Exploring the Impact of Competing Narratives on Financial Markets I: An Opinionated Trader Agent-Based Model as a Practical Testbed

Arwa Bokhari[1,2,a]
[1]Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.
[2]Information Technology Department, College of Computers and Information Technology, Taif University, Saudi Arabia

Keywords: Agent-Based Models, Narrative Economics, Opinion Dynamics, Financial Markets.

Abstract: This paper introduces a framework to empirically investigate the influence of competing narratives on financial market dynamics. We present an agent-based model of traders in a financial market, where traders are driven by opinion dynamics and are subject to self-reinforcement, herding behaviors, and an accumulative response to new information. Our systematic approach includes isolating these factors, enabling a parametric analysis within the collective opinion dynamics of the market. Our simulation provides a testbed to evaluate various market scenarios. While our findings are based on simulated data and thus warrant caution in real-world interpretation, they offer important insights into market fluctuations. This study lays groundwork for further research on trader behavior and market dynamics, and we have made the source-code publicly available for replication and extension.

1 INTRODUCTION

With the advent of instantaneous information dissemination, narratives play a significant role in shaping market dynamics (Shiller, 2019; Hirshleifer, 2020). Narrative economics, an emerging field championed by Nobel laureate Robert Shiller (Shiller, 2017; Shiller, 2019), investigates the influence of prevalent stories on individual economic actions. Rather than merely reporting economic events, these narratives actively shape market behaviors by influencing collective sentiment. Investing in speculative assets is not only a matter of individual psychology; it is fundamentally rooted in social activity (Kim et al., 2023). Individual traders often make investment decisions based on information shared by others. This behavior is recognized by investment experts, institutions, and fund managers (Kim et al., 2023). Shiller notes that stories blending truth and fiction can create uncertainty and varied interpretations. When they go viral, these stories can influence asset prices in ways that deviate from traditional market fundamentals. The impact of such narratives can be amplified by social media, where enthusiastic groups create and spread these stories, wielding significant market power. When such narrative-driven activities are combined with user-friendly trading platforms like Robinhood, stock prices can skyrocket to irrational heights and/or drop significantly. There is a belief that coordinated actions on social media played a role in the fluctuations seen in GameStop price and trading volume in January 2021 (Kim et al., 2023; Jakab, 2022; Aliber et al., 2015).

Narratives' influence on the market is not solely due to their existence but also depends on their interplay. Tesla's stock price rise in 2013 was driven by a prevailing narrative about the promise of a sustainable future powered by electric vehicles. However, a counternarrative cast doubts over Tesla's valuation (Liu, 2021). Similarly, Bitcoin's ascent in 2017 was powered by stories about the era of decentralized currencies but was counterbalanced by narratives pointing to potential misuse and volatility. Conflicting narratives about a company's business prospects can trigger pronounced market fluctuations, as illustrated by the 2021 GameStop episode, and the November 2022 collapse of the crypto-currency exchange FTX.

In this exploratory study, we use group opinion dynamics to construct a testbed model that elucidates the interplay between two competing narratives within social networks. The preliminary findings lay the foundation for subsequent exploration into the intricate dynamics of narrative competition and consensus formation in social networks and their effect on

[a] https://orcid.org/0000-0003-2987-4601

Bokhari, A.
Exploring the Impact of Competing Narratives on Financial Markets I: An Opinionated Trader Agent-Based Model as a Practical Testbed.
DOI: 10.5220/0012429500003636

In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART 2024) - Volume 1*, pages 127-137
ISBN: 978-989-758-680-4; ISSN: 2184-433X

financial markets.

My model examines the concurrent dynamics in a bifurcated financial system, marked by bearish sentiments anticipating market declines on one side and bullish expectations forecasting market rises on the other. We outline three pivotal mechanisms that shape the complex process of financial decision-making and the formation of underlying opinions. These mechanisms are detailed in the following section.

1.1 Key Drivers of Market Dynamics

In understanding financial market dynamics amid competing narratives, it is vital to pinpoint the behavioral and cognitive drivers that influence individual and collective behaviors within the market. In this paper, we report on an agent-based model (ABM) that allows for the exploration of the interplay of three dynamics between two groups of traders: one advocating positive narratives and the other emphasizing negative ones. We first examine the impact of the positive feedback mechanism on collective behavior, a concept from the literature that explains self-reinforcement among social groups. Next, we examine the influence of herding behavior, another well-documented characteristic in social contexts, on collective outcomes. Lastly, we address scenarios where an external factor additively sways group behavior, regardless of internal interactions.

- **Self-Reinforcement.** This behavioral phenomenon occurs when narratives, consistently repeated within a community or environment, amplify and intensify over time. This amplification can lead to escalating confidence in specific beliefs or behaviors, often resulting in a progressively entrenched stance.

  A tangible representation of self-reinforcement is observed within "echo chamber" dynamics. These enclosed environments, prevalent on digital platforms, facilitate the uninterrupted circulation of congruent viewpoints, largely shielded from external challenges or alternative perspectives. Consistent exposure to these conforming opinions within such chambers acts as a recursive feedback mechanism. Each reaffirmation serves to reinforce the pre-existing belief, making it more robust with each iteration.

  Many market phenomena exemplify the self-reinforcing logic in action. As market trends intensify, they can trigger a cascade of investor behavior aligning with the prevailing direction. This positive feedback loop, where market behaviors reinforce and intensify existing trends, further underscores the pervasive nature of self-reinforcement in socio-digital and economic contexts.

- **Herding Behavior.** Herd behavior refers to the tendency of individuals in a group to instinctively mimic each other's actions or beliefs, often influenced by mutual interactions rather than explicit instructions (Kameda et al., 2014). Herd behavior is particularly evident in financial markets, as investors frequently imitate the decisions of others, often presuming that those they follow have done their due diligence.

  The GameStop short squeeze event serves as a prime example of this behavior. Informed investors, such as Keith Gill, a financial advisor from Massachusetts [2], played pivotal roles. In January 2021, Keith Gill's bullish view on the GME stock and his subsequent gains were cited as key factors contributing to the GME short squeeze (Anand and Pathak, 2021). As Gill identified a potential profit opportunity in GameStop, noting its significantly high short interest. Based on this observation, these investors began amassing considerable shares. As the stock's price began to climb, a surge of other investors, motivated more by a fear of missing out than by understanding market intricacies, joined the fray. This initiated a feedback loop, propelling the stock price well beyond GameStop's intrinsic value. Amidst this surge, hedge funds and institutional investors that had shorted the stock felt the heat to buy back at higher prices, intensifying the rise. Numerous subsequent retail investors seemed influenced less by market fundamentals and more by these early participants, underscoring the influence of herd behavior in financial contexts (Andreev et al., 2022).

