. Kaisers, “Evolutionary dynamics of multi-agent learning: A survey,” *Journal of Artificial Intelligence Research*, vol. 53, pp. 659–697, 2015.
- [39] J. Nash, “Non-cooperative games,” *Annals of mathematics*, pp. 286–295, 1951.
- [40] R. B. Myerson, “Nash equilibrium and the history of economic theory,” *Journal of Economic Literature*, vol. 37, no. 3, pp. 1067–1082, 1999.

- [41] M. P. Wellman, “Methods for empirical game-theoretic analysis,” in *AAAI*, 2006, pp. 1552–1556.
- [42] M. P. Wellman, D. M. Reeves, K. M. Lochner, S.-F. Cheng, and R. Suri, “Approximate strategic reasoning through hierarchical reduction of large symmetric games,” in *AAAI*, vol. 5, 2005, pp. 502–508.
- [43] J. Eriksson, N. Finne, and S. Janson, “Evolution of a supply chain management game for the trading agent competition,” *AI Communications*, vol. 19, no. 1, pp. 1–12, 2006.
- [44] M. P. Wellman and B. Wiedenbeck, “An empirical game-theoretic analysis of credit network formation,” in *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2012, pp. 386–393.
- [45] M. H. DeGroot, “Reaching a consensus,” *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [46] D. Bindel, J. Kleinberg, and S. Oren, “How bad is forming your own opinion?” *Games and Economic Behavior*, vol. 92, pp. 248–265, 2015.
- [47] M. Ye, J. Liu, B. D. Anderson, C. Yu, and T. Başar, “On the analysis of the degroot-friedkin model with dynamic relative interaction matrices,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11 902–11 907, 2017.
- [48] D. Ferraioli, P. W. Goldberg, and C. Ventre, “Decentralized dynamics for finite opinion games,” *Theoretical Computer Science*, vol. 648, pp. 96–115, 2016.
- [49] A. Di Mare and V. Latora, “Opinion formation models based on game theory,” *International Journal of Modern Physics C*, vol. 18, no. 09, pp. 1377–1395, 2007.
- [50] M. Epitropou, D. Fotakis, M. Hofer, and S. Skoulakis, “Opinion formation games with aggregation and negative influence,” *Theory of Computing Systems*, vol. 63, no. 7, pp. 1531–1553, 2019.

- [51] K. Bhawalkar, S. Gollapudi, and K. Munagala, "Coevolutionary opinion formation games," in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 2013, pp. 41–50.
- [52] M. Wyart and J.-P. Bouchaud, "Self-referential behaviour, overreaction and conventions in financial markets," *Journal of Economic Behavior & Organization*, vol. 63, no. 1, pp. 1–24, 2007.
- [53] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.
- [54] G. HARRAS and D. Sornette, "Endogenous versus exogenous origins of financial rallies and crashes in an agent-based model with bayesian learning and imitation," *Swiss Finance Institute Research Paper*, no. 08-16, 2008.
- [55] O. J. Blanchard, "Speculative bubbles, crashes and rational expectations," *Economics letters*, vol. 3, no. 4, pp. 387–389, 1979.
- [56] O. J. Blanchard and M. W. Watson, "Bubbles, rational expectations and financial markets," National Bureau of economic research, Tech. Rep., 1982.
- [57] A. Johansen, D. Sornette, and O. Ledoit, "Predicting financial crashes using discrete scale invariance," *arXiv preprint cond-mat/9903321*, 1999.
- [58] D. Sornette and J. V. Andersen, "A nonlinear super-exponential rational model of speculative financial bubbles," *International Journal of Modern Physics C*, vol. 13, no. 02, pp. 171–187, 2002.
- [59] V. Lei, C. N. Noussair, and C. R. Plott, "Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality," *Econometrica*, vol. 69, no. 4, pp. 831–859, 2001.
- [60] S. S. Levine and E. J. Zajac, "The institutional nature of price bubbles," *Available at SSRN 960178*, 2007.

