complexities

MDPI

Article

Reproducing Stylized Facts in Artificial Stock Markets with Price-Data-Trained Neural Agents

Qi Zhang and **Yu Chen** *

Department of Human and Engineered Environment, Graduate School of Frontier Sciences, The University of Tokyo, Kashiwa-shi 277-8563, Japan
* Correspondence: chen@edu.k.u-tokyo.ac.jp

Abstract

Agent-based models of financial markets often rely on a small set of hand-crafted trading rules, making it difficult to relate model heterogeneity to information that is observable in market data. We take a different standpoint and treat the design of heterogeneity as a representation problem under limited observations. In our framework, each agent's decision rule is implemented as a neural-network mapping from recent price histories to order decisions, trained on historical index or stock price series. To describe and manipulate heterogeneity without pre-assigning mechanism labels, we introduce Fit Quality (FQ), an ex post effect-defined index summarizing how strongly each learned rule fits the price patterns it was trained on, and we use FQ solely as a coordinate for organizing agent populations and constructing controlled changes in agent composition, rather than as a measure of forecasting skill or economic performance. Using this representation, we examine whether simulations can reproduce several stylized features of return series. We also perform simple ablation experiments to assess how far the observed properties depend on the data-trained decision rules rather than on the market mechanism alone. Taken together, the framework is intended as a step toward more data-linked, representation-conscious agent-based models, in which alternative ways of organizing heterogeneity can be compared within a common market environment.

Keywords: agent-based modeling; data-driven market simulation; neural agents; maket microstructure; market stylized facts

Academic Editor: José F.F. Mendes

Received: 7 January 2026
Revised: 7 February 2026
Accepted: 11 February 2026
Published: 13 February 2026

Correction Statement: This article has been republished with a minor change. The change does not affect the scientific content of the article and further details are available within the backmatter of the website version of this article.

1. Introduction

Modeling financial markets from the bottom up remains a challenging problem in complex systems science. The central difficulty lies in capturing the heterogeneous, adaptive decision processes of market participants—processes that constitute a "black box" of evolving heuristics and strategies [1,2].

Agent-based modeling (ABM) has long been a widely used approach for studying such systems. Traditional financial ABMs categorize traders into archetypes—such as fundamentalists, chartists, and noise traders—each governed by predefined rules [2,3]. While this rule-based design has yielded valuable qualitative insights, it also imposes structural rigidity. Because behavioral logic is manually specified, model outcomes can be sensitive to subjective assumptions and parameter choices, and calibration can become challenging when multiple free parameters interact [4–6].

Within this tradition, a well-established branch—often built on Kirman-type herding or noisy-voter dynamics—develops parsimonious, mechanism-driven models and

 https://doi.org/10.3390/complexities2010004

validates them against empirical statistical regularities [7–10]. In parallel, a related methodological literature develops more principled estimation/calibration procedures for ABMs (e.g., likelihood-free or SMC-based approaches), improving reproducibility and parameter interpretability [11]. However, the representation of heterogeneity remains mechanism-first: each trader archetype is still realized by an explicit, hand-crafted decision rule, which compresses a potentially wide spectrum of individual decision logics into a small set of homogeneous templates. Moreover, the content of these rules can be contestable in realistic settings—for example, how "fundamental value" should be defined or estimated, and why different traders would share the same valuation and updating logic.

Beyond these mechanism-specified ABM designs, recent work has explored integrating machine learning into agent-based modeling (ABM) to mitigate some of the limitations of hand-crafted rule design [12,13]. A recent line of work explores representing agents' decision processes with neural networks, so that the mapping from observations to actions can be learned rather than manually specified. For example, Jäger proposes replacing explicit behavioral rules with neural decision functions in conceptual ABM settings [14,15]. Related work further examines whether neural agents can reproduce patterns observed in simplified social-interaction tasks and offers broader frameworks for integrating machine-learning components into ABM pipelines [12,16]. However, these studies are often demonstrated in toy models and outside the financial domain, and how to implement such "learned decision rules" in financial-market ABMs—while keeping the market mechanism coherent and evaluating canonical stylized facts—remains comparatively underexplored.

In the financial ABM literature, by contrast, machine learning has more commonly been used as an auxiliary tool—for calibration, inference, surrogate modeling, or likelihood-free estimation—while the underlying agents remain rule-based and mechanism-specified [17–20]. Meanwhile, reinforcement learning (RL) is widely used to learn trading policies through repeated interaction with simulators [21,22]. In many RL-based studies, the primary goal is to train an agent that achieves strong decision performance (e.g., profitability or execution quality) under a given environment, rather than reproduce the emergent stylized facts. Moreover, because these policies are often trained in synthetic environments, how well a strategy learned in simulation transfers to real markets can depend sensitively on simulator fidelity, reward design, and constraints, making real-world validity an open question in many settings [23].

