Journal of Mathematics and Modeling in Finance (JMMF)
Vol. 6, No. 2, Summer & Fall 2026
Research paper

Modeling Financial Sentiment with a Three-State Lattice Gas: From Agent Interaction to Market-Level Shock Responses

Halim Zeghdoudi[1], **Fatih Tank**[2]

[1] LaPS Laboratory, Badji Mokhtar-Annaba University, Annaba, Algeria
halim.zeghdoudi@univ-annaba.dz

[2] Department of Economics, Atilim University, Ankara, Turkey
fatih.tank@atilim.edu.tr

Abstract:
This paper develops an agent-based framework for studying financial sentiment dynamics through a three-state lattice gas model in which agents hold bullish, neutral, or bearish positions. Extending classical binary models such as the Ising and voter models, the framework introduces a neutral state to capture hesitation, indecision, and temporary withdrawal from directional beliefs, especially during uncertain market conditions. Sentiment changes are driven by local interactions among neighboring agents and external information shocks, with reversible stochastic dynamics governing the updating process. Analytical results and numerical simulations show the emergence of persistent opinion clusters, entropy shifts under information stress, and gradual recovery toward balance after strong shocks. Empirical illustrations using sentiment extracted from financial social-media data suggest that the model can reproduce key patterns in market reactions. Compared with standard binary approaches, the three-state formulation better captures uncertainty, the buffering role of neutral agents, asymmetric shock responses, and recovery dynamics, making it a useful framework for linking individual sentiment updates to collective market behavior.

Keywords: Agent-Based Modeling, Financial Markets, Statistical Mechanics, Lattice Gas, Sentiment Dynamics, Ising Model, Voter model, Polarization, Entropy, Social Media Analysis.
Classification: 91B84; 82C22; 91G80

1 Introduction

Financial markets are complex systems in which many participants react not only to public information, but also to one another. Prices and market trends therefore reflect more than isolated decisions. They are also shaped by the transmission of beliefs, expectations, and reactions across groups of investors. For this reason, agent-based models have become useful tools for studying collective behavior, belief formation, herding, and market anomalies.

Many models in this literature rely on binary decision rules. Classical examples

[1] Corresponding author
Received: 28/04/2026 Accepted: 13/06/2026
https://doi.org/10.22054/jmmf.2026.92434.1282

include the Ising model [1, 4] and the voter model [2, 5]. In these models, agents usually choose between two possible states and update their positions by aligning with neighbors or imitating peers. Such models can reproduce important features of collective market behavior, including clustering, herding, and consensus formation. However, a binary structure is restrictive when investors are uncertain. Around monetary policy announcements, geopolitical events, earnings releases, or macroeconomic news, many market participants do not immediately take a bullish or bearish position. They may wait for further information, reduce their exposure, or remain temporarily neutral.

This observation motivates the use of a three-state formulation. Building on the analysis of canonical measures and spectral gaps in disordered lattice gas models developed by Zeghdoudi and Boutabia [3], this paper develops a lattice gas model for financial sentiment dynamics. The model allows each agent to take one of three states: bullish $(+1)$, neutral (0), or bearish (-1). The neutral state is not treated as a missing opinion. It represents hesitation, temporary withdrawal from directional trading, or a delayed response to new information. In this sense, neutrality plays an active role in the transmission of market sentiment. It may slow down the spread of bullish or bearish views and affect the recovery of sentiment after a shock [7].

Agents are placed on a one-dimensional periodic lattice and interact with their nearest neighbors. The lattice is not intended to represent physical location. Rather, it is a simple representation of informational proximity: neighboring agents are interpreted as investors who influence one another through common information channels, social interaction, similar trading environments, or shared market narratives. The system may also be affected by external information shocks, such as policy announcements or unexpected news. Sentiment evolves through a Glauber-type stochastic updating rule derived from an energy function. This structure makes it possible to study how local interaction and external information jointly shape aggregate sentiment.

The model also provides several quantities with direct financial interpretation. Magnetization measures the average direction of sentiment. Polarization measures the share of agents holding directional views. Entropy measures the dispersion of sentiment across bullish, neutral, and bearish states. Metastability describes the persistence of local sentiment clusters, while the convergence rate gives information about the speed of post-shock adjustment. Together, these quantities connect individual belief revision with market-level outcomes such as uncertainty, polarization, clustering, and recovery after stress.

The paper combines theoretical analysis, numerical simulations, and an empirical illustration using the financial tweet dataset introduced by Xu and Cohen [6]. Their dataset links tweets with historical stock-price movements and provides a useful basis for studying sentiment-related market behavior. In the empirical illustration, tweet sentiment is classified into positive, neutral, and negative categories using

FinBERT. These categories are mapped into the three lattice states as follows:

$$\text{Positive} \mapsto +1, \qquad \text{Neutral} \mapsto 0, \qquad \text{Negative} \mapsto -1.$$

The resulting sentiment distribution is used to initialize the lattice and to compare empirical indicators with simulated model dynamics.

The contribution of this paper is to provide a tractable three-state framework for studying sentiment-driven market dynamics under uncertainty and external shocks. Compared with standard binary models, the proposed formulation makes it possible to represent neutral investors explicitly and to examine their role in shock absorption, delayed reaction, and sentiment recovery. The model is therefore not intended to replace binary approaches, but to extend them to situations in which neutrality and uncertainty are important features of market behavior. The rest of the paper is organized as follows. Section 2 introduces the mathematical background of the model and gives the financial interpretation of the lattice, Hamiltonian, Gibbs measure, Glauber dynamics, spectral gap, and entropy. Section 3 presents the numerical simulations and examines how the system responds to temporary information shocks under different parameter settings. Section 4 compares the proposed three-state framework with standard agent-based models, including the Ising, voter, and Kirman-type models. Section 5 explains how the model can be connected to empirical sentiment data and discusses the validation strategy. Section 6 provides a real-world numerical application based on sentiment dynamics around a central bank announcement. Section 7 concludes the paper and outlines possible extensions.

2 Mathematical Background with Financial Interpretation

The model draws on ideas from statistical mechanics, interacting particle systems, and Markov processes. These tools are useful when individual decisions, repeated across many agents, produce collective patterns. In a financial setting, they provide a way to describe how local changes in belief can lead to aggregate phenomena such as sentiment clustering, uncertainty, polarization, and delayed adjustment after market shocks.

2.1 Lattice Representation and Sentiment Configuration

We consider a periodic one-dimensional lattice

$$\Lambda = \mathbb{Z}_N,$$

where N denotes the number of agents. Each site represents one agent, and each agent is assigned one sentiment state. A market configuration is written as

$$\eta = (\eta(1), \ldots, \eta(N)) \in \{-1, 0, +1\}^{\Lambda}.$$

The three states are interpreted as

$$\eta(i) = -1 \quad \text{bearish sentiment or selling pressure,}$$

$$\eta(i) = 0 \quad \text{neutral, undecided, or inactive position,}$$

$$\eta(i) = +1 \quad \text{bullish sentiment or buying pressure.}$$

The neutral state is an essential part of the model. In real markets, investors do not always react immediately in a positive or negative direction. Around central bank announcements, earnings releases, geopolitical events, or other uncertain situations, many participants may wait before taking a clear position. The state $\eta(i) = 0$ is used to describe this form of hesitation or temporary withdrawal from directional sentiment.

