Exploring the Impact of Competing Narratives on Financial Markets II: An Opinionated Trader Agent-Based Model with Dynamic Feedback

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: Employing an agent-based trading model integrated with opinion dynamics, we conduct a systematic exploration of the factors potentially contributing to financial market frenzies. Applying our previously established testbed described in detail in a companion paper (part I), we examine the influence of two competing narratives on three hypotheses: self-reinforcement; herding; and an additive response to inputs. Utilizing a real-world dataset, we investigate these dynamics. Our findings reveal that although all three hypotheses affect price movements, herding behavior has the most substantial impact. The source code for these simulations is available on Github, allowing researchers to replicate and extend our work.

1 INTRODUCTION

In financial markets, frenzies are characterized by unusually high trading volumes, significant price volatility, and, often, divergent opinions about asset valuations. With the advent of instantaneous information dissemination, narratives can play an important role in precipitating these frenzies (Shiller, 2019; Hirshleifer, 2020).

The field of narrative economics, which explores the power of prevalent stories to affect economic decisions, has been notably advanced by Robert Shiller (Shiller, 2017; Shiller, 2019). These narratives do more than reflect collective sentiment; they are active agents in the marketplace, entwining truth and fiction in ways that can create ambiguity and divergent interpretations. As such, they play a critical role in investment decisions, acknowledged by experts and institutional investors alike (Kim et al., 2023). The propagation of these narratives can be substantially amplified through social media, which, in conjunction with trading platforms like Robinhood, can contribute to driving stock prices to extremes, as seen in the GameStop price surge in January 2021 (Kim et al., 2023; Jakab, 2022; Aliber et al., 2015).

It is not just the narrative's existence that makes it powerful but also its interaction with other narratives, sometimes opposing ones. Tesla's valuation surge in 2013, underscored by the promise of a sustainable electric future, faced skepticism through a counterpoint questioning its financial valuation (Liu, 2021). Similarly, Bitcoin's dramatic rise in 2017, supported by the decentralized currency narrative, encountered opposition, highlighting its potential for misuse and inherent volatility. Instances like the 2021 GameStop frenzy and the FTX collapse in November 2022 further illustrate the intricate interplay between competing narratives and their tangible effects on market movements.

In this study, we investigate the dynamics of conflicting narratives and assess how each of the three factors—self-reinforcement, herding, and additive response to inputs, which will be discussed in the following section—impacts their capacity to shape the intensity and orientation of market frenzies. Insight into these dynamics is pivotal for devising informed regulatory interventions aimed at controlling speculative activities and maintaining market stability.

1.1 Self-Reinforcement, Herding, and Additive Response Among Traders

In this paper, we extend my previous work (A, 2024) by incorporating market mood as a determinant of market dynamics in the presence of conflicting narratives. Using my agent-based model, we analyze the role of positive feedback loops in narrative reinforcement, herding behavior's effect on collective

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

Bokhari, A.
Exploring the Impact of Competing Narratives on Financial Markets II: An Opinionated Trader Agent-Based Model with Dynamic Feedback.
DOI: 10.5220/0012456500003636

In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART 2024) - Volume 1*, pages 138-148
ISBN: 978-989-758-680-4; ISSN: 2184-433X

decision-making, and the influence of external inputs on group behavior. By systematically comparing these dynamics—positive feedback (Hypothesis A), herding (Hypothesis B), and external influence (Hypothesis C)—my model explains their distinct and combined effects on market behavior under the context of bullish and bearish trader groups.

- **Self-Reinforcement:** Narratives within closed groups lead to stronger, more entrenched opinions over time. This is exemplified by digital "echo chambers," where the absence of opposing views strengthens beliefs through a feedback loop. Similarly, in financial markets, self-reinforcement is seen as trends gain momentum and influence investor behavior, thus reinforcing the prevailing market direction. This phenomenon demonstrates its impact in both digital and economic spheres.
- **Herding Behavior:** Herd behavior in financial markets is the propensity of individuals to mimic the actions or beliefs of their peers, influenced more by collective dynamics than individual decision-making (Kameda et al., 2014). This phenomenon is exemplified in the GameStop short squeeze event. Key figures such as Keith Gill, a financial advisor, played significant roles (Anand and Pathak, 2021). Gill's bullish view on GameStop, recognizing its high short interest, led many to follow his investment strategy, resulting in a feedback loop that significantly inflated the stock's price. This behavior, driven by a fear of missing out rather than a deep understanding of market fundamentals, led to a substantial increase in the stock price, especially as institutions that had shorted the stock were compelled to buy back at higher prices. This case underscores how herd behavior can lead to rational bubbles in the market, diverging from the "wisdom of crowds" principle, which relies on diverse, independent thinking (Surowiecki, 2004; Kim et al., 2023; Andreev et al., 2022).
- **Additive Response:** In financial contexts, this refers to investors' reactions based solely on external stimuli, independent of market data or collective sentiment. The GameStop frenzy provides a clear example of this. Influencer Keith Gill's decision to hold his stocks, despite significant unrealized profits, served as an external stimulus for many investors, who then mirrored his stance. This reaction was not based on market fundamentals but rather on additive stimuli like rallying phrases such as "diamond hands" and "YOLO", demonstrating the impact of such external signals in driving investor behavior contrary to standard market practices.