A shared constraint across the above directions is data. In equity markets, information that would directly characterize agent behavior—trader-level intentions, belief updates, decision rules, and the micro-level context behind order placement—is typically unavailable in academic datasets. Instead, the commonly available evidence is aggregated market outcomes (prices, returns, volumes), i.e., many-to-one projections of heterogeneous decision-making. This leaves an identification challenge: in the absence of micro-level behavioral records, how can limited aggregated data be used to build and evaluate agent-based representations of trading behavior?

Motivated by this data-to-behavior identification gap, we adopt a different standpoint. In real markets, trading behavior is often diverse, mixed, and time-varying. For an external observer, it is rarely possible to reliably identify which cues a trader relies on—fundamentals, a particular technical indicator, a combination of signals, or other idiosyncratic information. This makes an ex ante taxonomy based on a small set of mechanism labels inherently contestable: even within a label, multiple definitions are possible (e.g., what constitutes "fundamental value," or which indicator represents "momentum").

Nevertheless, trading decisions typically share a weaker and more general structure: they reflect some processing of past information into expectations about the future, which then drives action. Accordingly, instead of prespecifying the analytical form of trading heuristics, we represent each agent's decision rule as an individual-specific, high-dimensional mapping from observed price histories to actions, parameterized by a neural network.

To organize the resulting diversity, we use Fit Quality (FQ) as an ex post descriptive coordinate. FQ summarizes how strongly a learned rule is anchored to historical patterns in the data used for learning and serves as an analysis index for displaying heterogeneity and constructing controlled population-composition contrasts. Importantly, FQ is used only for organizing heterogeneous agents; it is not interpreted as a superior forecasting skill or a normative measure of real traders' performance. In other words, we shift from a mechanism-defined taxonomy ("what rule an agent follows") to an effect-defined taxonomy ("what statistical behavior a trained rule induces"), operationalized by FQ.

This stance shifts the modeling burden from selecting a small set of explicit behavioral equations to specifying a representation class for decision rules and a training protocol. We do not claim that this removes all modeling choices; rather, it reduces reliance on a priori assumptions about the analytical form of trading heuristics. Instead of tuning parameters that directly control the intensity of hypothesized mechanisms (e.g., switching rates, herding strength), we induce heterogeneity through training outcomes and evaluate the resulting market-level statistics under a fixed double-auction order-book matching engine. Accordingly, our novelty lies not in introducing neural networks per se, but in a concrete "Replace-and-Simulate" market-construction protocol plus an effect-defined heterogeneity axis (FQ), and in evaluating whether the resulting market reproduces canonical stylized facts under controlled ablations and compositions.

At the level of reduced-form time-series description, volatility clustering is commonly modeled by ARCH-family specifications, beginning with Engle's ARCH (1982) and Bollerslev's GARCH (1986), where the conditional variance evolves via low-dimensional recursions on past squared innovations and lagged conditional variances [24,25]. Such models are designed primarily as statistical descriptions for inference and forecasting of volatility, rather than as structural representations of interacting heterogeneous traders. In this sense, ARCH/GARCH provide complementary reduced-form references for characterizing volatility regularities in output return series, whereas our focus is on a generative market simulator in which these regularities emerge from, and can be studied through, the interaction and composition of heterogeneously learned decision rules.

The remainder of this paper is structured as follows. Section 2 details our methodology for agent design and market simulation. Section 3 presents empirical results, from agent-level validation to emergent market-level dynamics. Finally, Section 4 concludes the paper and outlines future research directions.

2. Method

The proposed framework combines (i) a neural-network-based agent architecture whose decision rules are trained from historical price data, and (ii) a fixed order-matching engine operating under a limit-order-book mechanism. Figure 1 illustrates the overall workflow of the proposed framework. Unlike archetype-based ABMs that define agents by a small number of explicit trading rules, our framework instantiates agents as trained decision mappings and characterizes heterogeneity via training-induced effects (FQ). The market mechanism is held fixed: trained agents submit orders, transactions occur through the matching engine, and a step-level price is defined from executed trades to form the simulated price series.