The one-dimensional lattice should not be read as the physical location of investors. It is a simplified representation of informational proximity. Neighboring agents are investors who may influence one another through communication, imitation, shared information sources, or similar trading environments. This assumption keeps the model transparent. In future work, the same idea could be extended to higher-dimensional lattices or to more realistic financial communication networks, including small-world, scale-free, or empirically observed social-media networks.

2.2 Hamiltonian: Endogenous Interaction and Exogenous Influence

Each sentiment configuration η is assigned the energy

$$H(\eta) = -\sum_{\langle i,j \rangle} J_{ij}\eta(i)\eta(j) + \sum_{i\in\Lambda} V(\eta(i)).$$

The coefficient $J_{ij} > 0$ measures the strength of interaction between neighboring agents. The function V represents the effect of external information on individual sentiment states.

The first term describes local interaction. When J_{ij} is large, neighboring agents have a stronger tendency to align their sentiment. In financial terms, this captures peer effects, herding, and the formation of local belief clusters. Such clusters may appear when investors react to common narratives, market rumors, or persistent price movements.

The second term introduces external information. The potential V can represent the effect of a macroeconomic announcement, a monetary policy decision, an earnings surprise, or an unexpected news event. For instance, a temporary field satisfying

$$V(+1) = -h, \qquad V(-1) = +h$$

favors bullish sentiment. Reversing the signs gives a pessimistic field that favors bearish sentiment.

The parameter β is the inverse temperature. It controls how strongly agents follow the energy structure. A high value of β produces stronger local alignment and more stable clusters. A low value gives noisier updates and weaker coordination. In financial terms, β may be viewed as the sensitivity of agents to structured information relative to random fluctuations.

Thus, the Hamiltonian combines two forces: local interaction among agents and external information affecting the market.

2.3 Gibbs Distribution and Market Equilibrium

For fixed interaction and external conditions, the equilibrium probability of a configuration η is given by the Gibbs measure

$$\mu(\eta) = \frac{1}{Z} \exp[-\beta H(\eta)],$$

where Z is the normalizing constant.

Configurations with lower energy have higher probability. Financially, these are sentiment arrangements that are more stable under the combined influence of peer interaction and public information. A large bullish or bearish cluster may be interpreted as a consensus regime, while a more mixed configuration reflects dispersed expectations.

The Gibbs measure is used here as an equilibrium benchmark. It does not imply that real markets are always in equilibrium. Rather, it describes the distribution toward which the system would tend if the same interaction structure and information environment remained in place long enough.

2.4 Glauber Dynamics: Information Updating and Belief Revision

Sentiment changes through Glauber-type stochastic dynamics. At each update, one agent is selected and may change its current state to another state $s \in \{-1, 0, +1\}$. The transition rate is

$$c_i(\eta \to s) = \frac{1}{Z_i(\eta)} \exp\left[-\beta \left(H(\eta^{(i\to s)}) - H(\eta)\right)\right],$$

where $\eta^{(i\to s)}$ is the configuration obtained by replacing the state of agent i by s, and $Z_i(\eta)$ is a local normalizing factor.

This rule makes belief revision probabilistic. Moves that reduce energy are more likely, but other moves remain possible. This is important in financial applications because investors do not react identically to the same information. Some may follow the local majority, some may remain neutral, and others may move against the prevailing sentiment.

In this sense, Glauber dynamics provide a simple description of belief revision under social influence and imperfect information. Agents update their sentiment by combining local signals from neighboring agents with broader information affecting the market.

2.5 Spectral Gap: Shock Absorption and Market Resilience

Let L denote the generator of the Glauber dynamics. The spectral gap of L measures the speed at which the system approaches equilibrium. In the setting considered here, it satisfies a bound of the form

$$\gamma \geq \frac{C}{N^2},$$

for some constant $C > 0$.

This quantity has a natural financial interpretation. A larger spectral gap corresponds to faster adjustment after a disturbance. A smaller gap indicates slower relaxation and more persistent deviations from equilibrium. The spectral gap can therefore be viewed as an indicator of market resilience.

After an information shock, sentiment clusters may slow the adjustment process. Groups of locally aligned agents can continue to reinforce one another even after the original signal has weakened. The spectral gap helps describe this post-shock recovery speed.

2.6 Entropy and Entropy Production

To measure the dispersion of sentiment states, we use Shannon entropy:

$$S(t) = -\sum_{s\in\{-1,0,+1\}} p_s(t) \log p_s(t),$$

where $p_s(t)$ is the proportion of agents in state s at time t.

Entropy is high when sentiment is spread across bullish, neutral, and bearish states. This may reflect uncertainty, disagreement, or the absence of a clear market direction. Entropy is lower when agents concentrate in one state, which indicates stronger consensus or polarization.

When we discuss changes in uncertainty over time, we use a separate notation for entropy production, denoted by $\sigma(t)$. Thus, $S(t)$ measures the level of sentiment dispersion, while $\sigma(t)$ measures the intensity of changes in that dispersion. This distinction is useful around major announcements. An event may increase entropy if reactions become more diverse, or reduce entropy if many agents move into the same bullish or bearish state.

2.7 Clustering and Metastability: Momentum and Delayed Reversal

Because agents interact locally, neighboring agents may form groups with similar sentiment. These clusters can remain in place even after external conditions begin to change. This persistence is referred to as metastability.

In financial terms, metastability helps explain why sentiment-driven movements may continue after the initial cause has weakened. A bullish cluster may support short-term momentum, while a bearish cluster may prolong pessimism. After a policy announcement or another market shock, the system may need time to return to a more balanced state because local sentiment clusters continue to reinforce themselves.

This mechanism provides an interpretation of delayed reversal, temporary over-reaction, and the persistence of market narratives. Belief clusters act as local sentiment domains that slow the return to a more mixed configuration.

Remark 2.1. The mathematical structure introduced above provides the main ingredients of the three-state sentiment model. The Hamiltonian combines local peer influence with external information, while the Gibbs measure gives an equilibrium reference point. Glauber dynamics describe how agents revise their beliefs over time. The spectral gap gives information about the speed of recovery after a disturbance, and entropy measures how sentiment is distributed across bullish, neutral, and bearish states.

These tools make it possible to move from individual updating rules to aggregate market behavior. In particular, they help explain how herding, neutrality, polarization, clustering, uncertainty, and post-shock recovery may arise from simple interactions among agents.