- [61] S. Lee and K. Lee, "Heterogeneous expectations leading to bubbles and crashes in asset markets: Tipping point, herding behavior and group effect in an agent-based model," *Journal of Open Innovation: Technology, Market, and Complexity*, vol. 1, no. 1, p. 12, 2015.
- [62] I. Giardina and J.-P. Bouchaud, "Bubbles, crashes and intermittency in agent based market models," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 31, no. 3, pp. 421–437, 2003.
- [63] S. J. Leal, M. Napoletano, A. Roventini, and G. Fagiolo, "Rock around the clock: An agent-based model of low-and high-frequency trading," *Journal of Evolutionary Economics*, vol. 26, no. 1, pp. 49–76, 2016.
- [64] D. Sornette, "Critical market crashes," *Physics reports*, vol. 378, no. 1, pp. 1–98, 2003.
- [65] J. Paulin, A. Calinescu, and M. Wooldridge, "Understanding flash crash contagion and systemic risk: A micro–macro agent-based approach," *Journal of Economic Dynamics and Control*, vol. 100, pp. 200–229, 2019.
- [66] M. Magdon-Ismail and A. F. Atiya, "Maximum drawdown," *Risk Magazine*, vol. 17, no. 10, pp. 99–102, 2004.
- [67] M. Magdon-Ismail, A. Atiya, A. Pratap, and Y. Abu-Mostafa, "The maximum draw-down of the brownian motion," in *2003 IEEE International Conference on Computational Intelligence for Financial Engineering, 2003. Proceedings.* IEEE, 2003, pp. 243–247.
- [68] R. P. C. Leal and B. V. de Melo Mendes, "Maximum drawdown: Models and applications," *The Journal of Alternative Investments*, vol. 7, no. 4, pp. 83–91, 2005.
- [69] H. Zhang and O. Hadjiliadis, "Drawdowns and the speed of market crash," *Methodology and Computing in Applied Probability*, vol. 14, no. 3, pp. 739–752, 2012.

- [70] G. Rotundo and M. Navarra, “On the maximum drawdown during speculative bubbles,” *Physica A: Statistical Mechanics and its Applications*, vol. 382, no. 1, pp. 235–246, 2007.
- [71] D. Braha, “Global civil unrest: contagion, self-organization, and prediction,” *PloS one*, vol. 7, no. 10, p. e48596, 2012.
- [72] F. Della Rossa, L. Giannini, and P. DeLellis, “Herding or wisdom of the crowd? controlling efficiency in a partially rational financial market,” *PloS one*, vol. 15, no. 9, p. e0239132, 2020.
- [73] J. B. De Long, A. Shleifer, L. H. Summers, and R. J. Waldmann, “Noise trader risk in financial markets,” *Journal of political Economy*, vol. 98, no. 4, pp. 703–738, 1990.
- [74] T. C. Chiang and D. Zheng, “An empirical analysis of herd behavior in global stock markets,” *Journal of Banking & Finance*, vol. 34, no. 8, pp. 1911–1921, 2010.
- [75] M. K. Brunnermeier and M. K. Brunnermeier, *Asset pricing under asymmetric information: Bubbles, crashes, technical analysis, and herding*. Oxford University Press on Demand, 2001.
- [76] G. Tedeschi, G. Iori, and M. Gallegati, “Herding effects in order driven markets: The rise and fall of gurus,” *Journal of Economic Behavior & Organization*, vol. 81, no. 1, pp. 82–96, 2012.
- [77] P. V. J. da Gama Silva, M. C. Klotzle, A. C. F. Pinto, and L. L. Gomes, “Herding behavior and contagion in the cryptocurrency market,” *Journal of Behavioral and Experimental Finance*, vol. 22, pp. 41–50, 2019.
- [78] W. Paul and J. Baschnagel, “Stochastic processes,” *From Physics to finance*, 1999.
- [79] J. Voit and R. W. Lourie, “The statistical mechanics of financial markets,” *Physics today*, vol. 55, no. 8, pp. 080 000–52, 2002.
- [80] R. E. Bailey, *The economics of financial markets*. Cambridge University Press, 2005.

- [81] P. Bak, M. Paczuski, and M. Shubik, "Price variations in a stock market with many agents," *Physica A: Statistical Mechanics and its Applications*, vol. 246, no. 3-4, pp. 430–453, 1997.
- [82] S. Maslov, "Simple model of a limit order-driven market," *Physica A: Statistical Mechanics and its Applications*, vol. 278, no. 3-4, pp. 571–578, 2000.
- [83] J. D. Farmer, P. Patelli, and I. I. Zovko, "The predictive power of zero intelligence in financial markets," *Proceedings of the National Academy of Sciences*, vol. 102, no. 6, pp. 2254–2259, 2005.
- [84] D. Friedman, *The double auction market: institutions, theories, and evidence*. Routledge, 2018.
- [85] M. Bartolozzi, "A multi agent model for the limit order book dynamics," *The European Physical Journal B*, vol. 78, no. 2, pp. 265–273, 2010.
- [86] M. Bartolozzi, C. Mellen, F. Chan, D. Oliver, T. Di Matteo, and T. Aste, "Applications of physical methods in high-frequency futures markets," in *Complex Systems II*, vol. 6802. SPIE, 2008, pp. 15–28.
- [87] J. Brogaard, T. Hendershott, and R. Riordan, "High-frequency trading and price discovery," *The Review of Financial Studies*, vol. 27, no. 8, pp. 2267–2306, 2014.
- [88] Z. Nasar and S. W. Jaffry, "Trust-based situation awareness: comparative analysis of agent-based and population-based modeling," *Complexity*, vol. 2018, 2018.
- [89] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *Math. Sociology*, vol. 15, no. 3-4, 1990.
- [90] B. Roehner and D. Sornette, "'thermometers' of speculative frenzy," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 16, no. 4, pp. 729–739, 2000.