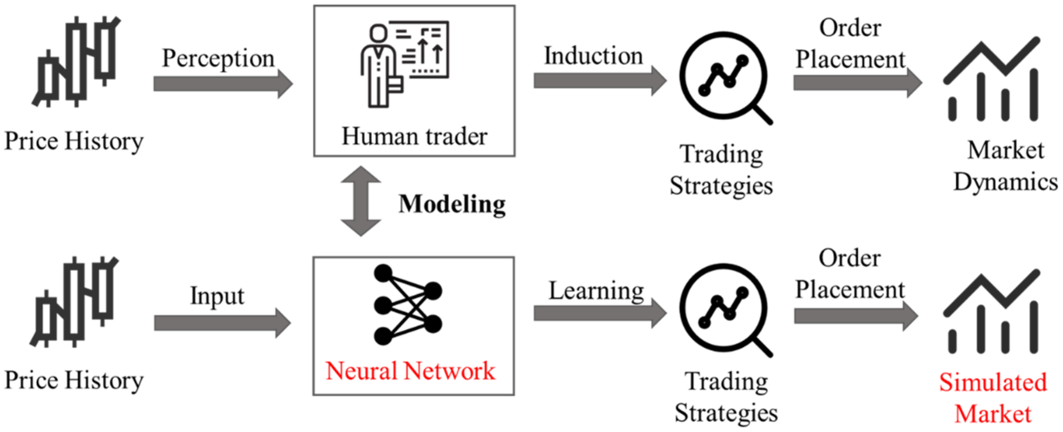

Figure 1. Schematic of the proposed data-trained agent-based market construction. The neural network (**bottom**) learns a mapping from recent price history to trading decisions, serving as an operational substitute for the decision component of a human trader (**top**). The resulting decisions are translated into orders and executed under a fixed market mechanism to generate simulated market dynamics.

2.1. Agent Architecture

Traditional ABMs manually specify traders' decision rules. In contrast, we replace this rule-design process with a trainable neural network that learns a decision rule from price histories under the specified input representation and training protocol. Formally, each agent observes a window of price history x_t of N past returns r_t,

$$\mathrm{x}_t = [r_{t-N}, r_{t-N+1}, \ldots, r_{t-1}] \quad (1)$$

and outputs a one-step-ahead prediction $\hat{r}_t = f_\theta(\mathrm{x_t})$, where f_θ is the neural network with learnable parameters θ. In this study, all agents share the same network architecture: a 4-layer feedforward MLP network (layer widths: 2048, 1024, 512, 256) with ReLU activations. We adopt this architecture as a simple, high-capacity baseline and do not aim to optimize model architecture; the specification is reported for reproducibility.

2.2. Endogenous Generation of Agent Heterogeneity

Having established a uniform neural architecture for all agents in Section 2.1, we now introduce heterogeneity not through structural variations, but through differentiating the "experiences" of each agent. Specifically, we fix the functional class (the MLP) and induce behavioral diversity solely through the distinct training datasets and the resulting rigor of historical fitting, quantified herein as Fit Quality (FQ).

2.2.1. Creating Agent "Personalities" via "Fit Quality"

Traditionally, fitting of a neural network evaluates how well the model generalizes to data similar to its training samples. In our framework, we reinterpret this fitting as an agent's behavioral characteristic. Because our goal is to simulate the financial markets, not predict the price, we discard the conventional validation step in training: the entire historical dataset representing the agent's lived experience is used to train the agent, and there is no test dataset. To control the depth of an agent's memorization of that experience, we define a continuous metric called Fit Quality (FQ), which serves as a proxy for its "personality" or trading style.

In practice, an agent's FQ is measured with its accuracy in predicting the direction of price movements on its own training set. For agent i with predicted return $\hat{r}_{i,t}$ and true return $r_{i,t}$ at time t, the sign-based accuracy is defined as

https://doi.org/10.3390/complexities2010004

$$Acc_i = \frac{1}{T}\sum_{t=1}^{T} \mathbb{I}(\text{sign}(\hat{r}_{i,t}) = \text{sign}(r_{i,t})) \quad (2)$$

where T is the number of historical observations used for evaluation, and I is the indicator function.

We emphasize that this FQ does not reflect an agent's predictive power on unseen future data. Rather, it is purely a measure of in-sample fit, quantifying how closely an agent has memorized the patterns within its specific historical training set. Consequently, a high-fit agent is not inherently more "skilled" or profitable than a low-fit one. The spectrum of FQ is anchored by illustrative archetypes at its extremes, rather than being limited to discrete categories.

- On one end, a low FQ (e.g., achieving 60% accuracy on its training data) represents a "Novice Trader." This agent has only grasped broad statistical tendencies, leading to more cautious and agnostic behavior.
- On the other hand, a high FQ (e.g., 90%+ accuracy) corresponds to an "Experienced Trader" who has effectively overfitted its training data. This results in a rigid and complex internal logic, mirroring a trader who is highly confident in familiar patterns but may react inflexibly to new market situations.

By treating FQ as a continuous, controllable parameter, we can generate an entire population of agents with finely graded differences in their trading styles.

