Ising Spin Models for Simulation of Stock Markets

Ege Sönmez
Department of Physics, Bilkent University, Ankara, Turkey

November 2023

Abstract

This paper aims to enhance a pre-existing financial market model based on statistical physics principles, particularly spin models. The foundational model characterizes traders as either fundamentalists or interacting traders, each contributing to market dynamics and price movements. Our goal is to refine this model, improving its representation of market behaviors, price evolution, and trading volumes, allowing for a more accurate portrayal and understanding of real-world financial market phenomena such as bubbles and crashes.

1 Introduction

The intriguing world of finance often borrows concepts from the realm of physics to describe and predict market behavior. A prime example of this interdisciplinary approach is the application of the Ising model, a physical model of ferromagnetism, to financial markets. By drawing an analogy between the interactions of spins and the decisions of traders in a market, researchers have aimed to understand complex market dynamics.

In this project, we're looking the Bornholdt's spin model [1] that draws inspiration from the Ising model, portraying each trader as behaving similarly to a spin, influenced by immediate peers (nearest neighbours) and the overall market sentiment (magnetization). This interaction can result in unstable market activity, especially in periods of low activity, similar to a low-temperature state in the Ising model.

We aim to use the this model to offer straightforward explanations for why stock prices sometimes experience dramatic increases or sharp declines. The model incorporates two types of traders: those who make decisions based on the fundamental value of stocks, and those who follow market trends. It also includes a mechanism to match buy and sell orders, just like a real market does.

According to the model, stock prices come from their true value along with the collective sentiment of traders, and the volume of trading is indicative of this overall sentiment. Impressively, this model can replicate real-life market phenomena, such as periods of intense trading and significant price changes. [2]

We intend to enhance this model to better mirror the intricacies of real market behavior, which will help us understand the intricate movements of financial markets.

2 Presentation of the Model

In the original study that our work is based on, a special model was developed by T. Kaizoji, S. Bornholdt and Y. Fujiwara. [2] First, presentation of their work is necessary to understand our further developments on their model.

Consider a stock market where a significant volume of a stock is traded at price $p(t)$. In this market, two distinct groups of traders operate, each employing different trading strategies: *fundamentalists* and *interacting traders*. The model assumes that the number of fundamentalists, m, and the number of interacting traders, n, remain constant. The model aims to portray the dynamic fluctuations of stock prices within brief periods, typically a single day, capturing the nuanced decision-making processes of each category of traders with higher precision. [2]

2.1 Fundamentalists

Fundamentalists in this model are traders who rely significantly on the intrinsic value of the stocks. They possess a substantial understanding of the fundamental value of the stock, $p^*(t)$. Fundamentalists tend to purchase stocks when they perceive them to be in a discount zone relative to their intrinsic value and sell when they appear to be in a premium zone. Thus, the buying propensity of a fundamentalist is described as:

$$x^F(t) = am(\ln p^*(t) - \ln p(t)) \tag{1}$$

where a and m are predefined constants, symbolizing the reaction coefficient and the number of fundamentalists, respectively.

2.2 Interacting traders

During each trading cycle, an interacting trader has two options: to buy or sell a specified quantity of stock, denoted as b. Interacting traders in this model are identified by an integer value ranging between 1 and n. Each interacting trader, i, is characterized by a variable, s_i, which reflects their trading stance. Specifically, the variable s_i is assigned a value of $+1$ if the trader opts to buy, and -1 if the trader chooses to sell the stock.

To elaborate on the decision-making process of interacting traders, the model integrates dynamics from the spin model. The trading inclination, s_i, of each interacting trader is updated using a heat-bath dynamics method. Mathematically, this is represented as:

$$s_i(t+1) = +1 \text{ with } p = \frac{1}{1+\exp(-2\beta H_i(t))}, \quad (2)$$

$$s_i(t+1) = -1 \text{ with } 1-p, \quad (3)$$

where $H_i(t)$ signifies the local field of the spin model, directing the strategic decision of each trader.

