Opinion Formation and Herding in Financial Markets

Chaoran Wang
King's College London
London, United Kingdom
k19076783@kcl.ac.uk

Ji Qi
King's College London
London, United Kingdom
qi.ji@kcl.ac.uk

Carmine Ventre
King's College London
London, United Kingdom
carmine,ventre@kcl.ac.uk

ABSTRACT

This paper focuses on opinion formation and herding effects in financial markets. We adopt a well-established opinion diffusion dynamics to model social connections between the traders in a hypothetical market environment. Opinions are translated to trading positions and market prices evolve accordingly. We relate the shape of the graph social network with the equilibria of the game, where traders can strategically decide whether to follow the wisdom of the crowd or act upon their own beliefs. We adopt Empirical Game-Theoretic Analysis to compute the Nash equilibrium of this game. We show that the larger their neighbourhood, 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.

KEYWORDS

Opinion formation; Herding effects; Empirical Game-Theoretic Analysis; Crashes and bubbles

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 equilibrium, and large price changes are simply the result of markets reacting to new external information or changing fundamentals. Thus, according to EMH, there is no room for asset prices to bubble or collapse [1]. However, history and past data show that 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 significant jumps 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 [7].

In this paper, 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 [9] building upon the seminal work of DeGroot [5], and studied extensively by the community working on incentives in multi-agent system, see, e.g., [3, 8]. 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 dynamics we want to study. To capture this, we link the opinion dynamics with trading decisions, which naturally define a 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 adopt Empirical Game-Theoretic Analysis to compute the equilibria of the game, where agents decide at each step whether to follow or steer, for a variety of social graphs. We set a parameter similar to the degree of 'stubbornness' of agents. A high level of agent 'stubbornness' means that the agent is more likely to believe in itself. 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. In absence of quantitative notions, we introduce a new definition of crashes and bubbles, inspired by the financial concept of Maximum Drawdown [11]. 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 few bumps along the road.

The paper unfolds as follows: Section 2 presents the literature review about this topic. Section 3 describes the design and general setup of the experiment. The main experimental results and analyses are presented in Section 4. The final section concludes our findings from the experiments.

2 LITERATURE REVIEW

Agent-based models (ABM) are a class of computational models that simulate the behavior and interactions of autonomous agents in order to assess their impact on the system as a whole. It combines elements of game theory, complex systems, computational sociology, and multi-agent systems. ABMs are microscopic models that simulate the simultaneous actions and interactions of multiple agents in an attempt to recreate and predict the emergence of complex phenomena. They are computational method that allows analysts to create, analyze, and experiment with artificial worlds composed of agents that interact in a particular environment [4]. ABM is applicable to the analysis of financial markets formed by heterogeneous entities and their interactions, as well as to the study of economic phenomena such as financial crises that are difficult to explain with mainstream economic tools. Some previous studies have analyzed financial markets based on asset management. One

Appears at the 1st Workshop on Learning with Strategic Agents (LSA 2022). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), N. Bishop, M. Han, L. Tran-Thanh, H. Xu, H. Zhang (chairs), May 9–10, 2022, Online.

of the earliest representative models is the Artificial Stock Market [12]. The authors developed a simple model of the stock market and showed that price bubbles, crashes and persistently high trading volumes could occur.

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. F. Slanina and H. Lavicka describe the Sznajd model of opinion formation and social influence [17]. The Sznajd model is an economic physics model proposed in 2000. Implementing a phenomenon called social validation [18]. 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. In a multi-person environment it has been pointed out that, in order to be realistic, the concept of a solution should be obtained through a decentralised, rational and simple process. Therefore, many authors investigated the concept of computing solutions through decentralised dynamics, see for example [8]. In these algorithms, participants take turns to reduce their costs in the hope that the system will reach equilibrium quickly. They note that finite opinion games are potential games [15]. That means these games admit pure Nash equilibria. They then studied the decentralised dynamics of finite opinion games, in which players can actively update their opinions without the intervention of a central authority. Morris considers a group of individuals who must act together as a team and assumes that each individual in the group has his or her own subjective probability distribution for an unknown value of a parameter [5]. He proposes a model which describes how the group can agree on a common subjective probability distribution for the parameter by collecting their individual opinions. He argues that the model can be applied to the problem of reaching consensus.