1.2 Structure of the Paper

This paper, the second in a two-part series, builds on the agent-based model (ABM) introduced in its predecessor (A, 2024). While the first paper provides a comprehensive detail of the ABM, this one extends the model and applies it to real-world data. For coherence and completeness, aspects of model design are reiterated here, mirroring the inclusion of illustrative results in (A, 2024). The structure of this paper is as follows: Section 2 reviews the related background; Section 3 describes the ABM's design and operation; Section 4 displays the results; and finally, Section 5 concludes the paper.

2 BACKGROUND

In the dynamic field of financial economics, a profound paradigm shift is unfolding, profoundly altering our comprehension of market mechanics. The Efficient Market Hypothesis (EMH), long revered as the foundational pillar in this domain, asserts that market prices are comprehensive reflections of all relevant information about an asset's intrinsic value. However, this hypothesis encounters substantial difficulties in accounting for certain anomalies within financial markets, notably the unpredictable behaviors observed in cryptocurrency markets, a challenge highlighted in Shiller's 2017 analysis of Narrative Economics.

At the forefront of this intellectual evolution is the emergence of narrative economics, a theory that posits a paradigmatic shift from the conventional reliance on empirical, quantifiable data, proposing instead that the narratives and stories pervading amongst market participants wield a formidable influence on economic trajectories (Shiller, 2017). In this context, narratives are not mere anecdotes but potent, contagious entities that disseminate through the intricacies of social networks, molding public sentiment in a manner akin to biological epidemics (Shiller, 2019). Grasping the essence and flow of these narratives is crucial for decoding the underlying currents driving market movements, particularly in instances where traditional economic theories offer inadequate explanations.

Parallel to this narrative-centric approach is the rapidly developing field of Opinion Dynamics (OD), a discipline dedicated to unraveling the formation and propagation of opinions within societal constructs. The first application to the domain of financial markets (Lomas and Cliff, 2021), OD elucidates the intricate nexus between socio-behavioral dynamics and economic phenomena, offering a nuanced lens through which market behaviors can be interpreted. The integration of OD principles with agent-based modeling in financial markets heralds a new era in economic analysis, elucidating how shifts in the perceptions and stances of trader-agents can manifest in tangible market price movements (Lomas and Cliff, 2021).

3 MODEL

3.1 BFL-PRDE Trader Model

In the evolution of trading strategies, PRZI (**P**arameterised-**R**esponse **Z**ero **I**ntelligence), as introduced in (Cliff, 2023), laid the groundwork for adaptive zero-intellegence (ZI) traders, paving the way for its successor, PRDE (**P**arameterized-**R**esponse **D**ifferential **E**volution) (Cliff, 2022). ZI traders have a long tradition of productive use as minimal models of human traders: see for example (Farmer et al., 2005) (Ladley, 2012), and (Axtell and Farmer, 2018).

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 (Bokhari and Cliff, 2022) extend the PRDE framework by incorporating a real-valued opinion variable, utilizing the opinion dynamics model proposed by **B**izyaeva, **F**ranci, and **L**eonard (BFL). 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 the ZI-trader strategies GVWY (Cliff, 2012; Cliff, 2018) and ZIC(Gode and Sunder, 1993), manifest heightened urgency, influencing their quote prices. Conversely, sellers lean towards a more relaxed position in the form of a hybrid between the ZI trader strategies 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, 2022), 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 is represented by 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 uses the standard "vanilla" DE/rand/1. 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 competing narratives. In this model, traders are 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 mathematically captured 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 hypotheses 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:

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

  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

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

  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.