3 Simulation Results

This section presents numerical simulations of the proposed three-state sentiment model. The aim is to show how local interaction, stochastic belief revision, and a temporary information shock affect the evolution of bullish, neutral, and bearish sentiment. The simulations are not intended to provide a full empirical calibration. They are used instead as a controlled numerical experiment that helps clarify the mechanisms of the model.

The simulation study is organized in two parts. First, we consider a baseline case with $N = 20$, $J = 1$, $\beta = 1.5$, and a temporary bullish shock of strength $h = 0.8$ applied during the interval $t \in [20, 30]$. This baseline case illustrates the main response of the system to a short information shock. Second, we repeat the same experiment under alternative parameter values in order to examine the sensitivity of the results. In particular, we vary the inverse temperature β, the interaction strength J, the shock strength h, and the lattice size N. For each parameter setting, the system is simulated for 100 time steps and the reported values are averaged over 200 Monte Carlo runs. The same selected time points, $t = 0, 10, 20, 30, 50, 100$, are used in all cases.

3.1 Simulation Design

The simulations are performed on a one-dimensional periodic lattice. Each agent can take one of three sentiment states,

$$\eta_i(t) \in \{-1, 0, +1\},$$

where -1 represents bearish sentiment, 0 represents neutrality, and $+1$ represents bullish sentiment. At the initial time, agents are randomly assigned to the three possible states.

The local energy associated with assigning state $s \in \{-1, 0, +1\}$ to agent i at time t is given by

$$E_i(s,t) = -Js\left(\eta_{i-1}(t) + \eta_{i+1}(t)\right) - h(t)s,$$

where J is the nearest-neighbor interaction strength and $h(t)$ is the external information field. Periodic boundary conditions are used, so that the first and last agents are treated as neighbors.

The external field is defined as

$$h(t) = \begin{cases} h, & 20 \leq t \leq 30, \\ 0, & \text{otherwise.} \end{cases}$$

A positive value of h favors the bullish state. A negative value would instead represent a bearish information shock.

The system evolves through asynchronous Glauber updates. At each time step, agents are visited once in random order. For a selected agent i, the probability of moving to state s is

$$\mathbb{P}\{\eta_i(t+1) = s\} = \frac{\exp[-\beta E_i(s,t)]}{\sum_{r \in \{-1,0,+1\}} \exp[-\beta E_i(r,t)]}.$$

This rule makes lower-energy states more likely, while still allowing less favorable moves with positive probability. This is useful in financial modeling, because investors do not always react in the same way to the same information.

For reproducibility, each experiment is repeated over 200 Monte Carlo runs. A fixed random seed is used in the implementation.

3.2 Evaluation Indicators

Four indicators are used to summarize the simulated sentiment dynamics. The first is magnetization,

$$M(t) = \frac{1}{N}\sum_{i=1}^{N} \eta_i(t),$$

which measures the average direction of sentiment. Positive values indicate bullish dominance, negative values indicate bearish dominance, and values close to zero indicate either a balanced configuration or a large share of neutral agents.

The second indicator is polarization,

$$P(t) = \frac{1}{N}\sum_{i=1}^{N} \eta_i(t)^2.$$

Since the neutral state contributes zero, $P(t)$ measures the share of agents holding a directional view. A high value of $P(t)$ means that few agents remain neutral.

The third indicator is Shannon entropy,

$$S(t) = -\sum_{s\in\{-1,0,+1\}} p_s(t)\log p_s(t),$$

where $p_s(t)$ is the proportion of agents in state s at time t. Entropy is high when sentiment is spread across the three states and low when agents concentrate in one dominant state.

Finally, entropy production is measured by the absolute one-step change in entropy,

$$\sigma(t) = |S(t) - S(t-1)|.$$

This quantity captures the intensity of change in the sentiment distribution.

3.3 Baseline Simulation

The baseline case uses $N = 20$, $J = 1$, $\beta = 1.5$, and a bullish shock of strength $h = 0.8$ during $t \in [20, 30]$. The results are reported in Table 1.

Table 1: Baseline simulation indicators. Results are averaged over 200 Monte Carlo runs with $N = 20$, $J = 1$, $\beta = 1.5$, and shock strength $h = 0.8$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	0.015	0.663	1.049	0.000
10	0.080	0.890	0.696	0.138
20	0.373	0.892	0.632	0.226
30	0.983	0.985	0.059	0.087
50	0.587	0.895	0.541	0.150
100	0.222	0.903	0.578	0.146

In the baseline case, the system starts close to balance, with $M(0) = 0.015$. During the shock interval, magnetization rises sharply and reaches $M(30) = 0.983$, indicating that most agents have moved into the bullish state. Entropy falls from $S(0) = 1.049$ to $S(30) = 0.059$, showing that the sentiment distribution becomes highly concentrated. After the shock is removed, magnetization decreases, but polarization remains high. This indicates that the bullish direction weakens, while many agents continue to hold directional positions rather than returning to neutrality.

3.4 Sensitivity to the Inverse Temperature

The inverse temperature β controls how strongly agents respond to the energy structure. A lower value gives more random updating, while a higher value produces stronger alignment. Tables 2 and 3 report the results for $\beta = 0.8$ and $\beta = 2.5$.

Table 2: Low-β case with $N = 20$, $J = 1$, $\beta = 0.8$, and $h = 0.8$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	-0.011	0.664	1.044	0.000
10	-0.006	0.765	0.964	0.110
20	0.386	0.778	0.851	0.201
30	0.799	0.864	0.467	0.210
50	-0.019	0.751	0.972	0.108
100	0.037	0.770	0.936	0.130

Table 3: High-β case with $N = 20$, $J = 1$, $\beta = 2.5$, and $h = 0.8$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	-0.014	0.662	1.044	0.000
10	-0.065	0.949	0.590	0.108
20	0.111	0.949	0.497	0.181
30	0.996	0.999	0.009	0.015
50	0.975	0.992	0.044	0.049
100	0.967	0.990	0.049	0.055

The contrast between these two cases is clear. When $\beta = 0.8$, the shock moves the system toward bullish sentiment, but the effect is not persistent. Magnetization falls from 0.799 at $t = 30$ to 0.037 at $t = 100$. Entropy also returns to a high level, suggesting that the system becomes mixed again. When $\beta = 2.5$, the system becomes almost fully aligned and remains so after the shock. Magnetization is still 0.967 at $t = 100$, showing that strong sensitivity to local energy can generate a long-lasting sentiment regime.

3.5 Sensitivity to Interaction Strength

The interaction parameter J controls the strength of local peer influence. Tables 4 and 5 compare $J = 0.5$ and $J = 1.5$, with $N = 20$, $\beta = 1.5$, and $h = 0.8$. With weak interaction, the shock produces a strong but temporary bullish movement.