- [91] V. Plerou, P. Gopikrishnan, X. Gabaix, and H. E. Stanley, “Quantifying stock-price response to demand fluctuations,” *Physical review E*, vol. 66, no. 2, p. 027104, 2002.
- [92] G. Barlevy and P. Veronesi, “Rational panics and stock market crashes,” *Journal of Economic Theory*, vol. 110, no. 2, pp. 234–263, 2003.
- [93] S. Jacob Leal, M. Napoletano, A. Roventini, and G. Fagiolo, “Rock around the clock: An agent-based model of low-and high-frequency trading,” *Journal of Evolutionary Economics*, vol. 26, pp. 49–76, 2016.
- [94] R. C. Radcliffe, *Investment: concepts, analysis, strategy*. Pearson Scott Foresman, 1990.
- [95] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, “Limit order books,” *Quantitative Finance*, vol. 13, no. 11, pp. 1709–1742, 2013.
- [96] M. Zare, O. Naghshineh Arjmand, E. Salavati, and A. Mohammadpour, “An agent-based model for limit order book: Estimation and simulation,” *International Journal of Finance & Economics*, vol. 26, no. 1, pp. 1112–1121, 2021.
- [97] S. H. Chan, K. A. Kim, and S. G. Rhee, “Price limit performance: Evidence from transactions data and the limit order book,” *Journal of Empirical Finance*, vol. 12, no. 2, pp. 269–290, 2005.
- [98] E. Yudovina, “A simple model of a limit order book,” *arXiv preprint arXiv:1205.7017*, 2012.
- [99] O. Kaya, J. Schildbach, and D. B. Ag, “High-frequency trading,” *Reaching the Limits, Automated Trader Magazine*, vol. 41, pp. 23–27, 2016.
- [100] B. Baldacci, D. Possamaï, and M. Rosenbaum, “Optimal make-take fees in a multi market-maker environment,” *SIAM Journal on Financial Mathematics*, vol. 12, no. 1, pp. 446–486, 2021.

- [101] C. E. Carpio, O. Isengildina-Massa, R. D. Lamie, and S. D. Zapata, “Does e-commerce help agricultural markets? the case of market maker,” *Choices*, vol. 28, no. 316-2016-7695, 2013.
- [102] K. Venkataraman and A. C. Waisburd, “The value of the designated market maker,” *Journal of Financial and Quantitative Analysis*, vol. 42, no. 3, pp. 735–758, 2007.
- [103] E. Theissen and C. Westheide, “Call of duty: Designated market maker participation in call auctions,” *Journal of Financial Markets*, vol. 49, p. 100530, 2020.
- [104] M. Zhu, C. Chiarella, X.-Z. He, and D. Wang, “Does the market maker stabilize the market?” *Physica A: Statistical Mechanics and Its Applications*, vol. 388, no. 15-16, pp. 3164–3180, 2009.
- [105] O. Guéant, “Optimal market making,” *Applied Mathematical Finance*, vol. 24, no. 2, pp. 112–154, 2017.
- [106] L. A. Veraart, “Optimal market making in the foreign exchange market,” *Applied Mathematical Finance*, vol. 17, no. 4, pp. 359–372, 2010.
- [107] X. Gao and Y. Wang, “Optimal market making in the presence of latency,” *Quantitative Finance*, vol. 20, no. 9, pp. 1495–1512, 2020.
- [108] K. Huang, D. Simchi-Levi, and M. Song, “Optimal market-making with risk aversion,” *Operations research*, vol. 60, no. 3, pp. 541–565, 2012.
- [109] K. Tuyls, J. Perolat, M. Lanctot, J. Z. Leibo, and T. Graepel, “A generalised method for empirical game theoretic analysis,” *arXiv preprint arXiv:1803.06376*, 2018.
- [110] A. Hefti, “Equilibria in symmetric games: Theory and applications,” *Theoretical Economics*, vol. 12, no. 3, pp. 979–1002, 2017.
- [111] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart, “Analyzing complex strategic interactions in multi-agent systems,” in *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002, pp. 109–118.

- [112] B. Nguyen, R. W. Faff, and M. Haq, "Pitching research lite: A reverse-engineering strategy for finding a new research direction," *Available at SSRN 2909549*, 2017.