3.3 Market Mood Input

The influence of aggregate market mood MM, as expressed through social media posts, on the market frenzy is profound. This assertion requires an understanding of the dynamics between social media sentiments and market behavior. With billions of users, social media platforms have become a crucial source of real-time collective sentiment on various subjects, including financial markets. When a substantial number of users express either positive or negative sentiments about a particular stock or the market as a whole, it creates an overarching market mood. This mood can either result from genuine financial news or a byproduct of widespread speculative opinions. Moreover, traders nowadays often turn to social media as a quick pulse-check on prevailing market sentiments before making decisions. If the aggregate sentiment is overwhelmingly positive, it can lead to heightened buying activity, potentially causing asset prices to surge. Conversely, a negative mood can lead to mass selling, driving prices down. This cascading effect, where social media sentiments bolster or dampen market activity, can trigger market frenzies.

The model of opinion dynamics is seeded with a market mood indicator that in turn can influence the three potential drivers of market frenzy. This is achieved by converting the $MM(t)$ into market mood inputs specific for two groups, $I_p(t)$ and $I_n(t)$, through the following generic nonlinearity (Leonard et al., 2021).

$$I_p(t) = f(MM(t) + I_0), \tag{9}$$

$$I_n(t) = f(-MM(t) - I_0), \tag{10}$$

where f is a function such that $f(0) = 0$ and $I_0 > 0$ is the basal opinion drive.

In this model adopted from (Leonard et al., 2021), the sentiment or opinion influence of the two distinct trader types at any given time t is derived from the market mood $MM(t)$. The model in (9) represents the bullish or positive trader's sentiment influence, $I_p(t)$, which is a function, f, of the market mood plus a constant basal opinion drive, I_0. The equation in (10) captures the sentiment influence of the bearish or negative trader, $I_n(t)$. This is described by the function, f, of the negative of the market mood subtracted from the negative of the basal opinion drive. The function f is such that its output for an input of zero is zero, ensuring a mean-neutral interpretation. The term I_0 represents a fundamental, inherent sentiment bias that is greater than zero, highlighting an intrinsic opinion drive irrespective of prevailing market conditions. Together, these equations offer a dynamic representation of how different trader groups might respond to fluctuations in market sentiment, while also accounting for a foundational bias in their opinion formation.

Traders tend to be inattentive to minor fluctuations in market mood, preferring more pronounced alterations. To capture this nuanced behavior, the model incorporates a "dead zone". This can be conceptualized as a threshold region where trivial market mood fluctuations remain largely ineffective at altering trader sentiments. The primary objective is to identify the precise boundaries where trader responsiveness becomes paramount. By introducing the conceptual "dead zone". The function f is defined as a nonlinear function (Leonard et al., 2021) given by:

$$f(x;U,L) = \begin{cases} x - U, & \text{if } x \geq U \\ 0, & \text{if } -L < x < U \\ x + L, & \text{if } x \leq -L \end{cases} \tag{11}$$

Where U and L are the upper and lower sensitivity thresholds, both of which are non-negative. To account for different types of traders, we denote U_p and L_p as sensitivity thresholds for bullish traders, and U_n and L_n for bearish traders. Hence, the market sentiment input for bullish traders, $I_n(t)$, incorporates $f(x;U_p,L_p)$ and for bearish traders, $I_p(t)$ is defined using $f(x;U_n,L_n)$.

3.3.1 Market Mood Drives Trader Self-Reinforcement Dynamics

We first explore the first hypothesis. To incorporate the adaptive behavior of traders to the prevailing market mood, we let the market mood inputs $I_n(t)$ and $I_p(t)$ influence trader dynamics through the self-reinforcing levels α_n and α_p. traders do not respond instantaneously to shifts in the market mood. Their reactions either involve adapting their strategies to resonate with the current mood or re-evaluating their investment portfolio. Both of these approaches necessitate time, making their responsiveness typically gradual.

To encapsulate the overall sensitivity of trader dynamics to the market mood, the rates of change of α_n and α_p are proportional to the market mood inputs with a common proportionality constant k_α (Leonard et al., 2021):

$$\frac{d\alpha_p}{dt} = k_\alpha I_p(t) \tag{12}$$

$$\frac{d\alpha_n}{dt} = k_\alpha I_n(t) \tag{13}$$

Building upon the foundational arguments presented in (Kim et al., 2023; Hirshleifer, 2020), it is evident that market narratives often commence as weak and heavily skewed to individual biases. Yet, through iterative community exchanges and continuous dissemination, these narratives intensify, evolving into increasing, self-reinforcing mechanisms that profoundly influence trader opinions.