  Surowiecki (Surowiecki, 2004) pointed out that these market trends can lead investors incorrectly, resulting in irrational bubbles where collective actions drive up asset prices. Even though the "wisdom of crowds" relies on diverse opinions, participants often end up imitating each other, favoring group consensus over individual, independent thought (Kim et al., 2023).

- **Additive Response.** Defined as a direct reaction to an external stimulus or input, an additive response remains independent of external influences, feedback mechanisms, or surrounding circumstances. In the context of financial markets, this response indicates that investors might adjust their positions based purely on these external signals, separate from prevailing market data or col-

[2] See the archived news story here

lective sentiment.

During the GameStop frenzy, this phenomenon was evident. Despite vast unrealized profits, influencer Keith Gill's decision not to sell acted as an external signal, prompting many to adopt the stance, "If he is holding, I am holding". Such behavior wasn't rooted in market fundamentals but was a market dynamic propelled by coordinated additive stimuli. Rallying cries like "diamond hands" and "YOLO" during this period further emphasize the power of these external signals, pushing many to act in ways perhaps contrary to standard market wisdom.

1.2 Structure of the Paper

This paper is the first of two, a pair of papers in which the first (this paper) presents my agent-based model (ABM) in full detail, and the second presents an extension of the model and application to real-world data. To make each paper self-contained, illustrative results are presented in this paper, and similarly, the second paper includes a brief summary of the main aspects of my model. The rest of this paper is organized as follows: Section 2 summarizes the theoretical background; Section 3 outlines the design and functioning of my ABM; Section 4 then presents illustrative results; and Section 5 concludes the discourse.

2 BACKGROUND

Experimental economics, established by the seminal works of (Smith, 1962), has been essential in understanding economic behaviors, particularly within continuous double auction (CDA) markets. In a CDA, participants can submit bids or offers at any time, creating a dynamic environment where trades execute as soon as matching bids and offers are found, without needing a centralized auctioneer. This market mechanism is continuous and asynchronous, allowing for the immediate execution of trades and making it a focal point for economic research (Cliff, 2012). Smith's experiments in 1963 laid the foundation for using controlled laboratory settings to analyze economic principles, contributing to his Nobel Prize in 2002. This empirical approach has significantly informed the development of agent-based computational economics (ACE), characterized by the simulation of market dynamics through autonomous trading agents.

These ABMs frequently implement zero-intelligence (ZI) and minimal-intelligence (MI) trader-agent algorithms. Despite their simplicity, these algorithms effectively replicate complex market behaviors. Prominent ZI trading strategies within the ACE framework include ZIC, SHVR, and GVWY:

- **Zero-Intelligence-Constrained (ZIC) Traders:** Functioning without any forecast of market trends or strategic intricacies, ZIC traders produce random price quotes within a predefined range to avoid loss-making trades, in compliance with their individual constraints (Gode and Sunder, 1993b).
- **Shaver (SHVR) Traders:** SHVR traders incrementally improve upon existing market quotes in a deterministic fashion. A SHVR buyer will issue a bid just above the current best, while a SHVR seller will set an ask just below the lowest present ask, with both adhering to their limit prices (Cliff, 2012).
- **Giveaway (GVWY) Traders:** GVWY traders passively match their quotes to their limit prices, foregoing active market spread exploitation. They can, nonetheless, realize gains if market fluctuations are more favorable than their quoted prices (Cliff, 2012).

In the context of ACE and market simulations, understanding the narratives that drive trader behavior becomes essential. This consideration is particularly relevant in the development of advanced algorithms like (**P**arameterised-**R**esponse **Z**ero **I**ntelligence) PRZI and (**P**arameterized-**R**esponse **D**ifferential **E**volution) PRDE, which are ZI strategies that can adapt to market conditions. Cliff introduced the PRZI (Cliff, 2023) algorithm, enabling adaptive strategy changes in response to market dynamics (Cliff, 2022a). PRZI traders adjust their bids based on a real-valued strategy parameter and are capable of operating as ZIC, SHVR, GVWY, or hybrid strategies. The subsequent PRDE algorithm, as introduced by Cliff (2022), further extends these capabilities through differential evolution optimization (Cliff, 2022b).

Integrating narrative economics into these models offers new insights into the interplay of market mechanics, trader behaviors, and prevailing narratives. The first work in this field is by Lomas and Cliff (Lomas and Cliff, 2021), combining opinion dynamics with ABM to understand how narratives impact financial market prices. However, they did not consider how price dynamics affect market narratives, a gap addressed by (Bokhari and Cliff, 2022) in the BFL-PRDE model: this involved an extension of the **B**izyaeva, **F**ranci, and **L**eonard (BFL) model, building on recent research (Bizyaeva et al., 2020) which integrates continuous-time opinion dynamics with the

PRDE (Cliff, 2022b) strategy. This approach dynamically represents the influence of narratives and opinions on traders' actions and market prices. Traders in the BFL-PRDE model have opinion variables influenced by other agents and market observations, offering a nuanced view of narrative dynamics interplay with market dynamics.

The evolution from PRZI to its successor PRDE, and the advanced BFL-PRDE, are elaborated in the following section.

3 MODEL

3.1 BFL-PRDE Trader Model

In the evolution of trading strategies, PRZI laid the groundwork for adaptive ZI traders, paving the way for its descendant, PRDE. While PRDE equipped traders with means for adapting to market fluctuations, it inherently lacked the capability to anticipate market trends. Traders, operating within the PRDE domain draw upon their specific strategies to determine quote prices. The dynamic nature of these strategies stems from a dual interaction: intrinsic strategy values and the prevailing strategies of peer traders in the market. (Bokhari and Cliff, 2022) extend the PRDE framework by incorporating a real-valued opinion variable, utilizing the opinion dynamics model proposed by Bizyaeva et al., 2020 (Bizyaeva et al., 2020). This integration yields a more sophisticated trading model called (BFL-PRDE) where buyers and sellers, informed by their opinions, demonstrate contrasting market behaviors. Under a bullish consensus, BFL-PRDE buyers, a hybrid of GVWY (Cliff, 2012; Cliff, 2018) and ZIC(Gode and Sunder, 1993a), manifest heightened urgency, influencing their quote prices. Conversely, sellers lean towards a more relaxed position in the form of a hybrid between ZIC and SHVR (Cliff, 2012; Cliff, 2018), especially when bearish sentiments dominate.