Wyart and Bouchaud study a general model of self-referencing behavior in financial markets [25]. But their model does not take into account the effects of imitation through social networks. An 'information cascade' occurs when a person trusts its neighbors' information too much and does not consider its own [2]. 'Information cascade' is a phenomenon described in behavioral economics and network theory, where many people make the same decision in a continuous manner. Georges Harras 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 [10]. There are three sources of information in the paper: public information (news), information from their neighbors' network 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. 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, i.e., the social graph is a grid. In their conclusion, they argue that the resulting market is approximately efficient when an agent has little faith in its neighbors. They think that when the agent is more trusting of its neighbor the market is more prone to crash and bubble. We will be applying our experiments to more neighbours for each agent and more type of social graphs. We focus not only on the detailed dynamics of the collapse and the development of the bubble but also on opinion formation and herding effects in financial markets.

Maximum Drawdown (MDD)Maximum Drawdown (MDD) is defined as the maximum loss that occurs from peak to trough over a specified time period. A sharp contraction may even indicate that an already successful trading system is deteriorating due to changes in the market regime. In Hongzhong Zhang and Olympia Hadjiliadis' paper, they study the probabilistic behavior of two quantities that are closely related to market crashes [26]. 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 [16]. 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.

Empirical game theory analysis is an emerging empirical methodology that bridges the gap of game theory and simulation in practical strategic reasoning [22]. 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 [6] 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 [22] EGTA is broken down into three basic steps which are: Parametrize Strategy Space, Estimate Empirical Game and Analyze Empirical Game. This approach to has been applied a variety of games, especially market-based scenarios, see, e.g., [19–21]. 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 [24].

3 GENERAL SETUP

In our model, public news is included in private information. More neighbors of the agent can make the agent's opinion formation more realistic and allow a more complete study of herd behavior. We set up a fixed number of $N = 100$ agents who are buying and selling a stock. We follow [10] and model 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. Therefore, we limit the authenticity of the model to a simple set of components.

The information that underlies an agent's decision is limited to two sources (neighbors and themselves) and the payoffs that have been realized. We let T denote the trading horizon and set $T = 10,000$ in our experiments. We denote the opinion for agent i at time step t, as $s_i(t) \in [0,1]$. The opinion will define the trading position. The graph $G = (V,E), V = N$, is undirected with $E = \{(i,j)|i,j \in N\}$. 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.

3.1 Two sources of information

The first source of information is private and 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. $b_i(t)$ is the belief of agent i about the stock at time t. It reflects each agent's particular subjective view of the stock's future performance. Intuitively, the bigger $b_i(t)$ the more the agent believes that the stock will increase in value (and then needs 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. The second source of information of agent i is provided by the past decisions of other agents at time t, $s_j(t-1)$, for $j \in N_i$ and $t \leq T$. With limited access to information and limited bounded rationality, some argue that imitating others is optimal [14]. 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 decision.

3.2 Opinion formation

The cost $C(s_i(t))$ of the opinion $s_i(t)$ for agent i at time step t 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 more reasonable it is. We follow [8] and define:

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

As we can see in (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 means that how much agents trust their neighbors, $\alpha \in [0,1]$. A bigger α represents more trusting other agents, 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 that we call 'stubbornness' in the following. 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. As discussed in the introduction, this model of opinion dynamics is well motivated and rooted in economics and sociology [5].

3.3 Trading decision

3.3.1 *Buy or sell.*

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 1.

Figure 1: Trading decision

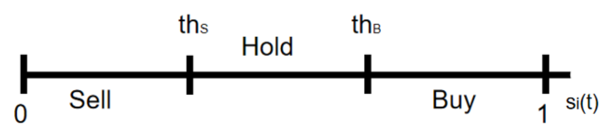

3.3.2 *Trading size.*

The assets of agent i consist of the amount of cash it holds, $cash_i(t)$, and the number of single stock 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 defined by:

$$A_i(t) = g \cdot \frac{cash_i(t)}{price(t)} \quad (2)$$

$$B_i(t) = +1(buying) \quad (3)$$

if $s_i(t) \geq th_B$, and

$$A_i(t) = g \cdot stock_i(t) \quad (4)$$

$$B_i(t) = -1(selling) \quad (5)$$

if $s_i(t) \leq th_S$.