The core proposition of this modeling framework is that a strong bullish sentiment in the market [represented by $MM(t) > 0$] leads to an intensification of self-reinforcing strategies among bullish traders when $MM(t) + I_0 \geq U_p$, and a diminution of the same among bearish traders when $MM(t) - I_0 \leq L_n$. Conversely, a pronounced bearish sentiment [indicated by $MM(t) < 0$] spurs an increase in the bullish traders' self-reinforcement when $-(MM(t) - I_0) \geq U_n$ and a decrease in the bearish traders' strategies when $-(MM(t) + I_0) > L_p$. These dynamics can culminate in an amplification of the prevailing market sentiment.

3.3.2 Market Mood Drives: Trader Herding Dynamics

To evaluate the second hypothesis, my attention is redirected to herding behavior, governed by the parameters γ and δ. In this context, we employ $I_p(t)$ and $I_n(t)$ as the driving forces behind dynamic alterations in σ, operating under the constraints $\alpha_p = \alpha_n = 0$ and the magnitudes $|b_p|$ and $|b_n|$ being minimal.

The tendency to herd, characterized by a parameter σ, is modulated based on the prevailing market mood. The parameter σ is introduced to influence and control herding behavior amongst the agents in the simulation. Given the market mood $MM(t)$, which ranges from -1 (representing extreme fear) to 1 (representing extreme greed), To ensure that σ approaches the value of $+1$ during extreme market moods (either positive or negative), we can model its rate of change as

Given:

$$f(x;\mathrm{U},\mathrm{L}) = \begin{cases} \frac{1}{1+e^{-x}} & x \in [-\mathrm{U},-\mathrm{L}] \cup [\mathrm{L},\mathrm{U}] \\ -\frac{1}{1+e^{-x}} & x \in (-\mathrm{L},\mathrm{L}) \\ 0 & \text{otherwise} \end{cases}$$

Then evolution of σ_p and σ_n are given by:

$$\frac{d\sigma_p}{dt} = K_\sigma \cdot f(I_p)$$

$$\frac{d\sigma_n}{dt} = K_\sigma \cdot f(I_n)$$

Subject to:

$$-1 \leq \sigma_p, \sigma_n \leq 1$$

This definition ensures that herding occurs at extremes of market sentiment, capturing periods when the market is either extremely fearful or extremely greedy.

3.3.3 Market Mood Drives Additive Response Dynamics

For the third hypothesis, we explore the application of the additive responses to the signals b_p and b_n. We utilize $I_p(t)$ and $I_n(t)$ to influence the dynamics of $|b_p|$ and $|b_n|$, in a manner similar to (12 and 13), with the constraints $\alpha_p = \alpha_n = \gamma_p = \gamma_n = 0$.

$$\frac{db_p}{dt} = K_b I_p(t) \tag{14}$$

$$\frac{db_n}{dt} = K_b I_n(t) \tag{15}$$

4 RESULTS

In this study, we employ the open-source BSE platform (Cliff, 2012) for simulating a financial market, specifically focusing on a market with a single commodity and 50 participants. These participants are split equally into 25 buyers and 25 sellers, with their roles being fixed throughout the experiment. Their trading strategy involves the BFL-PRDE method with a parameter setting of $NP = 5$. Participants' decisions revolve around setting their offer prices within a price range of \$60 to \$250, influenced by symmetrically shaped supply and demand curves. The BSE platform utilizes a discrete approach to simulate continuous trading, with a time-step of $\Delta t = \frac{1}{N}$, guaranteeing at least one transaction per trader per second.

In Figure 1, the market mood's temporal variation is depicted as two complete sine wave cycles over a seven-day period, with hourly intervals. Mood values are quantified on a spectrum from -1 ("Extreme Fear") to $+1$ ("Extreme Greed"), with intermediate values classified into "Fear" (-0.5 to 0), "Neutral" (0), and "Greed" (0 to 0.5). This quantification converts market sentiment into a measurable and structured format for analysis. The sine wave was chosen for its natural symmetry, ensuring an unbiased average mood value of zero. It portrays the mood's oscillation between extreme points and neutrality, as expressed by the function $MM(t)$ within the range of 0 to 4π, thereby encapsulating the complete mood dynamics within the depicted time frame. Please note that there is no suggestion here that mood in real markets follows nice, clean sinusoidal curves; rather, we are using these sine-wave mood functions to give maximum clarity in explaining /exploring the behaviour of my model.