This development requires a mapping function, translating trader opinion into its PRDE trading strategy. Elaborating on the intrinsic mechanics, as detailed in (Cliff, 2022a), Each PRDE trader holds a private local set of potential strategy-values with a population size $NP \geq 4$. For trader i, this set can be denoted as $s_{i,1}, s_{i,2}, ..., s_{i,NP}$. Given that PRDE traders rely solely on a singular real scalar value to characterize their bargaining approach, every individual in the differential evolution population represents a single value. Thus, the traditional differential evolution mechanism of crossover (i.e., selecting genes from a pair of parents, one gene for each genome dimension) isn't relevant: PRDE creates a genome exclusively based on the base vector. In its present version, PRDE adopts the standard "vanilla" DE/rand/1 (Storn and Price, 1997). After evaluating a strategy $s_{i,x}$, three distinct s-values are chosen at random from the population: $s_{i,a}, s_{i,b}$, and $s_{i,c}$ ensuring $x \neq a \neq b \neq c$. This results in the generation of a new candidate strategy $s_{i,y}$ defined as $s_{i,y} = \max(\min(s_{i,a} + F_i(s_{i,b} - s_{i,c}), +1), -1)$, where F_i symbolizes the trader's differential weight coefficient (in the outlined experiments, $F_i = 0.8; \forall i$). Utilizing the min and max functions, the candidate strategy's range is limited between $[-1.0, +1.0]$. Within BFL-PRDE, the trader's opinion $s_{i,o}$ emerges as an additional candidate strategy. The performances of $s_{i,y}$ and $s_{i,o}$ are then compared; the superior strategy becomes the new parent strategy $s_{i,x}$. If not, it's replaced with the subsequent strategy $s_{i,x+1}$.

3.2 BFL Opinion Dynamics Model

We use a social network model to represent traders with competing narratives interacting and potentially altering their opinions. In this model, traders can be categorized into two communities: those with positive opinions and those with negative ones. Negative traders will uniformly share one narrative, whereas positive traders will promote a contrasting narrative (Long et al., 2023). Consider a network of N_a trading agents forming opinions $x_1, \ldots, x_{N_a} \in \mathbb{R}$ about the price of a tradable asset. Let x_i be the opinion state of agent i. This real-valued scalar opinion variable indicates that a negative x_i indicates an expected decline in prices, while a positive x_i implies an anticipated increase. The vector $X = (x_1, ..., x_{N_a})$ represents the opinion state of the agent network. Agent i is neutral if $x_i = 0$. The origin $X = 0$ is called the network's neutral state. Agent i is *unopinionated* if its opinion state is small, i.e., $|x_i| \leq \vartheta$ for a fixed threshold $\vartheta \approx 0$. Agent i is *opinionated* if $|x_i| \geq \vartheta$. Agents can agree and disagree. When two agents have the same qualitative opinion state (e.g., they both favor the same option), they agree. When they have qualitatively different opinions, they disagree.

We utilize the BFL opinion dynamics from (Franci et al., 2019), which are simplified to the dynamics of N_c clusters or communities, as described in (Bizyaeva et al., 2020). Each cluster, indexed by $q = 1, ..., N_c$, comprises N_q agents out of a total of N_a, such that $\sum_{q=1}^{N_c} N_q = N_a$. These agents form opinions collectively.

Consider two clusters, p and n, representing communities of positive (bullish) and negative (bearish) traders, respectively. For a given cluster q, let the set of all agent indices in that cluster be denoted by I_q. With $q \in \{p,n\}$, $p \neq n$, then, each agent $i \in I_q$ has an opinion that evolves according to the dynamics presented in (Bizyaeva et al., 2020), which can be summarized by the following differential equation:

$$\dot{x}_i = -d_i x_i + u_i(\hat{S}_1(\alpha\hat{x}_p + \gamma_i\hat{x}_n) - \hat{S}_2(\beta\hat{x}_p + \delta_i\hat{x}_n)) + b_q \tag{1}$$

where $\hat{x}_q$ is the average opinion of cluster q:

$$\hat{x}_q = \frac{1}{N_q}\sum_{i\in I_q} x_i \tag{2}$$

and $\hat{S}_z(x)$, $z \in \{1,2\}$ are saturation functions defined as

$$\hat{S}_z(x) = \frac{1}{2}\left(S_z(x) - S_z(-x)\right), \text{where } S_z \text{ are odd sigmoids.} \tag{3}$$

The model in (1) is suitable for testing the aforementioned market's price drivers by considering the parameters as follows:

- $d > 0$ is a *resistance* parameter that drives the rate of change $\dot{x}_i$ towards the neutral point over time. Intuitively, a larger value of d implies the agent is less inclined to change its opinion. Within the context of social sciences, this parameter can symbolize an individual's level of "stubbornness".
- $u \geq 0$ is an *attention* parameter; it affects how $\dot{x}_i$ changes in response to social interactions. Intuitively, a larger value of u indicates greater attention or sensitivity of the agent to other agents' opinions. Thus, the two parameters d and u weigh the relative influence of the linear damping term and the opinion exchange term, respectively; when the influence of d outweighs that of u, the agent pays minimal attention to others. Conversely, if u dominates d, the agent becomes more attentive to others' opinions. The dynamics governing the evolution of the agent's attention parameter, as detailed in (Bizyaeva et al., 2020), are given by:

$$\tau_u\dot{u}_i = -u_i + S_u(\sum_{l=1}^{N_a}(\bar{a_{il}}x_l)^2) \tag{4}$$

let the feedback weight be $\bar{A}_i = \bar{a_{ij}} \in \{0,1\}$. If $\bar{a_{ij}} = 1$, it indicates that agent i is influenced by the status of agent j. The matrix $\bar{A}$ can either correspond to a predefined social network or be determined independently. The saturation function S_u is then decomposed as:

$$S_u(y) = u_f\left(F\left(g\left(y - y_m\right)\right) - F(-gy_m)\right) \tag{5}$$

S_u is defined with $F(x) = \frac{1}{1+e^{-x}}$.