3.4 Price clearing

When all the agents have taken the position, a new price is created. We denote $r(t)$ 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 seems to make sense for markets with many high frequency traders, providing liquidity to the market. The price is then determined by:

$$r(t) = \frac{1}{\gamma \cdot N}\sum_{i=1}^{N} A_i(t) \cdot B_i(t) \quad (6)$$

$$log[price(t+1)] = log[price(t)] + r(t) \quad (7)$$

In (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 [13]. This is similar to [10].

3.5 Cash and stock

As from above, we assume a friction-less market that has no transaction costs. When the return and new price of the stock are determined by (6) and (7), the amount of cash and stock held by agent i

will be determined by the following formulas:

$$cash_i(t) = cash_i(t-1) - A_i(t) \cdot price(t) \cdot B_i(t) \tag{8}$$

$$stock_i(t) = stock_i(t-1) + A_i(t) \cdot B_i(t) \tag{9}$$

3.6 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 change depending on the market environment we are studying. When the stock price trend downwards, δ_{in} is the percentage decline in stock price, used to determine the start 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 where t_{in} is the start point of the crash and t_{out} is the end point of the crash. 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 2. Bubbles are define similarly in Algorithm 2. When $\delta_3 > \delta_{in}$ and $\delta_4 > \delta_{out}$, there is a bubble, see Figure 2.

Algorithm 1: Definition of crash

```
Let δ1 = (P_max − P_min)/P_max, δ2 = (P_t − P_min)/P_min
while δ1 > δ_in do
    t_in ← t_max
    if δ2 > δ_out then
        t_out ← t_min
    end
end
Crash in [t_in, t_out]
Set P_max ← P_min
```

Algorithm 2: Definition of bubble

```
Let δ3 = (P_max − P_min)/P_min, δ4 = (P_max − P_t)/P_max
while δ3 > δ_in do
    t_in ← t_min
    if δ4 > δ_out then
        t_out ← t_max
    end
end
Bubble in [t_in, t_out]
Set P_min ← P_max
```

In our experiment, δ_1, δ_2, δ_3 and δ_4 are the proportion of changes in stock prices. Changing the value of δ_{in} and δ_{out} allows more or less sensitivity to market movements. For example, large value of δ_{in} and small value of δ_{out} focus on 'big' systematic events that are of general interest. Conversely, small value of δ_{in} and large value of δ_{out} are more suitable for industry professional and high frequency traders.

Figure 2: Illustration of the definitions of crash and bubble

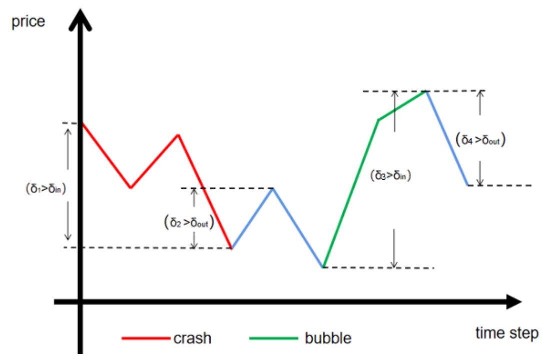

3.7 Three Strategies

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.

3.8 Payoff

The amount of cash and stock held by agent i will be determined by (8) and (9). So the total assets of agent i at time step t will be determined by the following formula:

$$Z_i(t) = cash_i(t) + stock_i(t) \cdot price(t) \tag{10}$$

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

$$\begin{aligned} R_i(t) &= (Z_i(1) - Z_i(0)) + \cdots + (Z_i(t) - Z_i(t-1)) \\ &= Z_i(t) - Z_i(0) \end{aligned} \tag{11}$$

In general our experiment has two steps of decision making at each time step. First agents form their opinion by a function of 'stubbornness' (1). Then they make their decision by the trading thresholds that we mentioned in section 3.5. The dynamics of the market are updated according to (6) and (7). Finally we get the payoff for each agent at this time step by (10) and (11).