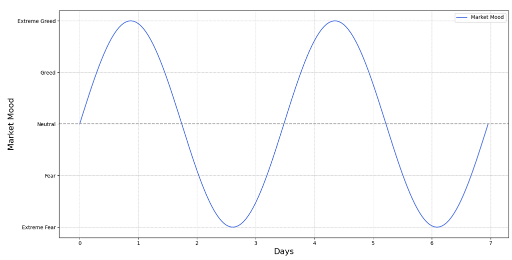

Figure 1: Market Mood. This time series forms two perfect cycles, transitioning through neutral, greed, extreme greed, back to neutral, followed by fear, and then extreme fear.

Figure 2 shows the results of the three hypotheses of driving market frenzy, using the market mood as the influence for each controlling parameter. The results illustrate how market mood dynamics influence the market dynamics.

In Figure 2 *A*, to produce α values using Equation (13) and (12), we set $\alpha(t_0) = 0.9$ and $k_{\alpha_p} = k_{\alpha_n} = 0.9$. When the self-reinforcement parameter is applied, it is evident that when $\alpha_p > \alpha_n$, the positive group's opinions are stronger, leading to higher prices. Conversely, when $\alpha_p < \alpha_n$, the negative group is more self-reinforcing, resulting in decreasing prices. This is highly correlated with the underlying market mood.

In Figure 2 *B*, when the market mood indicates extreme greed ($\sigma = 1$), both groups exhibit herding behavior. This is influenced by the opinion distribution and updates, as the γ and δ dynamics are highly dependent on the product $\hat{x}_p \hat{x}_n$. Transaction prices are highest when both groups are herding between days two and four and between days five and seven, which is indeed highly correlated with the market mood.

In Figure 2 *C*, we observe the opinion distribution when both groups receive additive inputs from an external source. Using Equation (14) and (13), we set $b(t_0) = 1.0$ and $k_p = k_n = 1.0$. The negative group receives its input as a negative value. When $|b_p| > |b_n|$, the positive group receives greater input, resulting in stronger opinions and increasing transaction prices. On the other hand, when $|b_p| < |b_n|$, the negative group receives greater input, leading to a decrease in transaction prices.

4.1 Null Hypothesis Testing

We consider and reject the null hypothesis as all three hypotheses do cause price fluctuations. In my analysis of transaction prices under Hypotheses A, B, and C, we observe distinct patterns essential in evaluating these prices' volatility. Both the mean and median values across the hypotheses exhibit a degree of similarity, suggesting comparable central tendencies. However, notable differences emerge in measures of dispersion. Hypothesis B is particularly prominent, displaying the highest variance (227.33) and standard deviation (15.04), indicative of more significant price fluctuations compared to Hypotheses A and C. This finding is further reinforced by the largest range (81.03) and the greatest number of outliers (12.02) in Hypothesis B, suggesting more extreme variations in price changes. Conversely, Hypotheses A and C present lower levels of variance (110.76 for A) and standard deviation (10.50 for A), with Hypothesis A being relatively less volatile.

To statistically confirm these differences in price fluctuations, we conducted the Kruskal-Wallis H test. The results of this test showed a statistic of 5155.54 and a p-value of 0.0. The exceedingly low p-value, effectively below the computational precision threshold, provides robust evidence against the null hypothesis, suggesting significant differences in the distribu-