Consider a scenario where the strategic choices of an interacting trader are influenced predominantly by two categories of information: local and global. Local information encompasses the behaviors and strategies adopted by nearest neighbour interacting traders, determining the immediate trading environment. In contrast, global information represents broader market trends, reflecting the aggregated stance of larger trader groups, whether they lean towards buying or selling.

This model assumes that each interacting trader's decisions are primarily influenced by the behaviors of their immediate neighbors and the overall market trends. One crucial aspect is the determination of whether a trader aligns with the majority or minority in terms of market stance, and understanding the magnitude of each group. This is quantified by the value of the magnetization $M(t)$, defined as:

$$M(t) = \frac{1}{n}\sum_{i=1}^{n} s_i(t). \quad (4)$$

Interacting traders aim to increase their profits through trading. They believe that aligning with the majority group is essential for earning profits. However, just being part of the majority isn't enough. The majority group needs to grow over each trading period. Traders in the majority group realize that as the size of the group increases, $|M(t)|$ gets larger, making it more challenging to maintain or increase the group's size further.

Hence, traders in the majority tend to switch to the minority group to avoid potential losses, like avoiding a significant market drop. In simple terms, as the majority group grows, traders in the group become more cautious. Conversely, traders in the minority group are more likely to take risks as the size of the majority group expands, switching their positions to join the majority and increase their profits.

To summarize, the magnitude of $|M(t)|$ significantly influences the propensity of traders, whether in the majority or minority group, to shift away from their current affiliations. The specific interactions involved are given by the local field $H_i(t)$:

$$H_i(t) = \sum_{j=1}^{m} J_{ij} s_j(t) - \alpha s_i(t) M(t), \quad (5)$$

where α is a positive constant that encompasses global coupling. The initial term is selected based on a localized Ising Hamiltonian that accounts for interactions between nearest neighbors, where $J_{ij} = J$ and $J_{ii} = 0$ for unrelated pairs.

The model further assumes that the excessive demand followed by interacting traders for the stock is approximated by

$$x^I(t) = bnM(t). \quad (6)$$

This equation is crucial in estimating the surplus demand manifested by interacting traders in the stock market.

2.3 Market Dynamics

This section focuses on explaining the decision-making processes of traders and determining the subsequent market price. The model introduces a market-clearing system in which a market maker facilitates trading, aligning the market price with the market-clearing values. This ensures that buy and sell orders are matched effectively.

The interplay between demand and supply is captured by the equation:

$$x^F(t) + x^I(t) = am[\ln p^*(t) - \ln p(t)] + bmM(t) = 0. \quad (7)$$

Consequently, the model computes the market price and trading volume as:

$$\ln p(t) = \ln p^*(t) + \lambda M(t), \quad \lambda = \frac{bn}{am} \quad (8)$$

$$V(t) = bn\frac{1+|M(t)|}{2}. \quad (9)$$

The model categorizes market situations based on the value of $M(t)$, determining whether the market is in a bullish ($|M(t)| > 0$) or bearish ($|M(t)| < 0$) state. A relative change in price, or the log-return, is then defined, enhancing our understanding of market price variations.

For a scenario focusing primarily on the participation of fundamentalists in trading ($p(t) = p^*(t)$), the model aligns with the Efficient Market Hypothesis [3], implying that prices would follow a random walk. This is supported by the assumption of a Gaussian process, which supports the random walk theory of the log-returns of asset prices, therefore offering insight into the fundamental price evolution.

2.4 Introducing different traders

Building upon the foundation laid by Kaizoji, Bornholdt, and Fujiwara, our study introduces a novel class of market participants: the *high beta traders* and the *normal interacting traders*. The distinction between these two types of traders lays in their respective levels of market sensitivity and information processing factors encapsulated in the parameter β, which denotes the inverse temperature in statistical mechanics.

2.4.1 High Beta Traders

High beta traders are defined by their heightened sensitivity to the market, modeled by an extremely high beta value, approaching the limit of market responsiveness.