4 EXPERIMENTS AND RESULTS

We built a model with a closed system. Our agent will invest a certain proportion of its 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. But 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 his or her own choice. So we define a game and in this game, each agent is considered to be a player. 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. But simulating all the circumstances is a very large game. It requires a very large amount of computation. A simulation-based approach to finding an approximate equilibrium was proposed by Wellman [22]. In this method, strategy profiles and corresponding payoffs are extracted from the payoff distribution by means of extensive systematic simulations. The payoff data is used to induce a normal-form game. In this way, we can refine the estimation of our game and get an approximate equilibrium rather than analysing the complete game. The equilibrium we obtain in the empirical game is considered likely to be relatively stable in a full game [23]. After we apply EGTA to solve this game, we then go through the simulation again to obtain the market performance at equilibrium.

4.1 Set up parameters

4.1.1 Basic parameters. Our experiments are as consistent as possible with the basic parameters given in [10] to allow a fair comparison. 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 $g = 2\%$. 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. This does not mean that the agents change their mind at every time step. The agents just get different information about the stock at every time step.

4.1.2 Trading threshold. Because the agents in our experiment had 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 stock unchanged.

4.1.3 δ_{in} and δ_{out}. As mentioned above δ_{in} and δ_{out} determine the market crash and bubble's in point and out point. In our experiment we set $\delta_{in} = 5\%$, $\delta_{out} = 2.5\%$.

4.1.4 Choice of the value of 'stubbornness'. These values are based on the results of preliminary experiments with agent-based model. When $\alpha = 0$, the 'stubbornness' of agent i reach maximum. That means agent i only trust itself, $s_i(t)$ depends entirely on its private information $b_i(t)$. The first strategy is 'No imitation' and it represents $\alpha = 0$. When $\alpha > 0.7$, few agents trade in the market. The price of the stock remains almost unchanged. 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, 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 third strategy, called 'Neutral'.

Figure 3: Dynamics for cliques

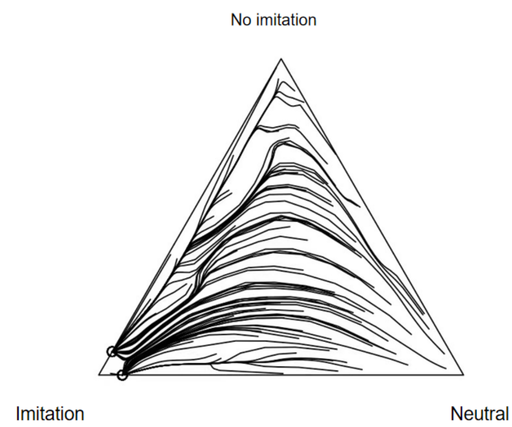

4.2 Results

We run 100 simulations and then we take the average of the prices at each time step. Recall that $T = 10,000$ as in [10]. Our experiment uses different social graphs. We can change the way agents contact each other so that the information they get about their neighbors changes. To make the experiment fairer and more convincing, we used the same $b_i(t)$'s for each different social graph. In our experiments, each agent can choose one of the three strategies per experiment. At the end of each experiment, according to (10) and (11), 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 returns. We then process the data analytically and find the Nash equilibrium by EGTA. We will then observe the crashes and bubbles that occur in markets at these Nash equilibrium points.

Figure 3 shows the Nash equilibrium that we obtained after experimentation for a clique. We use a two-dimensional triangle to represent a three-dimensional space. Any point in the triangle represents a mixed strategy. As an example, the apex of the triangle represents 100% of the agents choosing 'No imitation' strategy. The hollow dot in the figure which the trajectories converge to represents the stable state, i.e., a Nash equilibrium point. We choose the the equilibrium point where the majority of trajectories converge to. At this point we have that 92% of the agents choose the 'Imitation' strategy, 0% of the agents choose the 'Neutral' strategy and the remaining 8% of the agents choose the 'No imitation' strategy. With this equilibrium point, we obtain the results shown in Figure 4; 3 crashes and 4 bubbles occur in the market.