- $\alpha \geq 0$ is the *self-reinforcement* of the cluster's averaged opinions $\hat{x}_q$. This parameter quantifies the degree of dependency of an agent's temporal evolution in opinion on the average of its encompassing cluster. For an elevated α, there's a pronounced amplification of the intrinsic historical or mean consensus of the cluster, potentially driving the system towards a state of reduced external influence and susceptibility to becoming an echo chamber. Conversely, a diminished α results in a diminished anchoring to past consensus, rendering the system more susceptible to external influences.
- β is the *intra-agent* interaction weight, representing how an agent processes and weights opposing opinions within its own decision-making paradigm. This becomes particularly relevant when an agent is faced with multiple choices, such as when formulating opinions on a variety of tradable assets like different stocks, in this case the state of its opinion would be a vector $\hat{X}_p$. However, given that there's only one object for decision-making in my system, this parameter will take a value opposite to α as the second term of the dynamics in (1) is subtracted from the first, making the opinion more emphasized by $\hat{x_p}$.
- γ and δ are the *inter-agent* interaction weights, which determine whether cluster p and cluster n form a consensus $\gamma - \delta > 0$ or a dissensus $\gamma - \delta < 0$. The state feedback dynamics of these parameters take the form of a leaky nonlinear integrator (Bizyaeva et al., 2020):

$$\tau_\gamma\dot{\gamma}_i = -\gamma_i + \sigma_q S_\gamma(\hat{x}_p\hat{x}_n) \tag{6}$$

$$\tau_\delta\dot{\delta}_i = -\delta_i - \sigma_q S_\delta(\hat{x}_p\hat{x}_n) \tag{7}$$

where $\sigma_q \in \{1,-1\}$ is the *design parameter* and $\tau_\gamma, \tau_\delta > 0$ are time scales, and the saturation function is

$$S_c(y) = c_f \tanh\,(g_c y) \quad c \in \{\gamma, \delta\} \tag{8}$$

where $c_f, g_c > 0$. In any configuration of opinions where the product $\hat{x}_p\hat{x}_n$ is notably non-neutral and significantly large, it prompts γ_i to gravitate towards $\sigma\gamma_f$ and δ_i to move towards $-\sigma\delta_f$ (Bizyaeva et al., 2020).
- b is the *input parameter*, potentially derived from environmental factors, such as market fluctuations, or it could signify inherent biases. For traders with a bearish (or negative) opinion, we designate $b_n \leq 0$ to convey a predominant sentiment predicting a price decrease. Conversely, for those holding bullish (or positive) opinions, we

assign $b_p \geq 0$ to signify an expectant bias towards a price ascent.

For non-negative α and β, the terms $\alpha\hat{x}_p$ and $\beta\hat{x}_p$ in Equation (1) exemplify the self-reinforcement mechanism within cluster p. Similarly, the dynamics of cluster n can be described by interchanging $\hat{x}_p$ with $\hat{x}_n$ in the same equation, indicating analogous self-reinforcement. To understand the dynamics in the context of opposing narratives, consider that $\hat{x}_p > 0$ and $\hat{x}_n < 0$. In this situation, $\alpha\hat{x}_p > 0$ and $\beta\hat{x}_p > 0$ will reinforce cluster p to adopt a more positive view, while $\alpha\hat{x}_n < 0$ and $\beta\hat{x}_n < 0$ will push cluster n towards a more negative direction. Such parameterization effectively captures self-reinforcement's role in shaping opposing perspectives.

Herding behavior is modeled by embedding feedback mechanisms into the social influence parameters γ and δ. The system's tendency—towards consensus or dissensus—is dictated by the sign of the parameter σ in the dynamics of Equations (6) and (7). A reversal in σ's sign triggers a shift between consensus and dissensus states: $\sigma = 1$ aligns both clusters towards consensus, whereas $\sigma = -1$ drives them to dissensus.

In the case where $\alpha = \beta = \gamma = \delta = 0$, the dynamics in (1) are linear. Then, $\dot{x}_i$ responds additively to b_q, where b_q is interpreted as an environmental signal. We can model additive response by setting the value of b_q.

4 ILLUSTRATIVE RESULTS

4.1 Trading Dynamics

In the financial market simulation under discussion, we utilize the open-source BSE platform (Cliff, 2012) to model a single-commodity market with $N = 100$ traders, equally divided into buyers and sellers ($N_B = N_S = 50$ each). These traders are restricted to their initial roles, unable to switch between buyer and seller positions. Their sole decision-making capability is in setting their quoting price, and each adopts the BFL-PRDE strategy with a $NP = 5$ parameter. The simulation features a pricing schedule based on symmetric supply and demand curves, ranging from \$60 to \$250. To emulate continuous time, the BSE platform adopts a discrete time-slicing technique with a temporal step-size of $\Delta t = 1/N$, ensuring at least one market interaction per trader per second. The experiments run for one continuous week of round-the-clock, 24/7 trading, each experiment were repeated 50 times. In these experiments, traders are integrated into a fully connected network, facilitating direct interactions among them all.

In my analysis, I attempted to elucidate the underlying trend of transaction prices over time, which Ire initially represented as an agglomeration of individual data points from fifty trials. To achieve this, I implemented a polynomial regression model, recognizing its capacity to adapt to the non-linear nature of my data. Specifically, I chose at least second-degree polynomial regression, which allows for the curvature. This model fits the aggregated data, encapsulating the collective trends across all trials while balancing the need to reflect general tendencies without overfitting the noise inherent in the dataset. The resulting polynomial trend line, plotted against the transaction prices, provided a visual representation of the average directional movement, revealing a more nuanced trajectory of price changes over time and offering a compelling graphical narrative of the price's dynamics.

4.1.1 Self-Reinforcing Dynamics

In Figures 1, 2, 3 and 4, I use the opinion dynamics outlined in equation (1) to model the three key factors affecting traders' strategies in response to competing narratives and how these influence traders' decisions in the market, potentially leading to market price fluctuations.

Figures 1 and 2 illustrate the temporal evolution of $\hat{x_p}$ and $\hat{x_n}$. This aligns with studying the influence of the first factor, which investigates how a linear change in each group's self-reinforcement level, denoted by α_p for cluster p and α_n for cluster n, impacts the market price dynamic. For the context of this study, the parameters are configured with $\gamma = \delta = 0$, and the biases are set at $b_p = 0.05$ and $b_n = -0.05$.