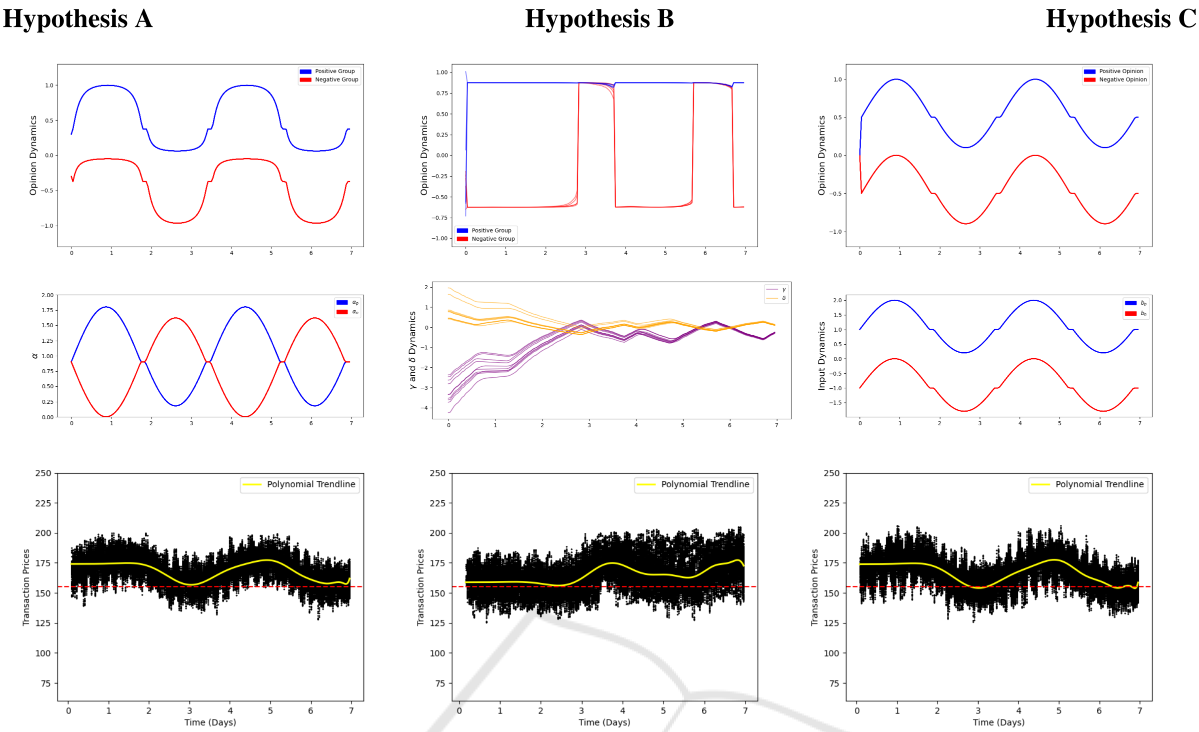

Figure 2: This figure demonstrates the comparative model dynamics as defined in equation (1) for three distinct hypotheses regarding competing narratives, in conjunction with the syntactic market mood in Figure 1. The data is depicted as a temporal function spanning a full week. Each column, labelled A, B, and C, showcases the results derived from implementing Hypotheses A, B, and C, in that order. The top row displays the opinion dynamics distributions, with $\hat{x_p}$ (blue) and $\hat{x_n}$ (red) representing the contrasting market sentiments. The second row captures the unique dynamics inherent to each hypothesis. The last row visualizes the transaction prices—(black dots, with the market's theoretical equilibrium price indicated by a dashed red line and the polynomial trendline is plotted as a yellow line)— originating from 50 IID experiments, which are overlaid with a 24-hour moving average and span an uninterrupted seven-day trading interval.

tion of transaction prices among the three hypotheses. These results corroborate my initial observations derived from descriptive statistics, especially highlighting the distinct behavior of Hypothesis B in terms of price fluctuations as compared to Hypotheses A and C. This supports the notion that herding behavior, as represented by Hypothesis B, is likely to cause the most pronounced price fluctuations.

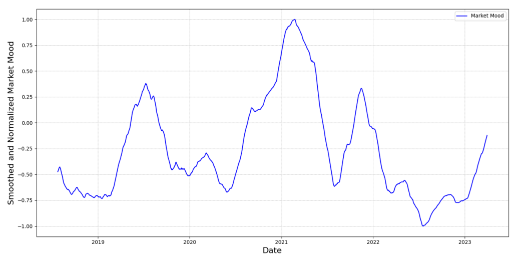

Figure 3: Real-world Market Mood. This time series is smoothed and normalized in $[-1, +1]$.

The dataset, sourced from [2], depicts market mood over time as shown in Figure 3, presenting market mood fluctuations. Each record includes a date and a corresponding value indicating market sentiment. For preprocessing, a simple moving average with a window size of 100 was applied to the data, effectively smoothing out short-term fluctuations and highlighting longer-term trends. Further normalization adjusted the smoothed values to fall within a range of -1

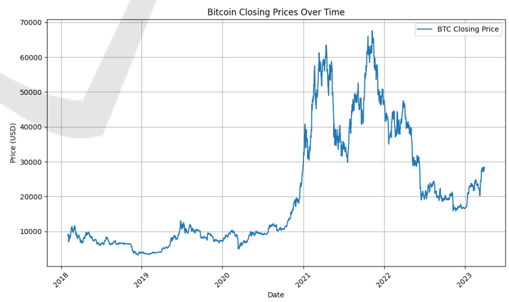

Figure 4: Real-world Transaction prices.