Figure 4: Crashes and Bubbles at equilibrium for cliques

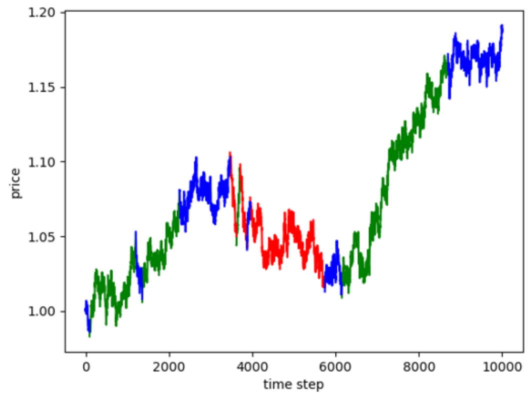

4.3 Changing the graph

In the above experiment, we set up 100 agents that are connected to each other. This means that each agent has 99 neighbors. To better understand opinion formation and herding effects in different social graphs, we change the structure of the social connections between agents. We obtain the Nash equilibrium points as shown in the figures below.

In Figure 5, 'x-a-group' represents a graph where each group has x agents and each agent has $x - 1$ neighbours. The graph is a ring where agent i has en edge with the nearest $x - 1$ agents. The results show that with each agent getting more of their neighbors' information, the Nash equilibrium point with the majority of trajectories converge from the top of the triangle to the bottom left corner of the triangle. When 90 agents are in a group ('90-a-group'), all trends are essentially concentrated at a Nash equilibrium point. Similarly to the case of clique for analysis we choose the Nash equilibrium point where the majority of trajectories converge in each graph.

Figure 6 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' strategy increases significantly.

At '30-a-group', as shown in Figure 7, 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'.

As we can see in Figure 8, 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 9 shows the agent's choice of strategy at '90-a-group' and '100-a-group'. In Figure 9 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 6 to Figure 9, 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.

Table 1: Crashes and Bubbles at equilibrium

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

Based on the data at Nash equilibrium of different social graph, we record the number of crashes and bubbles that have occurred in the market in Table 1. The data in Table 1 shows that the more neighbor information agents get the fewer crashes and bubbles there will be in the market. And 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 '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 '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 occur in the market at '100-a-group'.

As for the singularity of the '90-a-group', as we stated in the paper (just below Figure 9), when each agent is essentially endowed with progressively more information from '30-a-group' to '80-a-group' as shown in Figure 6 to Figure 8, 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 effect is maximized for