Figure 1*A* illustrates the scenario where both groups exhibit linearly increasing self-reinforcing with $\alpha_p = \alpha_n$, both groups exhibit equal self-reinforcement. This equality leads to a pronounced polarization in the opinion distributions, causing each group to distance from the other. Consequently, the transaction prices stabilize, reflecting the equal influence from both sides. Figure 1*B* captures the scenario where $\alpha_p > \alpha_n$. In this situation, the positive group assumes a dominant role, causing the negative group to lean towards a weaker negative stance. This dynamic results in an increase in transaction prices during the period where the negative group is gravitated by the positive one; however, as both α_p and α_n are linearly increasing, both groups will eventually move away from each other. In contrast, Figure 1*C* presents the case of $\alpha_p < \alpha_n$. Here, the negative group's influence is more significant, drawing the positive group towards a weaker positive position. This leads to a decrease in transaction prices during the period where

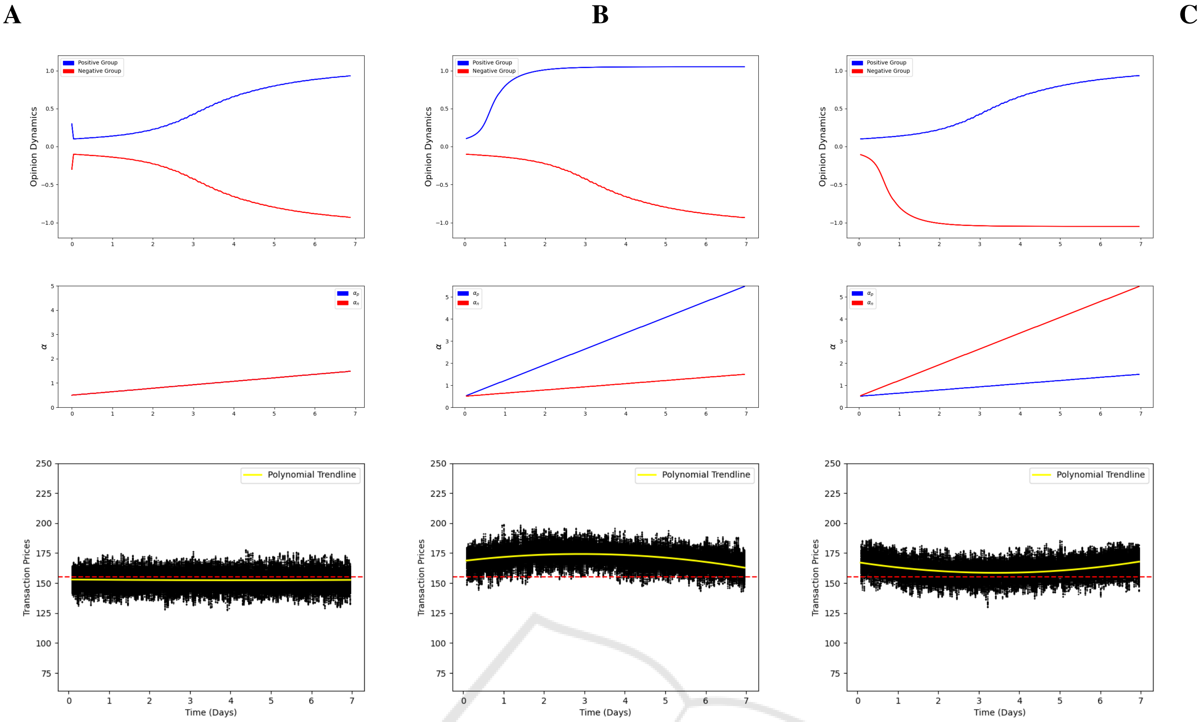

Figure 1: Comparison of system dynamics over three scenarios based on differing self-reinforcement levels (α) in the model (1) of two competing groups. Over a seven-day around-the-clock trading period, the plots display: (A-C, Top) Opinions $\hat{x_p}$ (blue) and $\hat{x_n}$ (red); (A-C, Middle) Self-reinforcement levels α_p (blue) and α_n (red); (A-C, Bottom) Transaction prices from 50 IID experiments (black dots, with the market's theoretical equilibrium price indicated by dashed red line). Initial conditions: $x_p(t0) = 0.3$, $x_n(t0) = -0.3$, $\gamma_p = \gamma_n = \delta_p = \delta_n = 0$, with biases $b_p = 0.05$ and $b_n = -0.05$. (A) Equal self-reinforcement: $\alpha_p(t) = \alpha_n(t) = 0.5 + (1.0)(t - t_0/D)$. This results in symmetrical opinion dynamics, and transaction prices remain close to equilibrium. (B) The positive group is more self-reinforcing with $\alpha_p(t) = 0.5 + (5.0)(t - t_0/D)$ and $\alpha_n(t) = 0.5 + (1.0)(t - t_0/D)$. Here, the negative group's opinions are swayed by the stronger positive group, leading to transaction prices above equilibrium. (C) Greater self-reinforcement in the negative group with $\alpha_n(t) = 0.5 + (5.0)(t - t_0/D)$ and $\alpha_p(t) = 0.5 + (1.0)(t - t_0/D)$ leads to the positive group's opinions being influenced more by the negative group, resulting in transaction prices below equilibrium.

the negative group is attracting the positive one.

Based on this model of self-reinforcement, I can conclude that when one group is at least five times more self-reinforcing than another, it can influence transaction prices. As both groups increasingly reinforce their own opinions, their views tend to polarize over time

Figure 2*A* illustrates the scenario where both groups exhibit linearly decreasing self-reinforcing with $\alpha_p = \alpha_n$, both groups exhibit equally decreasing self-reinforcement. This equality leads to both groups moving toward the neutral point in the opinion distributions, causing each group to get close to the other. Consequently, the transaction prices stabilize, reflecting the equal influence from both sides. Figure 2*B* captures the scenario where the negative group becomes linearly less self-reinforcing $\alpha_p > \alpha_n$. In this situation, the positive group assumes a dominant role, causing the negative group to lean towards a weaker negative stance. This dynamic results in an increase in transaction prices as both α_p and α_n are linearly decreasing, and both groups will eventually move close to each other in the direction of the less decreasing self-reinforcement group. In contrast, Figure 2*C* presents the case of $\alpha_p < \alpha_n$. Here, the negative group's influence is more significant, drawing the positive group towards a weaker positive position. This leads to a decrease in transaction prices as the negative group attracts the positive one.

The model indicates that a group with at least a five times lower self-reinforcement rate exerts a disproportionate influence on transaction prices by being more susceptible to the opposing group's opinion.