[2]https://www.kaggle.com/datasets/adilbhatti/bitcoin-and-fear-and-greed

to 1, rendering the dataset apt for comparative analysis and interpretation of market mood dynamics. Figure 4 displays the corresponding closing prices, highlighting a significant spike between the years 2021 and 2022

Figure 5 displays the outcomes for three hypotheses underlying market frenzy, with market mood acting as the influencing factor for the controlling parameters. These outcomes depict the impact of market mood on market dynamics.

In subfigure *A* of Figure 5, we compute α values using Equations (13) and (12) with initial conditions $\alpha(t_0) = 0.9$ and $k_{\alpha_p} = k_{\alpha_n} = 0.9$. The application of the self-reinforcement parameter reveals that a dominant α_p over α_n strengthens the positive group's opinions, resulting in rising prices during these periods, while a dominant α_n enhances the negative group's influence, leading to decreasing prices. As shown in α dynamics, the largest difference between α_p and α_n is during the period from year 2021 to 2022. The transaction prices are showing the maximum increase during the same period.

Subfigure *B* in Figure 5 illustrates herding behavior in both groups under extreme greed conditions ($\sigma = 1$), shaped by the opinion distribution and updates, with γ and δ dynamics being highly responsive to the product $\hat{x}_p\hat{x}_n$. It can be seen that when both groups are herding toward the positive because we assume a basis toward the positive when the market mood reports an increased positive sentiment $b_p = b_n = 0.5$, a significant spike in transaction prices accrues during the period from year 2021 to 2022.

In subfigure *C* of Figure 5, the opinion distribution is shown under the influence of external additive inputs to both groups, as described by Equations (14 and 15), setting $b(t_0) = 1.0$ and $k_p = k_n = 1.0$, with the negative group receiving inputs negatively. When $|b_p| > |b_n|$, the positive group's stronger inputs lead to more robust opinions and increasing transaction prices, while $|b_p| < |b_n|$ implies the negative group's greater input, resulting in more influential negative opinions and declining prices. It can be seen in input dynamics that during the period from year 2021 to 2022 the positive group receives a stronger positive input, and the negative group receives a weaker negative input.

4.2 Comparative Analysis of Model Predictions and Bitcoin Closing Prices

The predictive performance of three hypotheses was quantitatively assessed against actual Bitcoin closing prices. Data normalization was conducted using the MinMaxScaler, aligning the scales of the model predictions with the Bitcoin price data for a direct comparison. The assessment utilized three error metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and the Pearson correlation coefficient. The results for each hypothesis are presented as follows:

- **Hypothesis A:** Exhibited an inverse relationship with the actual price trends, evidenced by a MAE of 0.3432, RMSE of 0.3865, and a negative correlation coefficient of -0.1711.
- **Hypothesis B:** Demonstrated a marginal positive linear relationship, with a MAE of 0.4366, RMSE of 0.4717, and a correlation coefficient of 0.0564.
- **Hypothesis C:** Showed a slightly stronger inverse relationship than H1, with a MAE of 0.3362, RMSE of 0.3808, and a correlation coefficient of -0.1863.

The negative correlation coefficients for Hypothesis A and Hypothesis C suggest a counterintuitive relationship, where an increase in the actual Bitcoin prices is associated with a decrease in the model's predicted transactions, and vice versa. The negligible correlation in Hypothesis B indicates an absence of a significant linear relationship. These outcomes underscore a potential misalignment between the predictive models and the actual price behavior, suggesting that the models may not adequately capture the influential market dynamics. Refinements in the models' assumptions, parameters, and the inclusion of more descriptive features could potentially enhance predictive accuracy.

5 CONCLUSION

In conclusion, our investigation sheds light on the complexities of financial market frenzies, particularly within the context of competing narratives. Our approach, which integrates an agent-based trading model augmented by opinion dynamics, allows for a meticulous evaluation of three main hypotheses: self-reinforcement, herding behavior, and responsiveness to additive external inputs. Empirical analysis, anchored in actual market data, suggests that while each hypothesis has its role in influencing price movements, it is the herding behavior that predominantly dictates market behavior. This revelation highlights the complex fusion of individual decision-making and collective market sentiment in the financial domain.