Figure 5: Nash Equilibrium of different graph

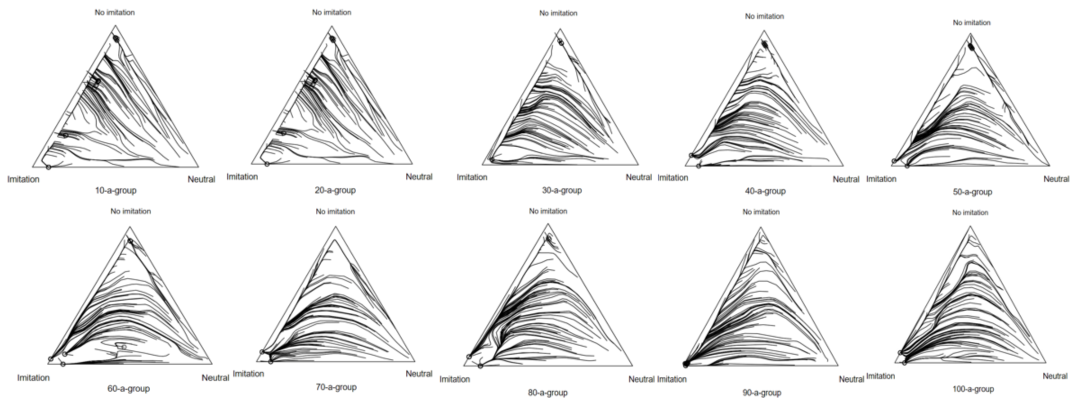

Figure 6: 10-a-group and 20-a-group

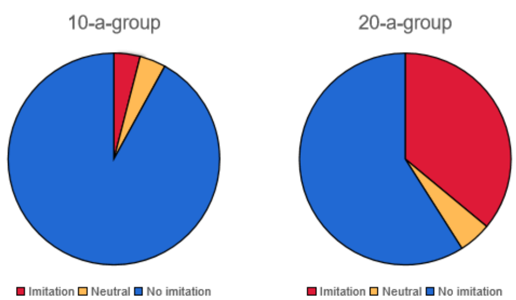

Figure 7: 30-a-group and 40-a-group

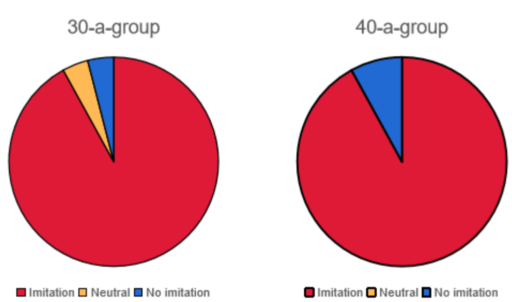

Figure 8: 50-a-group and 80-a-group

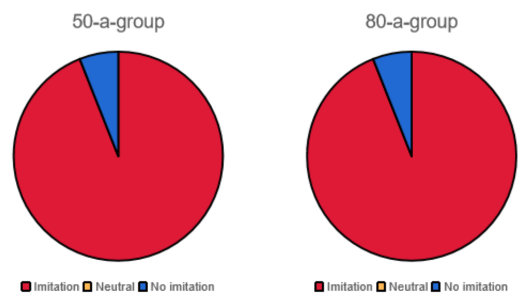

Figure 9: 90-a-group and 100-a-group

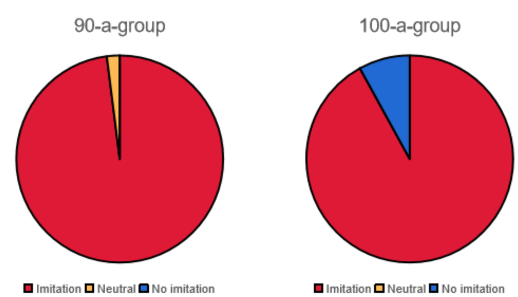

this setting. Basically, all agents in the market choose the 'Imitation' strategy, but few prefer to keep a positive weight for their belief. This changes in the '100-a-group', where every agent is given the information of the other 99 agents. In this full information setup, some agents remain stubborn.

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'.

4.4 Changing risk attitude

As from Section 3.3, 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 that there are 4% agents choosing 'Imitation' strategy, 4% agent choosing 'Neutral' strategy and 92% agents choosing 'No imitation' strategy. This means that when the trading risk is high enough, most agents will not want to imitate their neighbours no matter how much neighbors' information they are given. The number of crashes and bubbles that occur in the market at any given social graphs remains virtually unchanged. This is because when the vast majority of agents choose '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 graph. We got 8 crashes and 4 bubbles in this experiment, which is higher than the data in Table 1. This confirms that high risk makes markets unstable.

Then we try to 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 4.3. They have the same pattern which is at '10-a-group' and '20-a-group', the majority of agents in the market choose '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 'Imitation' strategy and 6% agents choose 'No imitation' strategy. The result of '90-a-group' and '100-a-group' is similar to Figure 9. 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 4.3, especially the singularity of the '90-a-group'.

4.5 About the trading size

In our experiment, the fraction of their cash or stocks that investors trade per action g as mentioned in section 3.3 is the same for all agents following [10] and $g = 2\%$. When we change the g to 1%, crashes and bubbles in the market have become less frequent. When we try to slowly increase g to 5%, the number of crashes and bubbles will increase with g. Then we random g from 1% to 5% for each agent of different graph and get the Nash equilibrium. The results are nearly the same as the Figure 5. That means the trading size g only change the stability of the market and can not affect the Nash equilibrium of different social graph.

5 CONCLUSIONS

We have defined a novel ABM 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; with our definition of crashes and bubbles we show that herding leads to low volatility unless the social network is a clique. In this case, the small minority of traders who are not part of the herd seem to cause the market to jump a bit; 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 questions 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 continuous double auction market or different market trading mechanisms. We could also consider different notions of crashes and bubbles, where the length is limited in time as in flash crashes. Moreover, other independent variables, social network structures and richer strategy spaces can be used in follow-up experiments. To complement our results with risk-seeking traders, one could imagine to consider risk aversion as well as mixed populations of traders (both by risk attitude and maybe investment style).