For high values of β, we consider the limit as $\beta \to \infty$, simplifying the probability p to $\frac{1}{2}$, assuming that $H_i(t)$ is finite. Thus, we have:

$$p = \frac{1}{1 + \exp(-2\beta h_i(t))} \approx \frac{1}{2}, \tag{10}$$

$$1 - p = \frac{1}{2}. \tag{11}$$

Consequently, the updated states of the spins for high beta traders become:

$$s_i(t+1) = +1 \text{ with } p = \frac{1}{2}, \tag{12}$$

$$s_i(t+1) = -1 \text{ with } 1 - p = \frac{1}{2}. \tag{13}$$

These traders represent a class of market participants with maximum information entropy, indicative of an essentially random behavior. Their decisions are not swayed by the typical influence of market trends or the actions of their peers. Consequently, their trading behavior is highly unpredictable and serves as a proxy for the most extreme speculative forces in the market.

2.4.2 Normal Interacting Traders

In contrast, normal interacting traders possess a moderate level of responsiveness, defined by a standard beta value ($\beta = 2$). These traders are influenced by both local and global market information, making decisions that balance individual insights with broader market sentiments. Their behavior can be considered more representative of the average market participant, who is neither completely random nor perfectly rational.

3 Market Initialization and Simulation Dynamics

The simulation begins with an equitable distribution of high beta traders and normal interacting traders, totaling n, throughout a two-dimensional, $10 * 10$ lattice. Initially, the traders are randomly distributed, reflecting a market free of spatial or informational segmentation.

To assess the impact of spatial and informational homogeneity on market dynamics, we support the random distribution scenario with a second initialization configuration: a market divided into two distinct zones. On the left side of the lattice reside all normal interacting traders, while the right side is solely occupied by high beta traders. This segregation provides an investigation into the interaction and diffusion of trading strategies across a market that is initially polarized in terms of trader behavior and market outlook.

3.1 Simulation Objectives and Expected Outcomes

The core objective of our simulation is to analyze how the presence and distribution of different types of traders affect market dynamics, price formation, and volatility. We aim to understand whether the introduction of high beta traders worsen market fluctuations, leading to greater extremes in price movements, or whether their randomness simply injects noise into the market without significantly altering its overall behavior.

Moreover, by starting the simulation with segregated groups of traders, we intend to explore how information and trading strategies percolate through a divided market. We hypothesize that the interface between the two trader groups will be a critical zone, giving rise to complex trading dynamics as the two contrasting strategies (random versus responsive) consolidate.

The outcomes of this simulation have the potential to provide deeper insights into the consequences of trader heterogeneity on financial markets. Specifically, we aim to contribute to the existing knowledge on how extreme speculative behavior, modeled by high beta traders, interacts with and influences the more rational strategies employed by normal interacting traders.

3.2 Simulation Methodology

The simulation will be executed over a series of time steps, with each step allowing traders to update their positions based on their individual beta values and the computed local field (Eq. 5), as per the heat-bath dynamics method. For our simulation, parameters will be adjusted as following: $J = 1$ and $\alpha = 2.8$. These parameters play a pivotal role in shaping market dynamics because they directly influence the local field, as described by Equation 5. The parameter J controls the pace at which the market responds to the interactions between neighboring traders. When J is large, the market adapts slowly, and this tends to produce more distinct, step-like changes in market behavior. Conversely, a smaller J leads to more gradual market responses, resulting in smoother, more continuous price movements. The parameter α, on the other hand, controls the strength of the market's overall tendency to move from premium and discount zones towards an equilibrium point, under the guiding influence of magnetization. The process for updating the trading stance (s_i) will remain consistent with the model of Kaizoji et al. for normal beta valued interacting traders, therefore maintaining the structure of the underlying theoretical framework while introducing the new element of trader diversity.

Market prices will evolve according to the established equations, incorporating the distinct demand functions of the fundamentalists and the speculative pressures exerted by the interacting traders. The resulting time series of prices, log-returns, and other relevant financial indicators will be analyzed to distinguish the effects of the introduced trader heterogeneity.