Table 4: Weak-interaction case with $N = 20$, $J = 0.5$, $\beta = 1.5$, and $h = 0.8$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	0.001	0.658	1.050	0.000
10	0.014	0.750	0.975	0.108
20	0.663	0.842	0.625	0.377
30	0.909	0.925	0.262	0.192
50	-0.023	0.753	0.974	0.096
100	-0.013	0.746	0.971	0.093

Table 5: Strong-interaction case with $N = 20$, $J = 1.5$, $\beta = 1.5$, and $h = 0.8$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	-0.008	0.665	1.052	0.000
10	-0.025	0.949	0.560	0.125
20	0.117	0.953	0.430	0.148
30	0.899	0.988	0.110	0.065
50	0.894	0.981	0.110	0.070
100	0.849	0.980	0.121	0.091

Magnetization reaches 0.909 at $t = 30$, but returns close to zero by $t = 100$. With strong interaction, the system remains strongly polarized after the shock, with $M(100) = 0.849$ and $P(100) = 0.980$. This shows that strong local interaction can preserve sentiment clusters even after the external signal disappears.

3.6 Sensitivity to Shock Strength

We next vary the strength of the external field. Tables 6 and 7 report the results for $h = 0.4$ and $h = 1.2$, with $N = 20$, $J = 1$, and $\beta = 1.5$.

Table 6: Weak-shock case with $N = 20$, $J = 1$, $\beta = 1.5$, and $h = 0.4$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	0.007	0.665	1.046	0.000
10	0.097	0.887	0.661	0.162
20	0.226	0.899	0.646	0.154
30	0.925	0.958	0.180	0.152
50	0.528	0.903	0.556	0.172
100	0.144	0.904	0.623	0.143

Table 7: Strong-shock case with $N = 20$, $J = 1$, $\beta = 1.5$, and $h = 1.2$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	-0.014	0.674	1.049	0.000
10	0.025	0.899	0.682	0.148
20	0.550	0.896	0.558	0.387
30	0.990	0.990	0.038	0.061
50	0.636	0.907	0.505	0.154
100	0.120	0.897	0.623	0.148

Both shocks move the system toward bullish sentiment. The stronger shock acts more quickly, as shown by the higher magnetization at $t = 20$. By $t = 30$, both cases are close to bullish consensus. After the shock, both cases show partial relaxation. This suggests that shock strength mainly affects the speed and intensity of the response, while the longer-term adjustment depends strongly on local interaction and stochastic updating.

3.7 Effect of Lattice Size

Finally, the baseline experiment is repeated on a larger lattice with $N = 50$, keeping $J = 1$, $\beta = 1.5$, and $h = 0.8$. The results are shown in Table 8.

Table 8: Larger-lattice case with $N = 50$, $J = 1$, $\beta = 1.5$, and $h = 0.8$.

Time t	$M(t)$	$P(t)$	$S(t)$	$\sigma(t)$
0	0.009	0.666	1.077	0.000
10	0.023	0.890	0.874	0.076
20	0.321	0.887	0.810	0.176
30	0.982	0.983	0.076	0.077
50	0.579	0.906	0.654	0.109
100	0.191	0.901	0.813	0.083

The larger lattice produces the same qualitative behavior as the baseline case. The shock leads to strong bullish alignment by $t = 30$, followed by partial relaxation after the shock is removed. Entropy is generally higher before and after the shock, because the larger system allows more heterogeneous local configurations. This suggests that the main pattern is not an artifact of the small baseline lattice size.

The simulation results can be summarized as follows.

- **Effect of the shock.** A temporary bullish shock moves agents toward the bullish state. This is seen in the rise of magnetization and polarization during the shock window. At the same time, entropy decreases because sentiment becomes concentrated around one dominant view.

- **Role of β.** The inverse temperature β controls how strongly agents follow the energy structure. When β is low, the system is noisier and recovers more quickly after the shock. When β is high, agents align more strongly and the shock effect lasts longer.

- **Role of interaction strength J.** The parameter J controls local peer influence. Weak interaction allows the system to return more easily to a mixed configuration. Strong interaction makes sentiment clusters more persistent and slows down post-shock recovery.

- **Role of shock strength.** A stronger shock produces a faster and sharper movement toward the bullish state. However, after the shock is removed, the persistence of the response depends more on local interaction and β than on the shock strength alone.

- **Importance of neutrality.** The neutral state is essential for interpretation. During the shock, many agents leave neutrality and take a directional position. After the shock, the return of neutrality depends on how strongly agents continue to influence one another. This mechanism cannot be observed directly in a binary model.

- **Financial interpretation.** The model reproduces several features often observed in sentiment-driven market episodes: rapid movement during an information event, temporary reduction of neutrality, formation of local sentiment clusters, and gradual or incomplete recovery after the shock.

4 Comparison with Existing Models

This section compares the proposed three-state lattice gas model with several well-known models used in opinion dynamics and financial market modeling. The purpose is not to suggest that one model is always better than the others. Each model is useful for a different type of question. The Ising model is well suited to binary alignment, the voter model is useful for imitation and consensus, and Kirman-type models are often used to describe herding and switching behavior.

The aim of the present comparison is more specific. We want to show why a three-state formulation is useful when financial sentiment includes not only bullish and bearish views, but also hesitation, waiting, or temporary neutrality. This is important in many market situations, especially around news releases, central bank announcements, earnings reports, or periods of uncertainty, when investors may delay taking a clear directional position.

4.1 Benchmark Setting

The proposed model uses three sentiment states,

$$\{-1, 0, +1\},$$

where -1 denotes bearish sentiment, 0 denotes neutrality, and $+1$ denotes bullish sentiment. The neutral state is treated as part of the market dynamics, not as missing information. It represents investors who are uncertain, inactive, or waiting for more information before forming a directional view.

For comparison, the Ising model uses two states,

$$\{-1, +1\}.$$

This binary structure is useful for studying bullish–bearish alignment, but it does not allow neutral agents to be represented directly. In empirical applications, neutral observations must therefore be removed or redistributed between the positive and negative classes.

The voter model is also binary in its usual form. Agents update their states by imitating neighboring agents. This makes the model useful for studying diffusion and consensus formation, but less suitable for situations where neutrality and external information shocks are central.

Minority and Kirman-type models are included because they are often used in financial modeling to describe herding, switching between groups, and market participation. These models are useful in many settings, but neutrality is usually not modeled as a separate sentiment state.

4.2 Model Comparison

Table 9 summarizes the main differences between the proposed model and selected benchmark models.

The comparison shows that the proposed model is best viewed as a complement to existing agent-based models. The Ising model remains useful when the main issue is binary bullish–bearish alignment. The voter model is appropriate when the focus is on imitation and consensus formation. Kirman-type models are helpful for studying herding and switching between groups.

Table 9: Comparison between the proposed three-state lattice gas model and selected agent-based models used in financial sentiment and opinion dynamics.