4.1.2 Herding Dynamics

Figure 3 demonstrates the impact of the relationship between the opinion network weights γ and δ on market price dynamics, as per the second factor. For this illustration, the setup is divided into two phases. Dur-

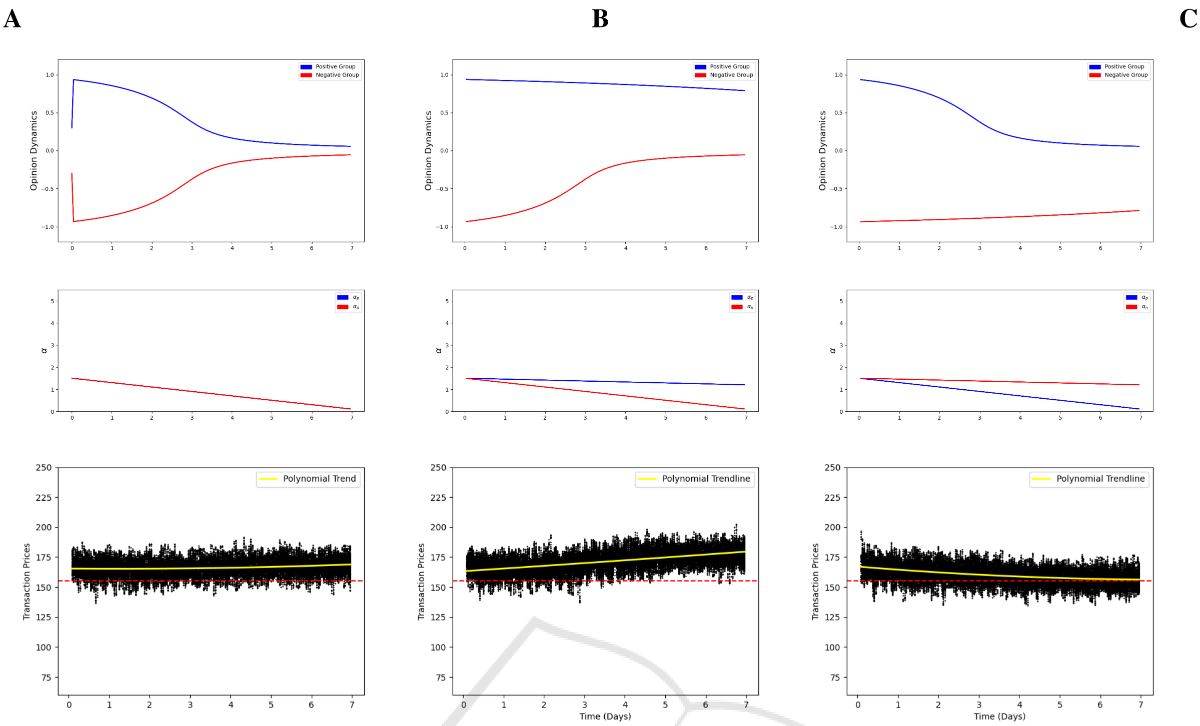

Figure 2: Comparison of system dynamics over three scenarios based on differing self-reinforcement levels (α) in the model (1) of two competing groups. Over a seven-day around-the-clock trading period, graph format and initial conditions are as in figure 1. (A) Equal linearly decreasing self-reinforcement: $\alpha_p(t) = \alpha_n(t) = 1.5 - (1.4)(t - t_0/D)$. This results in symmetrical opinion dynamics, and transaction prices remain close to equilibrium. (B) The positive group is more self-reinforcing with $\alpha_p(t) = 1.5 - (1.4)(t - t_0/D)$ and $\alpha_n(t) = 1.5 + (0.3)(t - t_0/D)$. Here, the negative group's opinions are swayed by the stronger positive group, leading to transaction prices above equilibrium. (C) Greater self-reinforcement in the negative group with $\alpha_n(t) = 1.5 - (1.4)(t - t_0/D)$ and $\alpha_p(t) = 1.5 - (0.3)(t - t_0/D)$ leads to the positive group's opinions being influenced more by the negative group, resulting in transaction prices below equilibrium.

ing the first half of the period, $\gamma - \delta < 0$, leading the group to a state of dissensus (or anti-herding). In the latter half, $\gamma - \delta > 0$ drives the two groups towards consensus (herding). The additional parameters are configured as $\alpha = \beta = 0$ with biases set at $b_p = 0.05$ and $b_n = -0.05$.

From Figure 3*(a)* and *(d)*, it's evident that the opinion distribution either leans positively or negatively, an outcome closely tied to the opinion formation process. Specifically, when herding gravitates towards the positive, transaction prices are above equilibrium. Conversely, when the trend is negative, transaction prices are below the equilibrium. Particularly, in *(b)* and *(c)*, herding towards a positive consensus leads to transaction prices that exceed the equilibrium value. In contrast, *(e)* and *(f)* show that negative herding results in transaction prices falling below equilibrium.

Figures 3*(b)* and *(e)* portray the transaction prices within a market consisting of 100 trades. Despite the existence of price shifts, their magnitude appears minimal. On the other hand, Figures 3*(c)* and *(f)* demonstrate transactions in a smaller market of 10 traders, where the shifts in prices are more pronounced, underscoring the influence of market size on price volatility.

To determine whether there was a statistically significant effect between the herding and anti-herding phases of both experiments displayed in Figure 3, we used the Wilcoxon-Mann-Whitney U test to determine whether the distribution of transaction prices in the first three days of the 7-day experiment had a different central tendency from the distribution of prices over days 5, 6 and 7. As my system is inherently stochastic, we repeated 50 i.i.d. trials at each set of initial conditions. Typically, around 40 of the 50 trials would result in highly significant distributional differences ($p \leq 0.001$), 3 or 4 would be moderately significant ($p \leq 0.01$) and the remainder would show no significant difference. In doing this analysis, we noticed that the p-values decline as the number of traders in the market increases, but even when N= 100 the p-values are in order of 1×10^{-10}. In further work we plan to explore this 'fall-off' in more detail, looking at ever-large populations of traders.