This methodology will enable us to compare the resultant market behavior under two initial conditions: one reflecting a random mixing of trader types and the other representing a market mixed into two ideologically disparate halves. Through this comparative analysis, we aim to shed light on the complex dynamics that govern financial markets and the role that trader diversity plays in shaping them.

4 Results

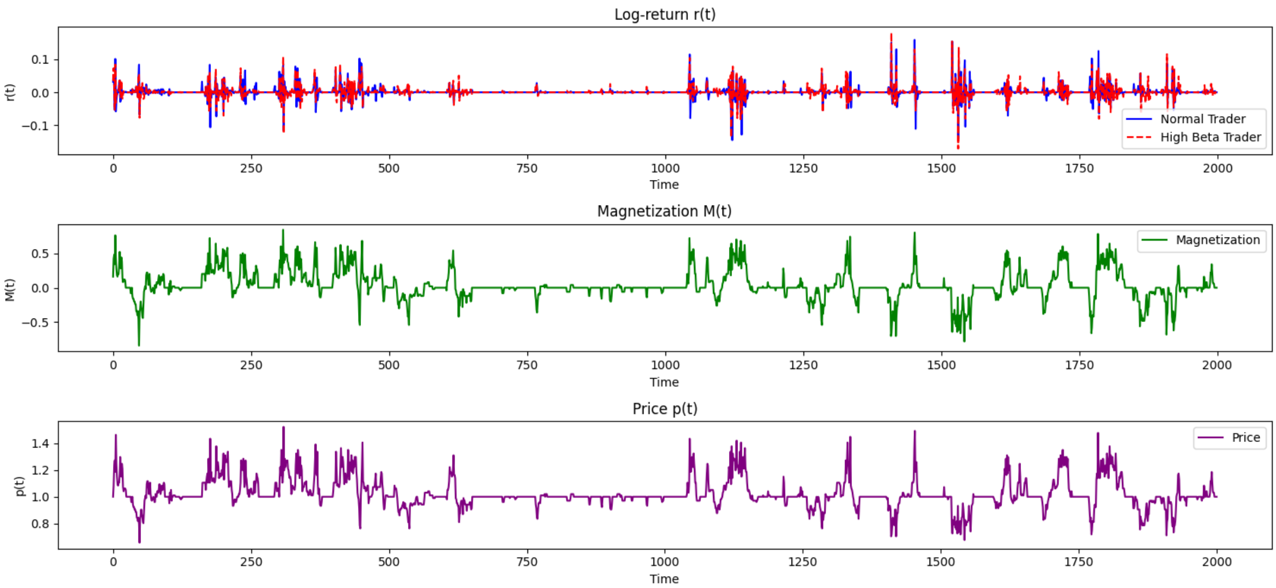

Figure 1: Market initialization with random-shuffled traders.