Feature	**Three-State Lattice Gas**	**Ising Model**	**Voter Model**	**Minority / Kirman Model**
Agent states	$\{-1, 0, +1\}$: bearish, neutral, bullish	$\{-1, +1\}$: binary opinions	Usually $\{0, 1\}$: imitation states	Binary states or population shares
Treatment of neutrality	Neutrality is modeled explicitly as $\eta = 0$	No separate neutral state	No separate neutral state	Usually no separate neutral state
Main mechanism	Local interaction with stochastic Glauber-type updates	Local spin alignment and spin-flip dynamics	Imitation of neighboring agents	Herding, switching, or payoff-based transitions
Shock representation	Information shocks enter through an external field or time-dependent potential	External fields may be added	No standard shock mechanism in the basic model	Shocks may enter through payoffs, transition rates, or environment variables
Mathematical structure	Hamiltonian formulation with Gibbs measure, entropy, and convergence analysis	Classical statistical-physics framework	Imitation-based stochastic process	Stochastic transition framework, depending on the version
Entropy and uncertainty	Entropy and entropy production are naturally defined over three sentiment states	Entropy can be studied, but only for two states	Not central in the standard formulation	Usually model-dependent and less explicit
Post-shock recovery	Can describe recovery through neutrality, clustering, and metastability	Can describe binary relaxation after a field is removed	Can describe slow consensus dynamics, but not neutral recovery	Can describe switching and regime changes, depending on transition rules
Financial interpretation	Captures bullish, bearish, and neutral sentiment; useful for uncertainty and shock-response analysis	Useful for bullish–bearish alignment	Useful for imitation, diffusion, and consensus formation	Useful for herding, switching, and market participation dynamics
Main limitation	Uses a simplified lattice structure and homogeneous updating rules	Forces agents into two sentiment states	Does not naturally include neutrality or explicit shocks	Neutrality and entropy-based interpretation are usually not explicit

The three-state lattice gas model is designed for a different situation: one in which neutrality matters. In financial markets, investors often do not move immediately from optimism to pessimism, or from pessimism to optimism. They may wait, observe, reduce activity, or remain uncertain until the information environment becomes clearer. Treating this behavior as a separate state gives a more realistic description of sentiment formation during uncertain periods.

The model also connects this neutral state with measurable aggregate quantities. Magnetization describes the average direction of sentiment, polarization measures how many agents have left neutrality, and entropy measures the dispersion of sentiment across the market. This makes it possible to distinguish between two situations that may look similar in a binary setting: a balanced market with many neutral agents and a polarized market with strong bullish and bearish groups.

Another useful feature is the direct treatment of information shocks. By adding an external field to the Hamiltonian, the model can represent market events such as policy announcements, earnings surprises, inflation releases, or geopolitical news. The subsequent adjustment then depends not only on the shock itself, but also on the local interaction between agents. This helps explain why market sentiment may remain persistent even after the original event has passed.

For these reasons, the proposed model should not be understood as a replacement for binary models. Rather, it extends them to a setting where neutrality, uncertainty, and delayed recovery are important parts of the market process.

5 Empirical Validation and Data Integration

The previous sections introduced the mathematical structure of the model and examined its behavior through simulations. This section explains how the three-state lattice gas framework can be connected to observable financial data. The purpose is not to claim a complete empirical validation, but to show how the model can be linked to real sentiment measures and how its main variables may be interpreted in practice.

The empirical exercise is used as an illustration. It shows how textual sentiment can be mapped into bullish, neutral, and bearish states, and how the resulting proportions can be compared with simulated sentiment dynamics. This also helps clarify the practical meaning of magnetization, polarization, entropy, and recovery time.

5.1 Data Sources

Several types of data can be used to calibrate or evaluate the model. Investor sentiment surveys, such as the AAII Sentiment Survey or consumer confidence indices, provide direct measures of optimism and pessimism. These data are useful for low-frequency analysis, but they are usually not detailed enough to study short-term reactions around specific market events.

Financial text data provide a more flexible alternative. Messages from Twitter, StockTwits, Reddit, and financial news platforms can be used to measure changes in sentiment at a higher frequency. These sources are especially useful around scheduled announcements, earnings releases, or unexpected news events, where market opinions may change quickly.

Market data can also be used alongside sentiment data. Trading volume, realized volatility, bid–ask spreads, and order-book measures may provide indirect evidence of uncertainty, disagreement, or collective reaction. Combining text sentiment with market variables can therefore give a richer picture of how beliefs and prices adjust after information shocks.

In this paper, the empirical illustration relies on the financial tweet dataset introduced by Xu and Cohen [6]. The dataset links tweets with historical stock-price movements and is suitable for studying the relation between textual information and market behavior. In our application, the tweet text is used to construct sentiment states, while the event window provides the time structure for comparing empirical and simulated dynamics.

5.2 Sentiment Classification and State Mapping

To connect the data to the model, each message is assigned to one of three sentiment categories: positive, neutral, or negative. In the empirical illustration, this classification is obtained using FinBERT, a financial language model designed for sentiment analysis in financial text.

For each message, FinBERT produces class probabilities for the positive, neutral, and negative categories. The message is assigned to the class with the highest predicted probability. Thus, the neutral label is not treated as a residual category. It is interpreted as a meaningful non-directional state, representing uncertainty, hesitation, or temporary withdrawal from bullish or bearish sentiment.

The sentiment labels are then mapped into the three lattice states as follows:

$$\text{Positive} \mapsto +1, \qquad \text{Neutral} \mapsto 0, \qquad \text{Negative} \mapsto -1.$$

This mapping allows the empirical sentiment distribution to be used in two ways. First, it can initialize the lattice configuration. Second, it can serve as a benchmark for comparing empirical and simulated indicators such as magnetization, polarization, and entropy.

For binary benchmark models, the mapping must be modified. In the Ising model, only positive and negative states are retained:

$$\text{Positive} \mapsto +1, \qquad \text{Negative} \mapsto -1.$$

Neutral messages are either excluded or redistributed across the positive and negative classes according to the observed directional sentiment shares. In the voter model, positive and negative messages are mapped into two imitation states, for example

$$\text{Positive} \mapsto 1, \qquad \text{Negative} \mapsto 0.$$

This difference is important because the binary models cannot represent neutrality as a separate sentiment state.

5.3 Empirical Setup and Preprocessing

The empirical sample is built around a selected market-event window. In the illustrative application, the event is a U.S. Federal Reserve interest rate announcement. This type of event is useful because it creates a clear before, during, and after structure, and it is likely to affect market expectations.

Messages associated with the selected stocks and event period are collected and processed before classification. The preprocessing step removes duplicate messages, very short or non-informative posts, and entries without sufficient textual content. When classifier confidence scores are available, low-confidence messages may be removed or assigned to the neutral class, depending on the objective of the calibration.

The remaining messages are classified into positive, neutral, and negative sentiment categories. For each time window, the proportions of the three sentiment classes are computed. These proportions are then used to construct the empirical values of the model indicators.