Despite this, there are noticeable divergences between our model's forecasts and the empirical Bitcoin price data, emphasizing the challenges inherent

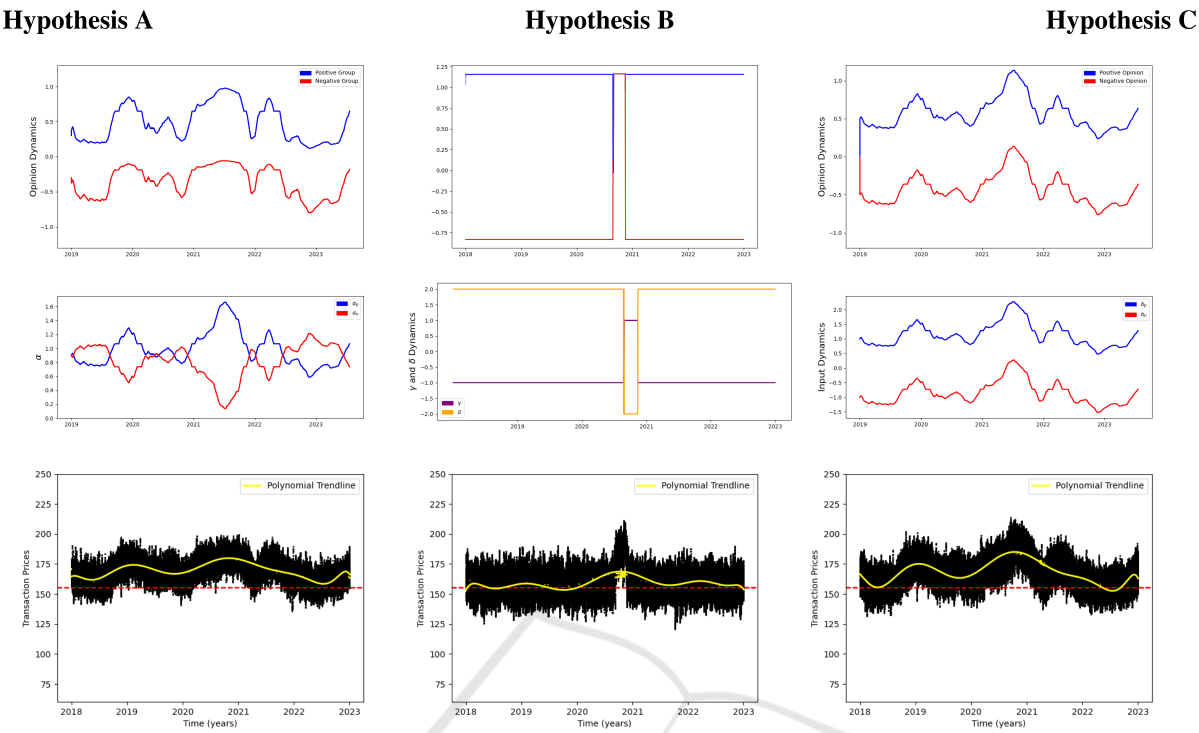

Figure 5: Comparison of the model dynamics (1) for the three hypotheses on the transaction prices using the real-world market mood in Figure 3. Plotted as a function of time during the time from 2018 to 2023.

in precisely simulating financial markets. Recognizing the multifaceted nature of such markets, future endeavors will be directed towards refining the modeling framework. The aim is to iteratively enhance its sophistication, thereby improving its capacity to encapsulate the subtleties and volatility of market price trajectories. This refinement process is expected to advance our understanding of the psychological and sociological factors that drive market behaviors, potentially leading to more robust predictive models. 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

A, A. (2024). Exploring the Impact of Competing Narratives on Financial Markets I: An Opinionated Trader Agent-Based Model as a Practical Testbed. In *Submitted to the 24th International Conference on Agents and Artificial Intelligence (ICAART)*, volume 2, pages First Page–Last Page. INSTICC. Available at Linktothepaperifavailable.

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

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.

Axtell, R. and Farmer, J. D. (2018). Agent-based modeling in economics and finance: Past, present, and future. Technical report.

Bizyaeva, A., Franci, A., and Leonard, N. E. (2020). A 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. (2022). 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. (2023). Parameterised Response Zero Intelligence Traders. *Journal of Economic Interaction and Coordination*, pages 1–54.

Farmer, J. D., Patelli, P., and Zovko, I. (2005). The Predictive Power of Zero Intelligence in Financial Markets. *Proceedings of the National Academy of Sciences*, 102(6):2254–2259.

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. (1993). 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.

Ladley, D. (2012). Zero Intelligence in Economics and Finance. *The Knowledge Engineering Review*, 27(2):273–286.

Leonard, N. E., Keena Lipsitz, Anastasia Bizyaeva, Alessio Franci, and Yphtach Lelkes (2021). The nonlinear feedback dynamics of asymmetric political polarization. *Proceedings of the National Academy of Sciences*, 118(50):e2102149118.

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.

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.