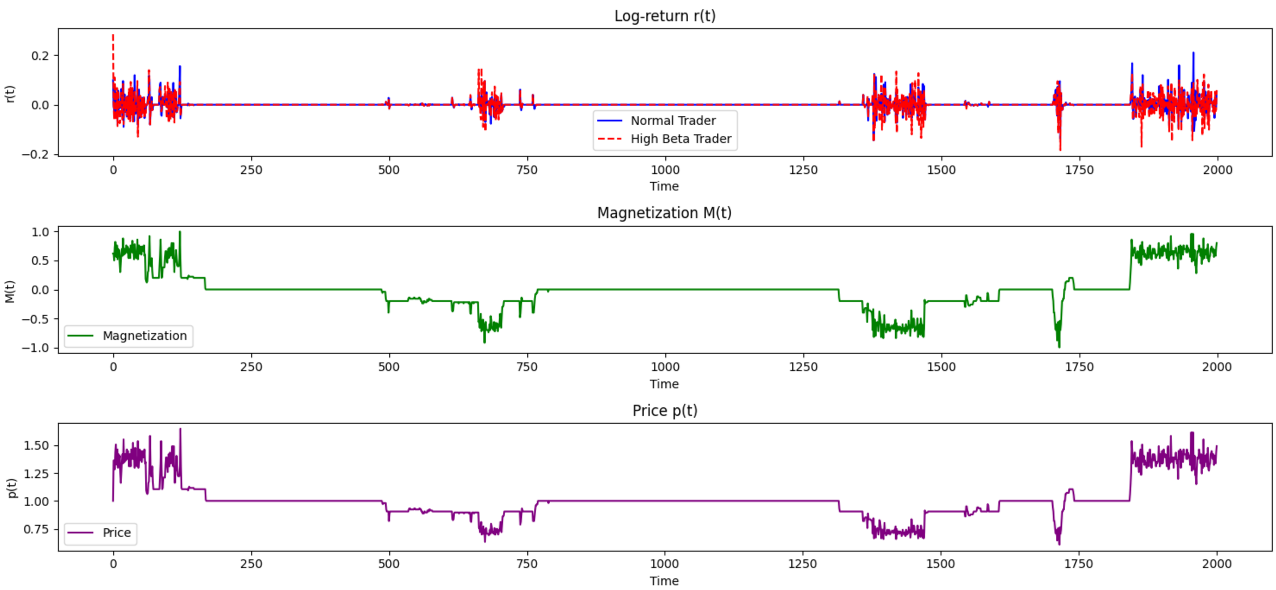

Figure 2: Market initialization with distinctly distributed traders.

The comparative analysis between the random shuffled case (Fig. 1) and the distinctly distributed case (Fig. 2) gives a great understanding about the influence of trader types on market dynamics. The findings highlight the impact of trader interactions and the consequences of collective behavior on financial indicators such as log-returns and market magnetization.

Average Log-returns		
	Random-Shuffled	Distinctly Distributed
Normal Traders	0.001334	0.002618
High Beta Traders	0.001522	0.000984

In the random shuffled case, the market dynamics are characterized by a higher degree of fluctuation, which is attributed to the presence of high beta traders randomly spread among normal traders. The extreme trading behaviors of the high beta traders seem to exert a disruptive influence on the normal traders, leading to a decrease in their log-returns. This scenario underlines the significance of one's financial neighbors and suggests that the proximity to less rational, high-risk takers can negatively affect the performance of more moderate, rational market participants. Additionally, the random distribution of high beta traders across the market implies a higher temperature state of the system, indicative of more volatile market behavior.

Conversely, in the distinctly distributed case, the market experiences less severe fluctuations. The normal traders, segregated from the high beta traders, demonstrate higher returns, unaffected by the extreme behaviour of their high beta counterparts. This observation suggests that a community of like-minded traders can achieve more stable and potentially more profitable outcomes when insulated from extreme traders. High beta traders, now in a homogeneous group, exhibit lower returns compared to when they are mixed with normal traders. This could reflect the diminishing returns of high-risk strategies in an echo chamber-like environment without the balance provided by more risk-averting trading strategies.

The most significant market activity occurs at the boundary between the two distinct groups. It is at this boundary where the different strategies meet and interact, causing potential fluctuations. Unlike the constant market-wide fluctuations seen in the random shuffled case, the distinctly distributed case presents fluctuations that are more localized to the regions of ideological transition between the two trader types.

The outcomes of the study make it clear that the market behaves very differently depending on the mix and placement of traders with various strategies. When traders with all sorts of strategies are thrown together, the market gets pretty shaky, which can mess with everyone's profits. But when traders are sorted into their own zones, things tend to be less wild, and regular traders often do better because they're not caught up with the high-risk takers. This shows that having a variety of trader behaviors and where they are in the market can really change how stable the market is and how much money people can make. It's a heads-up that how traders are arranged and who's near who can really shape what happens in the market, not just for one kind of trader but for everyone involved.

5 Correlation with Real-World Markets

In the real-world financial markets, the dynamics observed in simulations often correlate with complex interactions among different market participants. The impact of financial contagion and network theory suggests that the structure of connections among agent significantly influences systemic risk. For instance, Allen and Gale's research highlights how disturbances can spread through a financial system, affecting various entities differently based on their interconnectedness. [4]

Furthermore, the influence of high-frequency trading, parallel to the simulated high-beta traders, is well-documented, with Brogaard, Hendershott, and Riordan illustrating how such activities can lead to increased price discovery and market volatility. [5] These traders, operating at high speeds, can worsen market movements, similar to the disruptive influence noted between high beta and normal traders in the simulation.