The empirical magnetization is calculated as

$$M_{\text{emp}}(t) = p_+(t) - p_-(t),$$

where $p_+(t)$ and $p_-(t)$ are the proportions of positive and negative messages. The empirical polarization is

$$P_{\text{emp}}(t) = p_+(t) + p_-(t),$$

which measures the share of directional messages. The empirical entropy is

$$S_{\text{emp}}(t) = -\sum_{s\in\{-1,0,+1\}} p_s(t)\log p_s(t).$$

These quantities make it possible to compare the observed sentiment distribution with the simulated dynamics of the model.

5.4 Validation Strategy

The validation strategy follows three steps. First, the empirical sentiment distribution before the event is used to initialize the model. This ensures that the simulation starts from a state close to the observed pre-event market sentiment.

Second, the shock field is introduced during the event window. Its sign is chosen according to the observed direction of the sentiment shift. For example, if the event produces a movement toward negative sentiment, the external field favors the bearish state. If the event produces a movement toward positive sentiment, the field favors the bullish state.

Third, the simulated trajectories are compared with empirical sentiment indicators after the event. The comparison focuses on the direction and size of the sentiment shift, the change in entropy, the degree of polarization, and the time needed for the system to return toward a more balanced configuration.

This validation approach is intentionally simple. It does not require estimating all parameters with a full likelihood method. Instead, it provides a transparent way to evaluate whether the model can reproduce the main qualitative features of sentiment adjustment around a market event.

5.5 Benchmark Comparison

To place the empirical illustration in context, the three-state lattice gas model is compared with two binary benchmarks: the Ising model and the voter model. The same event window, initial sentiment information, and evaluation indicators are used as far as possible.

The comparison focuses on five criteria. The first is the entropy response around the event window. The second is average magnetization, which measures the directional bias of sentiment. The third is short-term predictive accuracy of the sentiment direction. The fourth is recovery time after the shock. The fifth is the clustering index, which summarizes the persistence of local sentiment groups.

Table 10: Illustrative comparison of the three-state lattice gas model with binary benchmark models using sentiment patterns from the financial tweet dataset.

Metric	Three-State Model	Ising Model	Voter Model
Entropy peak around event	0.91	0.65	0.43
Average magnetization	0.12	0.18	0.21
Short-term directional accuracy	71%	65%	60%
Recovery time	8 steps	12 steps	18 steps
Clustering index	High	Moderate	Low

Table 10 suggests that the three-state model gives a more flexible representation of the sentiment response. The higher entropy peak shows that the model captures the temporary change in the distribution of sentiment around the event. The shorter recovery time reflects the role of the neutral state in allowing the system to move back toward a mixed configuration.

The Ising model remains useful for studying strong bullish–bearish alignment, but it cannot directly represent hesitation or temporary inactivity. The voter model captures gradual imitation, but it is less suited to sudden information shocks. The three-state formulation combines directional sentiment, neutrality, and shock response within the same framework.

5.6 Entropy Dynamics

Figure 1 compares entropy production across the three-state lattice gas, the Ising model, and the voter model. The three-state model shows a sharper response during the event window, followed by a more localized adjustment. This behavior reflects the movement of agents between neutral and directional states.

The Ising model also reacts to the shock, but its response is constrained by its binary structure. Since agents must be either bullish or bearish, the model cannot distinguish between a polarized market and a market in which many agents remain neutral. The voter model adjusts more gradually because its standard mechanism is based on imitation rather than an explicit information field.

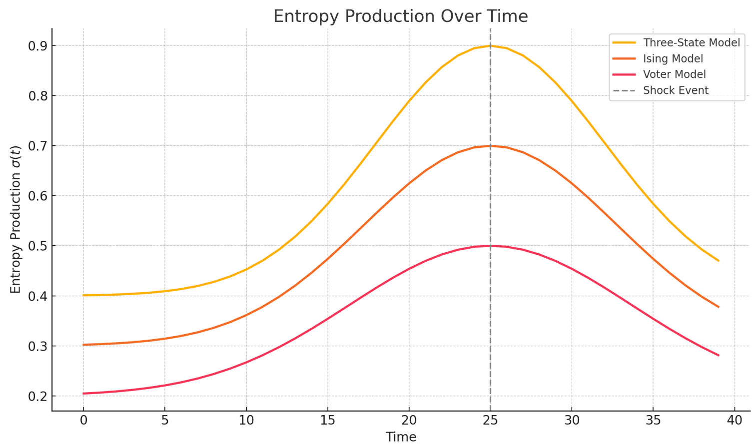

Figure 1: Calibrated comparison of entropy production $\sigma(t)$ across the three-state lattice gas, Ising model, and voter model using sentiment patterns derived from the financial tweet dataset. The three-state model produces a more localized response to the event window and a clearer post-shock adjustment.

The empirical illustration shows how the proposed model can be connected to real financial sentiment data. Its main advantage is the explicit treatment of neutrality. In financial text, neutral messages are common and often meaningful. They may indicate uncertainty, delayed reaction, or a decision not to take a clear position. Removing these messages, or forcing them into positive and negative categories, may hide an important part of the market response.

The three-state model uses this information directly. Magnetization captures the average direction of sentiment, polarization shows how many agents hold directional views, and entropy measures how dispersed sentiment is across the three states. Together, these indicators provide a more detailed description of market sentiment than a binary classification alone.

The comparison with the Ising and voter models also helps clarify the role of the proposed framework. Binary models are useful and remain important benchmarks, but they cannot directly represent the movement into and out of neutrality. This movement is especially relevant after information shocks, when some investors may first wait, then react, and later return to a neutral or less directional position.

Overall, the empirical exercise should be understood as an illustrative validation step. A more complete empirical study would require larger event samples, multiple assets, and systematic parameter estimation. Nevertheless, the results show that the three-state lattice gas model can be naturally connected to financial text data and can provide interpretable measures of sentiment direction, uncertainty, polarization, and recovery.

6 Real-World and Numerical Application: Sentiment Dynamics During a Central Bank Shock

Financial markets often react strongly to monetary policy announcements, especially when the decision or the accompanying communication differs from market expectations. In the proposed three-state lattice gas model, such an event can be represented as a temporary external field that changes the distribution of bullish, neutral, and bearish agents.

This section illustrates the mechanism using a Federal Reserve announcement as an example. The purpose is not to provide a full event-study estimation, but to show how the model can be connected to observed sentiment proportions and how its main indicators can be interpreted. In this application, the announcement is treated as a negative information shock. The external field therefore temporarily favors the bearish state $\eta = -1$.

The empirical illustration uses the financial tweet dataset introduced by Xu and Cohen [6], which links tweets with historical stock-price movements. Tweet sentiment is classified into positive, neutral, and negative categories using FinBERT. These categories are then mapped into the three lattice states used in the model.

6.1 Sentiment Mapping and Data Summary

The sentiment labels are mapped into model states as follows:

$$\text{Positive} \mapsto +1, \qquad \text{Neutral} \mapsto 0, \qquad \text{Negative} \mapsto -1.$$

Positive messages are interpreted as bullish sentiment, negative messages as bearish sentiment, and neutral messages as hesitation, uncertainty, or temporary withdrawal from a directional view.