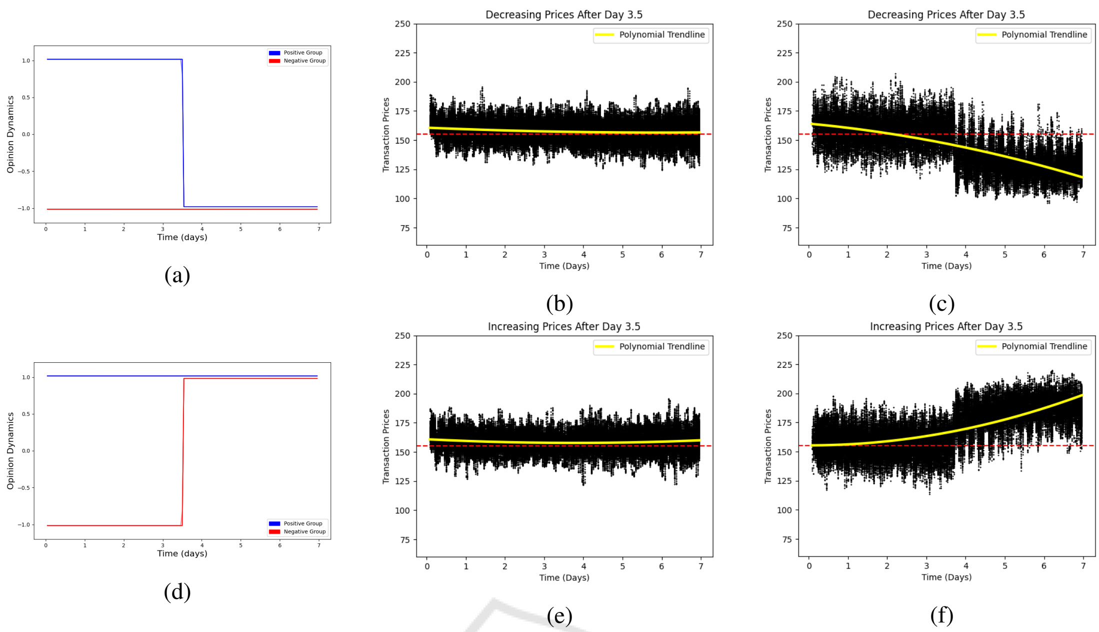

Figure 3: Comparison of system dynamics over two scenarios of herding behavior among the two competing groups is presented. In models 6 and 7, the control parameter σ is set to $\sigma = -1$ for the first half of the period and $\sigma = +1$ for the second half. Over a seven-day around-the-clock trading period, the plots (a) and (d) display herding trends in opinions $\hat{x_p}$ (blue) and $\hat{x_n}$ (red), with negative and positive herding respectively; Plots (b) and (e) show transaction prices —(black dots, with the market's theoretical equilibrium price indicated by dashed red line)— for 50 IID experiments in a market with 100 traders, while (c) and (f) depict transaction prices for 50 IID experiments in a market consisting of 10 traders; all parameters were held constant at $\alpha = \beta = 0$, with biases $b_p = 0.05$ and $b_n = -0.05$ for each respective group.

4.1.3 Additive Response Dynamics

Figure 4 illustrates simulations of the effects of the additive response, denoted by b, where input magnitudes linearly increase over time. Each group's input is directionally represented, with $b_p(t) > 0$ and $b_n(t) < 0$. These simulations were conducted with the self-reinforcement and herding parameters set to zero ($\alpha = \beta = \gamma = \delta = 0$) and initial opinions at equilibrium.

Subfigure *A* of Figure 4 depicts the results when $|b_p(t)| = |b_n(t)|$, meaning both inputs have identical magnitudes but opposite directions, increasing linearly with time. The simulation indicates that when inputs have the same magnitude but are diametrically opposed, they impact the opinion dynamics, resulting in both groups moving away from one another. This creates an equilibrium with equal forces, leading to the stabilization of transaction prices.

Subfigure *B* of Figure 4 demonstrates the dynamics when the positive group is influenced by a stronger additive sentiment, denoted by $|b_p|$, which surpasses that of the negative group, $|b_n|$. The simulation results show a pronounced intensification in the positive group's opinions, paralleled by a noticeable increase in transaction prices. This signifies that dominant positive sentiments bolster the group's collective stance, pushing the market trends upward.

Conversely, Subfigure *C* of Figure 4 illustrates a scenario where the negative group is affected by a prevailing additive negative sentiment. This situation leads to a marked decrease in transaction prices. The graphical results underscore the impact of negative group sentiment in swaying price dynamics, which manifests as a downward pressure on market valuations, reflecting a pessimistic outlook by the group.

5 DISCUSSION AND CONCLUSION

The GameStop episode is an illustrative case of the multifaceted interplay among diverse narratives within financial markets. It clearly demonstrates the potent impact of varied market sentiments, particularly when amplified through social media platforms such as Reddit. This confluence can give rise to abnormal market phenomena, such as short squeezes, underscoring the substantial influence of collective narratives on market functioning. This scenario presents a significant challenge, necessitating the introduction of rational regulatory measures to