Additionally, the clustering of trading behaviors or strategies can have a profound impact on market liquidity and price discovery. J.P. Morgan Asset Management has noted that the liquidity cost tradeoff is an essential consideration where the concentration of trading actions can affect market dynamics, echoing the simulation results where the separation of traders into distinct clusters led to varying outcomes. [6]

These real-world studies support the idea that the arrangement and behavior of traders can indeed shape market dynamics significantly. While the simulation models provide a ideal representation, empirical evidence suggests that such models can capture essential aspects of market behavior, though it's critical to recognize that real markets also contend with factors beyond the scope of these models.

6 Future Work

For future work, we plan to look into the critical moments when market volatility—reflected in sharp movements in magnetization—peaks or plummets. By zooming in on these time-steps, we aim to unravel the intrinsic dynamics of traders' states and decisions during periods of heightened activity. Understanding the triggers and responses of these fluctuations could offer insights into the collective market behavior and individual trading strategies.

Furthermore, we will explore the concept of adaptive trader strategies. Recognizing that markets often exhibit periodic patterns for realistic cases, there may be potential to come up with trading philosophies that adapt to these periodicities. The goal of this adaptive approach would be to finesse the balance between risk and return, optimizing traders' positions in accordance with evolving market conditions and emergent trends. By doing so, traders could potentially enhance their performance, achieving greater returns while mitigating risks associated with volatile market phases and other traders' behaviour.

References

[1] T. Yamano, "Bornholdt's spin model of a market dynamics in high dimensions," *International Journal of Modern Physics C*, vol. 13, no. 01, pp. 89–96, 2002.

[2] T. Kaizoji, S. Bornholdt, and Y. Fujiwara, "Dynamics of price and trading volume in a spin model of stock markets with heterogeneous agents," *Physica A: Statistical Mechanics and its Applications*, vol. 316, no. 1, pp. 441–452, 2002.

[3] W. Zhou and D. Sornette, "Self-organizing ising model of financial markets," *European Physical Journal B*, vol. 55, pp. 175–181, 2007. [Online]. Available: https://doi.org/10.1140/epjb/e2006-00391-6

[4] F. Allen and D. Gale, "Financial contagion," *Journal of Political Economy*, vol. 108, no. 1, pp. 1–33, 2000.

[5] J. Brogaard, T. Hendershott, and R. Riordan, "High-frequency trading and price discovery," *Review of Financial Studies*, vol. 27, no. 8, pp. 2267–2306, 2014.

[6] J.P. Morgan Asset Management, "The liquidity cost tradeoff," Global Liquidity Series, Tech. Rep., 2012.