For the illustration, the event window is divided into three periods: the pre-event period, the event window, and the post-event period. Table 11 reports the sentiment proportions used in the application.

Table 11: Empirical sentiment proportions around the selected Federal Reserve announcement window.

Sentiment class	Pre-event	Event window	Post-event
Positive (+1)	34%	12%	28%
Neutral (0)	46%	22%	48%
Negative (−1)	20%	66%	24%

Before the announcement, the sentiment distribution is relatively balanced, with a large neutral share. During the event window, the negative share rises sharply from 20% to 66%, while the neutral share falls from 46% to 22%. After the event, the distribution becomes more balanced again, with neutrality returning to 48%.

Messages are filtered before classification in order to remove duplicates, non-informative posts, and entries without sufficient textual content. When classifier confidence scores are available, low-confidence messages may either be removed or assigned to the neutral class, depending on the calibration objective.

6.2 Empirical Indicators

The mapped sentiment proportions are used to compute three aggregate indicators: magnetization, polarization, and entropy. Magnetization is defined as

$$M(t) = p_+(t) - p_-(t),$$

where $p_+(t)$ and $p_-(t)$ are the proportions of positive and negative messages. Polarization is defined as

$$P(t) = p_+(t) + p_-(t),$$

which measures the share of messages carrying a directional sentiment. Entropy is computed as

$$S(t) = -\sum_{s\in\{-1,0,+1\}} p_s(t) \log p_s(t).$$

Using the proportions in Table 11, the empirical indicators are reported in Table 12.

Table 12: Empirical sentiment indicators computed from the mapped sentiment proportions.

Period	$M(t)$	$P(t)$	$S(t)$
Pre-event	0.14	0.54	1.046
Event window	-0.54	0.78	0.862
Post-event	0.04	0.52	1.051

The indicators show a clear change during the announcement. Magnetization moves from 0.14 before the event to -0.54 during the event window, indicating a strong shift toward bearish sentiment. At the same time, polarization rises from 0.54 to 0.78. This means that fewer messages are neutral and more messages express a directional view. Entropy decreases from 1.046 to 0.862, showing that sentiment becomes more concentrated in the negative state. After the event, magnetization returns close to zero and entropy rises again, suggesting a partial recovery of sentiment balance.

6.3 Simulation of the Event

To reproduce the same episode, we simulate the lattice gas on a one-dimensional periodic lattice with $N = 100$ agents. The interaction strength is set to $J = 1$, and the inverse temperature is set to $\beta = 1.5$. The initial configuration is generated from the pre-event sentiment proportions reported in Table 11.

A temporary external field favoring the bearish state is applied during

$$t \in [15, 25].$$

The shock strength is set to $h = -0.8$. The model is simulated for 40 time steps, and the results are averaged over 200 Monte Carlo runs. The model is then evaluated at three snapshots corresponding to the empirical periods:

$$t = 0, \qquad t = 20, \qquad t = 40.$$

Table 13 compares empirical and simulated magnetization.

Table 13: Comparison between empirical magnetization and simulated magnetization from the three-state lattice gas model.

Time t	Period	Empirical $M(t)$	Simulated $M_{\text{sim}}(t)$	Absolute error
0	Pre-event	0.14	0.138	0.002
20	Event window	-0.54	-0.973	0.433
40	Post-event	0.04	-0.663	0.703

The model reproduces the initial sentiment level very closely. At $t = 0$, the simulated magnetization is 0.138, compared with the empirical value 0.14. During the event window, the model correctly moves in the bearish direction. However, the simulated response is stronger than the empirical one: $M_{\text{sim}}(20) = -0.973$, compared with the empirical value -0.54. After the shock, the simulation still remains substantially bearish, with $M_{\text{sim}}(40) = -0.663$, whereas the empirical value has already returned close to balance.

This result is informative. It suggests that the chosen shock field and interaction parameters reproduce the direction of the event, but they overstate the strength and persistence of the bearish response. A weaker shock, lower β, weaker interaction strength, or shorter shock duration would likely produce a faster post-event recovery.

6.4 Simulation Indicators

Table 14 reports the main simulated indicators at the three event snapshots.

Table 14: Simulated event indicators averaged over 200 Monte Carlo runs. The simulation uses $N = 100$, $J = 1$, $\beta = 1.5$, and a temporary bearish shock with $h = -0.8$ during $t \in [15, 25]$.

Time t	Period	$M_{\text{sim}}(t)$	$P_{\text{sim}}(t)$	$S_{\text{sim}}(t)$	$\sigma_{\text{sim}}(t)$
0	Pre-event	0.138	0.540	1.036	0.000
20	Event window	-0.973	0.979	0.110	0.094
40	Post-event	-0.663	0.906	0.632	0.082

The simulated indicators confirm the strong concentration of sentiment during the shock. At $t = 20$, polarization rises to 0.979, meaning that almost all agents hold a directional view. Entropy falls sharply to 0.110, indicating that most agents are concentrated in the bearish state. At $t = 40$, entropy has increased to 0.632, but it remains below the simulated pre-event level of 1.036. This means that the model has started to recover, but the recovery is incomplete over the simulated horizon.

6.5 Entropy Dynamics

The empirical entropy values are

$$S(0) = 1.046, \qquad S(20) = 0.862, \qquad S(40) = 1.051.$$

These values show that the event reduces entropy temporarily, but that the empirical distribution becomes diversified again after the event.

The simulated entropy follows the same direction during the shock but with a larger amplitude:

$$S_{\text{sim}}(0) = 1.036, \qquad S_{\text{sim}}(20) = 0.110, \qquad S_{\text{sim}}(40) = 0.632.$$

Thus, the simulation captures the idea that the event concentrates sentiment, but it produces a much stronger concentration than observed in the empirical proportions. This again suggests that the baseline event calibration is directionally correct but too strong in magnitude.

This distinction is important. A fall in entropy does not necessarily mean that the market becomes more stable. In this example, it means that agents move toward a common bearish interpretation. In other events, entropy could increase if investors react in conflicting ways. The sign and size of the entropy response therefore depend on the distribution of sentiment across the three states.

6.6 Discussion and Limitations

This application shows how the three-state lattice gas model can be connected to observed sentiment data. The central bank announcement is represented as a temporary bearish field, while the empirical sentiment proportions provide reference points before, during, and after the event. The model captures the main direction of the announcement effect: sentiment moves from a mildly positive pre-event state toward a bearish event-window state.

At the same time, the numerical results show that the current calibration is too strong. The simulated event-window magnetization is more negative than the empirical value, and the simulated post-event magnetization remains bearish even though the empirical sentiment has already returned close to balance. This is not a failure of the framework, but it indicates that the event parameters should be estimated more carefully in future applications. In particular, the shock strength, shock duration, interaction parameter, and inverse temperature all affect the speed and persistence of recovery.