REFERENCES

[1] W Brian Arthur, John H Holland, Blake LeBaron, Richard Palmer, and Paul Tayler. 1996. Asset pricing under endogenous expectations in an artificial stock market. *The economy as an evolving complex system II* 27 (1996).

[2] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy* 100, 5 (1992), 992–1026.

[3] David Bindel, Jon Kleinberg, and Sigal Oren. 2015. How bad is forming your own opinion? *Games and Economic Behavior* 92 (2015), 248–265.

[4] Mark Buchanan. 2009. Meltdown modelling: could agent-based computer models prevent another financial crisis? *Nature* 460, 7256 (2009), 680–683.

[5] Morris H DeGroot. 1974. Reaching a consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121.

[6] Joakim Eriksson, Niclas Finne, and Sverker Janson. 2006. Evolution of a supply chain management game for the trading agent competition. *AI Communications* 19, 1 (2006), 1–12.

[7] J Doyne Farmer and Duncan Foley. 2009. The economy needs agent-based modelling. *Nature* 460, 7256 (2009), 685–686.

[8] Diodato Ferraioli, Paul W Goldberg, and Carmine Ventre. 2016. Decentralized dynamics for finite opinion games. *Theoretical Computer Science* 648 (2016), 96–115.

[9] N. E. Friedkin and E. C. Johnsen. 1990. Social influence and opinions. *Math. Sociology* 15, 3-4 (1990).

[10] Georges Harras and Didier Sornette. 2008. Endogenous versus exogenous origins of financial rallies and crashes in an agent-based model with Bayesian learning and imitation. *Swiss Finance Institute Research Paper* 08-16 (2008).

[11] Malik Magdon-Ismail and Amir F Atiya. 2004. Maximum drawdown. *Risk Magazine* 17, 10 (2004), 99–102.

[12] Richard G Palmer, W Brian Arthur, John H Holland, Blake LeBaron, and Paul Tayler. 1994. Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear Phenomena* 75, 1-3 (1994), 264–274.

[13] Vasiliki Plerou, Parameswaran Gopikrishnan, Xavier Gabaix, and H Eugene Stanley. 2002. Quantifying stock-price response to demand fluctuations. *Physical review E* 66, 2 (2002), 027104.

[14] BM Roehner and D Sornette. 2000. "Thermometers" of speculative frenzy. *The European Physical Journal B-Condensed Matter and Complex Systems* 16, 4 (2000), 729–739.

[15] Robert W Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 1 (1973), 65–67.

[16] Giulia Rotundo and Mauro Navarra. 2007. On the maximum drawdown during speculative bubbles. *Physica A: Statistical Mechanics and its Applications* 382, 1 (2007), 235–246.

[17] Frantisek Slanina and Hynek Lavicka. 2003. Analytical results for the Sznajd model of opinion formation. *The European Physical Journal B-Condensed Matter and Complex Systems* 35, 2 (2003), 279–288.

[18] Katarzyna Sznajd-Weron and Jozef Sznajd. 2000. Opinion evolution in closed community. *International Journal of Modern Physics C* 11, 06 (2000), 1157–1165.

[19] Elaine Wah and Michael P Wellman. 2013. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 855–872.

[20] Elaine Wah and Michael P Wellman. 2016. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Finance* 5, 3-4 (2016), 69–93.

[21] Elaine Wah, Mason Wright, and Michael P Wellman. 2017. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research* 59 (2017), 613–650.

[22] Michael P Wellman. 2006. Methods for empirical game-theoretic analysis. In *AAAI*. 1552–1556.

[23] Michael P Wellman, Daniel M Reeves, Kevin M Lochner, Shih-Fen Cheng, and Rahul Suri. 2005. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *AAAI*, Vol. 5. 502–508.

[24] Michael P Wellman and Bryce Wiedenbeck. 2012. An empirical game-theoretic analysis of credit network formation. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 386–393.

[25] Matthieu Wyart and Jean-Philippe Bouchaud. 2007. Self-referential behaviour, overreaction and conventions in financial markets. *Journal of Economic Behavior & Organization* 63, 1 (2007), 1–24.

[26] Hongzhong Zhang and Olympia Hadjiliadis. 2012. Drawdowns and the speed of market crash. *Methodology and Computing in Applied Probability* 14, 3 (2012), 739–752.