7 Appendix

Python code for the simulation.

```
    import numpy as np
import matplotlib.pyplot as plt

# Constants and Parameters
M, N = 10, 10
T = 0.5
BETA = 1 / T
HIGH_BETA = 1e5
J = 1
ALPHA = 2.8
P_STAR = 1
a = 1
m = 100
b = 1
n = M * N

class Trader:
    def __init__(self, beta):
        self.s = 1 if np.random.rand() < 0.5 else -1
        self.beta = beta

    def update_spin(self, h):
        bh = self.beta * h
        if bh > 100:
            p = 1
        elif bh < -100:
            p = 0
        else:
            p = 1 / (1 + np.exp(-2 * bh))
        self.s = 1 if np.random.rand() < p else -1

class HighBetaTrader(Trader):
    def __init__(self):
        super().__init__(HIGH_BETA)

class NormalTrader(Trader):
    def __init__(self):
        super().__init__(BETA)

def initialize_lattice():
    traders = [NormalTrader() for _ in range(n // 2)] + [HighBetaTrader() for _ in range(n //
2)]
    np.random.shuffle(traders)
    return [traders[i*N:i*N+N] for i in range(M)]


#def initialize_lattice():
#    half = N // 2  # half the lattice width
#    left_traders = [NormalTrader() for _ in range(n // 2)]
#    right_traders = [HighBetaTrader() for _ in range(n // 2)]

#    trader_lattice = []

#    for _ in range(M):
#        row = left_traders[:half] + right_traders[:half]
#        trader_lattice.append(row)
#        left_traders = left_traders[half:]
```

```
#        right_traders = right_traders[half:]

#    return trader_lattice


def compute_local_field(trader_lattice, i, j):
    h = 0
    M_t = np.mean([t.s for row in trader_lattice for t in row])
    for x, y in [(i-1, j), (i+1, j), (i, j-1), (i, j+1)]:
        x = x % M
        y = y % N
        h += J * trader_lattice[x][y].s
    h -= ALPHA * trader_lattice[i][j].s * abs(M_t)
    return h

def fundamentalist_demand(p_t):
    return a * m * (np.log(P_STAR) - np.log(p_t))

def simulate(steps=2000):
    trader_lattice = initialize_lattice()
    magnetizations = []
    prices = [P_STAR]
    trader_log_returns, high_beta_log_returns = [], []

    for step in range(steps):
        # Initialize current returns for the time step
        current_trader_returns, current_high_beta_returns = [], []

        for i in range(M):
            for j in range(N):
                old_spin = trader_lattice[i][j].s
                h = compute_local_field(trader_lattice, i, j)
                trader_lattice[i][j].update_spin(h)

        # Calculate the market return after all traders potentially change spin
        if step > 0:  # Ensure there is a previous price to compare
            r_t = np.log(prices[-1]) - np.log(prices[-2])

            # Now, we iterate over all traders to calculate their individual returns
            for i in range(M):
                for j in range(N):
                    if isinstance(trader_lattice[i][j], NormalTrader):
                        current_trader_returns.append(r_t if trader_lattice[i][j].s == 1 else
-r_t)
                    elif isinstance(trader_lattice[i][j], HighBetaTrader):
                        current_high_beta_returns.append(r_t if trader_lattice[i][j].s == 1 else
-r_t)

        M_t = np.mean([t.s for row in trader_lattice for t in row])
        magnetizations.append(M_t)

        xF = fundamentalist_demand(prices[-1])
        xI = b * 0.5 * n * M_t
        new_price = prices[-1] * np.exp((xF + xI) / (a * m))
        prices.append(new_price)

        # Record average return for each type of trader during this time step
        if step > 0:  # No returns to record for the first step since there is no previous price
            trader_log_returns.append(np.mean(current_trader_returns))
            high_beta_log_returns.append(np.mean(current_high_beta_returns))
```

```
    return trader_log_returns, magnetizations, prices, high_beta_log_returns


trader_log_returns, magnetizations, prices, high_beta_log_returns = simulate()

plt.figure(figsize=(15, 7))

# Log Returns
plt.subplot(3, 1, 1)
plt.plot(trader_log_returns, label='Normal Trader', color='blue')
plt.plot(high_beta_log_returns, label='High Beta Trader', color='red', linestyle='--')
plt.title('Log-return r(t)')
plt.xlabel('Time')
plt.ylabel('r(t)')
plt.legend()

# Magnetization
plt.subplot(3, 1, 2)
plt.plot(magnetizations, label='Magnetization', color='green')
plt.title('Magnetization M(t)')
plt.xlabel('Time')
plt.ylabel('M(t)')
plt.legend()

# Price
plt.subplot(3, 1, 3)
plt.plot(prices, label='Price', color='purple')
plt.title('Price p(t)')
plt.xlabel('Time')
plt.ylabel('p(t)')
plt.legend()

plt.tight_layout()
plt.show()

# Statistics
print("Normal Trader Stats:")
print(f"Average Log-return: {np.mean(trader_log_returns)}")
print(f"Average Price: {np.mean(prices)}")

print("\nHigh Beta Trader Stats:")
print(f"Average Log-return: {np.mean(high_beta_log_returns)}")
print(f"Average Price: {np.mean(prices)}")
```