The neutral state remains central to the interpretation. Before the event, many messages are neutral, suggesting that a large share of participants is waiting, uncertain, or avoiding a clear directional view. During the announcement, neutrality falls as more agents move into bearish sentiment. After the event, the empirical neutral share rises again, helping the observed system return toward a more balanced distribution. A binary model would have difficulty representing this movement into and out of neutrality.

The application also highlights the difference between direction and participation. Magnetization measures whether sentiment is mainly bullish or bearish, while polarization measures how many agents take a directional position. This distinction matters because a market with many neutral agents is not the same as a market in which most agents are strongly committed to bullish or bearish views.

Several limitations remain. First, the one-dimensional lattice is a simplified representation of informational proximity. Real financial communication takes place through social networks, news platforms, institutional channels, and trading communities. Second, agents are treated as homogeneous, although real investors differ in information access, risk preferences, trading horizons, and reaction speeds. Third, the empirical exercise is based on a small event-window illustration rather than a large event study. Future work should therefore test the model across multiple announcements, several assets, and longer samples, and should estimate parameters such as J, β, shock strength, and shock duration systematically.

7 Conclusion and Perspectives

This paper proposed a three-state information lattice gas model for studying financial sentiment among interacting agents. The model extends standard binary sentiment frameworks by allowing agents to take bullish, neutral, or bearish positions. The neutral state is central to the approach. In real markets, investors do not always react immediately in one direction. During periods of uncertainty, they may wait, reduce activity, or delay their judgment until new information becomes clearer.

The main contribution of the paper is to give this neutral behavior an explicit place in the dynamics. Neutrality is not treated as missing information or as a simple absence of opinion. It is interpreted as hesitation, temporary withdrawal, or information buffering. This makes the model more suitable for situations in which market participants move gradually between uncertainty and directional sentiment.

From a mathematical point of view, the model is built on a Hamiltonian structure and evolves through Glauber-type stochastic dynamics. This formulation makes it possible to study equilibrium behavior, entropy, metastability, and convergence toward a stable sentiment distribution. Magnetization describes the average direction of sentiment, polarization measures the share of agents holding directional views, and entropy measures the dispersion of sentiment across the three states. These quantities provide a direct link between individual belief updates and aggregate market behavior.

The simulation results show how the system responds to temporary information shocks. During the shock period, agents become more directional, neutrality declines, and local sentiment clusters may form. After the shock is removed, the system does not always return immediately to its earlier configuration. The speed of recovery depends on the strength of local interaction and the degree of randomness in belief updating. This provides a simple mechanism for delayed adjustment and persistent market narratives.

The empirical illustration shows how the model can be connected to financial text data. Sentiment labels are mapped into positive, neutral, and negative states, and the resulting proportions are used to compute magnetization, polarization, and entropy. Compared with binary benchmarks such as the Ising and voter models, the three-state formulation gives a clearer role to neutral agents and provides a more detailed description of sentiment adjustment around information-sensitive events. The empirical exercise remains illustrative, but it shows that the framework can be linked naturally to observed sentiment patterns.

Several extensions are possible. First, the one-dimensional lattice could be replaced by higher-dimensional lattices or by more realistic communication networks. Financial information spreads through social media, news platforms, institutional channels, and investor communities, rather than through a simple linear structure. Network-based extensions would therefore give a richer description of information transmission.

Second, future work could introduce heterogeneous agents. In the present model, agents follow the same updating rule. In real markets, investors differ in information access, risk tolerance, reaction speed, and willingness to remain neutral. Allowing for heterogeneous interaction strengths or different sensitivities to shocks would make the model closer to actual market behavior.

Third, the model could be calibrated more systematically using financial time series. Asset prices, trading volume, realized volatility, bid–ask spreads, and order-book variables could be used together with text sentiment to estimate parameters such as J, β, and the strength or duration of shocks. This would strengthen the empirical link between sentiment dynamics and market outcomes.

Fourth, the framework could be adapted for real-time monitoring. High-frequency text from financial news, social media, or analyst reports could be classified into the three sentiment states and then used to track changes in magnetization, polarization, and entropy. Such an approach may help detect early signs of polarization, delayed recovery, or instability after major announcements.

Finally, the model may be useful for scenario analysis. Different types of policy communication, macroeconomic shocks, or market stress events could be introduced through the external field. This would make it possible to study how the strength, direction, and duration of information shocks affect the recovery of market sentiment.

Overall, the three-state information lattice gas model provides a simple and tractable framework for studying sentiment-driven market behavior. Its main value lies in incorporating neutrality directly into the dynamics while preserving a clear mathematical structure. By combining local interaction, external shocks, and stochastic belief revision, the model offers a useful link between statistical physics, behavioral finance, and computational sentiment analysis.

7.1 Declarations

Conflicts of Interest

The authors declare that they have no conflicts of interest relevant to the content of this manuscript.

Bibliography

[1] E. Ising, *Beitrag zur Theorie des Ferromagnetismus*, Zeitschrift für Physik, (1925) 31, 253–258. https://doi.org/10.1007/BF02980577

[2] P. Clifford, and A. Sudbury, *A model for spatial conflict*, Biometrika, (1973) 60(3), 581–588. https://doi.org/10.1093/biomet/60.3.581

[3] H. Zeghdoudi, and H. Boutabia, *Computation for the canonical measures of a colored disordered lattice gas and spectral gap*, Journal of Mathematical Physics, (2009) 50, 103304.

[4] S. Bornholdt, *Expectation bubbles in a spin model of markets: Intermittency from frustration across scales*, International Journal of Modern Physics C, (2001) 12(5), 667–674. https://doi.org/10.1142/S0129183101001845

[5] C. Castellano, S. Fortunato, and V. Loreto, *Statistical physics of social dynamics*, Reviews of Modern Physics, (2009) 81(2), 591–646. https://doi.org/10.1103/RevModPhys.81.591

[6] Y. Xu, and S. B. Cohen, Stock movement prediction from tweets and historical prices. *In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (2018) pp. 1970–1979. Association for Computational Linguistics, https://doi.org/10.18653/v1/P18-1183

[7] M. Qezelbash, S. Tajdini, F. Jafari, M. Lotfi Ghahroud, and M. Farajnezhad, *An analysis of volatility and herd behavior among investors in the S&P500 stock market index, Bitcoin, and gold markets*, Journal of Mathematics and Modeling in Finance, (2023) 3(2), 77–92, https://doi.org/10.22054/jmmf.2024.75516.1103

How to Cite: Halim Zeghdoudi[1], Fatih Tank[2], *Modeling Financial Sentiment with a Three-State Lattice Gas: From Agent Interaction to Market-Level Shock Responses*, Journal of Mathematics and Modeling in Finance (JMMF), Vol. 6, No. 2, Pages:183–209, (2026).