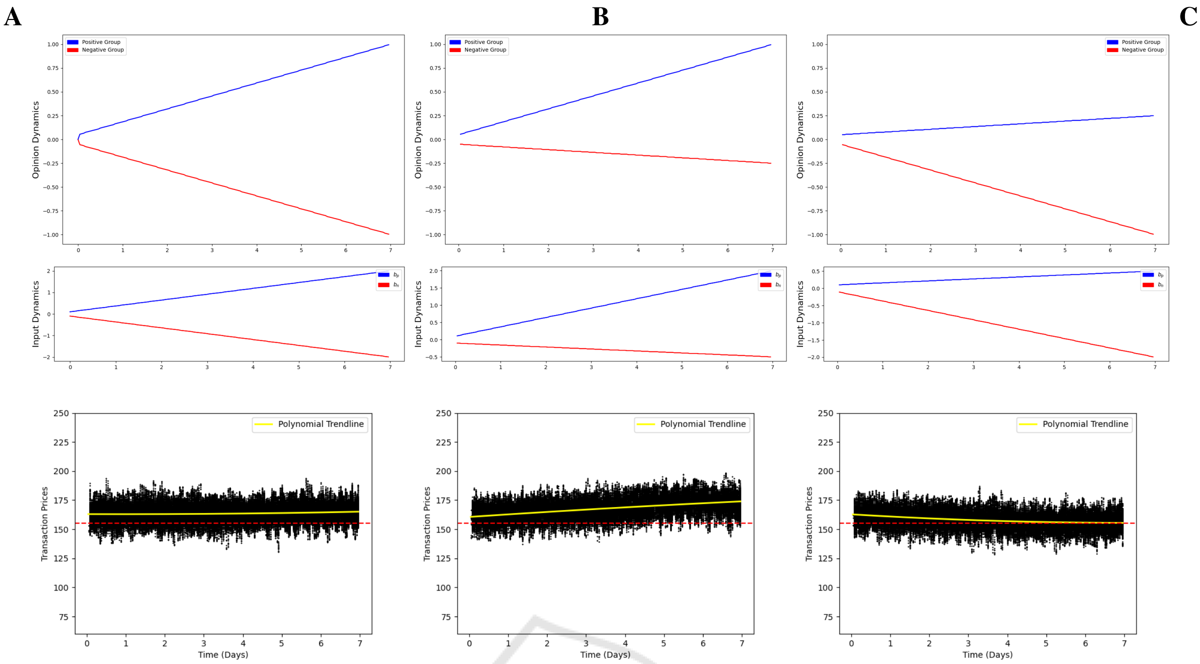

Figure 4: Comparison of system dynamics over three scenarios based on differing the levels of inputs (b) in the model (1) of two competing groups from 50 IID experiments running over a seven-day around-the-clock trading period, the plots display: (A-C, Top) Opinions $\hat{x_p}$ (blue) and $\hat{x_n}$ (red); (A-C, Middle) b values b_p (blue) and b_n (red); (A-C, Bottom) Transaction prices (black dots, with the market's theoretical equilibrium price indicated by dashed red line). Initial conditions: $x_p(t0) = x_n(t0) = 0$, $\alpha_p = \alpha_n = \beta_p = \beta_n = \gamma_p = \gamma_n = \delta_p = \delta_n = 0$. (A) Equal additive inputs: $|b_p(t)| = |b_n(t)| = 0.1 + (1.0)(t - t_0/D)$. This results in symmetrical opinion dynamics, and transaction prices remain close to equilibrium. (B) The positive group is receiving more input with $b_p(t) = 0.1 + (1.0)(t - t_0/D)$ and $b_n(t) = 0.1 + (0.15)(t - t_0/D)$. Here, the positive group is receiving a stronger input, leading to transaction prices above equilibrium. (C) Greater inputs to the negative group with $b_n(t) = 0.1 + (1.0)(t - t_0/D)$ and $b_p(t) = 0.1 + (0.15)(t - t_0/D)$ leads the negative group to have stronger opinions, resulting in transaction prices decreasing.

mitigate market imbalances.

In this paper I have developed a testbed to examine the principal factors influencing group dynamics amidst conflicting narratives in financial markets: collective self-reinforcement; herding behavior; and the assimilation of new information. With its adaptable parameters, this testbed provides a valuable tool for assessing the impact of these factors on market behavior in various scenarios.

The research presented here is significant in that it offers a method to quantify and model the effects of competing narratives on financial markets. I introduce a framework that integrates qualitative narrative dynamics with quantitative decision-making. Our proposed model enhances the financial analyst's toolkit by providing an empirical approach to anticipate the market volatility resulting from narrative shifts. With financial markets becoming ever more sensitive to the rapid flow of information, the ability to understand the role of narratives is crucial for comprehensive market analysis, highlighting the importance and relevance of our research. To facilitate replication and further advancement of this work, I will provide the system's source code as an open-source repository on GitHub[3]. I look forward to the diverse applications and enhancements the research community will derive from this resource.

REFERENCES

Aliber, R. Z., Kindleberger, C. P., and McCauley, R. N. (2015). *Manias, Panics, and Crashes: A History of Financial Crises*. Springer.

Anand, A. and Pathak, J. (2021). WallStreetBets against Wall Street: The Role of Reddit in the Gamestop Short Squeeze. Technical Report IIMB Working Paper-644, Indian Institute of Management Bangalore.

Andreev, B., Sermpinis, G., and Stasinakis, C. (2022). Modelling Financial Markets during Times of Extreme Volatility: Evidence from the GameStop Short Squeeze. *Forecasting*, 4(3):654–673.

Bizyaeva, A., Franci, A., and Leonard, N. E. (2020). A

[3]See my NarrativeEconomics Repository for source code and more details.

general model of opinion dynamics with tunable sensitivity. *arXiv preprint arXiv:2009.04332*.

Bokhari, A. and Cliff, D. (2022). Studying Narrative Economics by Adding Continuous-Time Opinion Dynamics to an Agent-Based Model of Co-Evolutionary Adaptive Financial Markets. In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence*.

Cliff, D. (2012). *Bristol Stock Exchange: open-source financial exchange simulator*. https://github.com/davecliff/BristolStockExchange.

Cliff, D. (2018). BSE: A Minimal Simulation of a Limit-Order-Book Stock Exchange. In Bruzzone, F., editor, *Proc. 30th Euro. Modeling and Simulation Symposium (EMSS2018)*, pages 194–203.

Cliff, D. (2022a). Metapopulation Differential Co-Evolution of Trading Strategies in a Model Financial Market. In *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1600–1609.

Cliff, D. (2022b). Metapopulation Differential Co-Evolution of Trading Strategies in a Model Financial Market. *SSRN: 4153519*.

Cliff, D. (2023). Parameterised Response Zero Intelligence Traders. *Journal of Economic Interaction and Coordination*, pages 1–54.

Franci, A., Golubitsky, M., Bizyaeva, A., and Leonard, N. E. (2019). A Model-Independent Theory of Consensus and Dissensus Decision Making. *arXiv preprint arXiv:1909.05765*.

Gode, D. and Sunder, S. (1993a). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1):119–137.

Gode, D. K. and Sunder, S. (1993b). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1):119–137.

Hirshleifer, D. (2020). Presidential Address: Social Transmission Bias in Economics and Finance. *The Journal of Finance*, 75(4):1779–1831.

Jakab, S. (2022). *The Revolution That Wasn't: GameStop, Reddit, and the Fleecing of Small Investors*. Penguin.

Kameda, T., Inukai, K., Wisdom, T., and Toyokawa, W. (2014). The concept of herd behaviour: its psychological and neural underpinnings. *Sciences*, 420(8).

Kim, K., Lee, S.-Y. T., and Kauffman, R. J. (2023). Social Informedness and Investor Sentiment in the GameStop Short Squeeze. *Electronic Markets*, 33(1):23.

Liu, S. (2021). Competition and Valuation: A Case Study of Tesla Motors. *IOP Conference Series: Earth and Environmental Science*, 692(2):022103.

Lomas, K. and Cliff, D. (2021). Exploring Narrative Economics: An Agent-Based Modeling Platform that Integrates Automated Traders with Opinion Dynamics. In Rocha, A.-P., Steels, L., and van den Herik, J., editors, *Proceedings of the 13th International Conference on Agents and Artificial Intelligence (ICAART2021)*, volume 1, pages 137–148. SciTePress.

Long, S., Lucey, B., Xie, Y., and Yarovaya, L. (2023). "I just like the stock": The role of Reddit sentiment in the GameStop share rally. *Financial Review*, 58(1):19–37.

Shiller, R. (2017). Narrative Economics. *American Economic Review*, 107(4):967–1004.

Shiller, R. (2019). *Narrative Economics: How Stories Go Viral & Drive Major Economic Events*. Princeton.

Smith, V. L. (1962). An Experimental Study of Competitive Market Behavior. *Journal of Political Economy*, 70(2):111–137.

Storn, R. and Price, K. (1997). Differential Evolution–A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11:341–359.

Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. Doubleday, Anchor Press.