2.2.2. Creating Agent "Expertise" via Data Composition

Because neural networks develop internal logic based on the patterns inherent within their specific training data, a powerful avenue for cultivating heterogeneity lies in varying the data provided to each agent. Our second mechanism, therefore, involves creating agents with different "domain knowledge" or "expertise" by training them on unique, targeted subsets of the available historical data. Agents are trained on intraday JPX price series over 2019–2023 [26]. For each training run, we first select a target industry (the industry categorization is based on the official Nikkei 225 component classifications [27]), then randomly sample stocks within that industry, and extract a continuous 10-trading-day segment. Returns are computed as log differences of consecutive 1 min prices.

This variation in the composition of the training data serves to endow each agent with a unique informational background or "upbringing." Consequently, agents learn different statistical patterns, develop distinct market perspectives, and form specific biases. These differences in learned internal logic naturally translate into diverse behavioral responses and decision-making patterns when the agents interact within the simulated market environment.

2.2.3. Formal Training Process

The mechanisms for cultivating "personality" (via Fit Quality) and "expertise" (via Data Composition) are grounded in a formal training process, specifically designed to achieve a predetermined level of memorization depth for each agent. Agents are trained offline on their assigned historical market datasets (as described in Section 2.2.2), and their parameters are fixed during the subsequent market simulation. We emphasize that the agents learn directly from the raw time-series data themselves, without any pre-assigned labels regarding the "trader type" (e.g., fundamentalist, chartist, etc.) that might have generated those patterns. The agents simply learn to map past returns to future returns based on the patterns present in their specific data sample.

Let $(x_t, y_t)_{t=1}^{T}$ denote an agent's training examples, where x_t is the input vector of past returns up to time t as defined in Equation (1), and y_t is the actual next-period return (the target). We train the agent's neural network by minimizing the mean squared error (MSE) between predictions and actual outcomes:

$$L(\theta) = \frac{1}{T}\sum\nolimits_{t=1}^{T}(\hat{r}_t - y_t)^2 \quad (3)$$

where $\hat{r}_t = f_\theta(x_t)$ is the network's prediction. We use log-returns as the prediction target for simplicity, but the framework can accommodate other targets by adjusting the loss function.

For each agent, a target FQ (a targeted in-sample accuracy, as defined in Equation (2)) is specified beforehand. The training process continues iteratively, updating the parameters θ to minimize $L(\theta)$, until the agent's performance on its training data reaches this predefined target FQ. This "train-to-target" methodology allows us to control the degree of memorization for each agent, consistent with the philosophy outlined in Section 2.2.1.

2.3. Prediction-to-Order Mapping

Once an agent has generated its prediction of the one-step-ahead return $\hat{r}_{i,t}$, this internal forecast must be translated into a concrete market order. Since we do not employ the microstructural data at this stage, we designed this component to generate a rich diversity of trading behaviors from a simple predictive signal. Each order $o_{i,t}$ is defined by three key attributes: limit price, direction and size.

2.3.1. Order Price

Agent i's limit price, $L_{i,t}$, is directly informed by its prediction of return. Specifically, it is sampled from a uniform distribution between the current market price, P_t, and the agent's forecasted next price, $P_t(1 + \hat{r}_{i,t})$, i.e.,

$$L_{i,t} \sim \begin{cases} \mathrm{Uniform}(P_t\ ,\ P_t(1+\hat{r}_{i,t})), & \hat{r}_{i,t} > 0, \\ \mathrm{Uniform}(P_t(1+\hat{r}_{i,t}),\ P_t), & \hat{r}_{i,t} < 0, \end{cases} \quad (4)$$

This formulation implies that the magnitude of the predicted return, $|\hat{r}_{i,t}|$, reflects the agent's "confidence" in its strategy. Indeed, a more confident prediction leads to a wider price range and thus a potentially more aggressive order placement.

2.3.2. Order Direction

In financial markets, motives of a trader may include trend following, profit-taking, and contrarian trading. To reflect this diversity, we decouple the order direction of agent i, $d_{i,t}$, from the sign of the agent's predicted return. At each simulation step, the agent is randomly assigned to submit either a buy or a sell order. Formally, this random assignment is achieved by sampling the order direction from a Bernoulli distribution with equal probability:

$$d_{i,t} \sim \mathrm{Bernoulli}(0.5),\ d_{i,t} \in \{\mathrm{buy}, \mathrm{sell}\} \quad (5)$$

This simple mechanism, when combined with the limit price logic above, generates four distinct and commonly observed behaviors, as shown in Table 1:

Table 1. Mapping predictions to strategic behaviors.

Price Prediction	Random Action	Resulting Behavior	Strategic Interpretation
RISE $\hat{r}_{i,t} > 0$	BUY	Places a bid near the current price P_t	Speculative Buy: "Buy now before the price goes up."
RISE $\hat{r}_{i,t} > 0$	SELL	Places an ask above P_t (near $P_t(1 + \hat{r}_{i,t})$)	Profit-Taking Sell: "Sell high if the rally happens."

Table 1. *Cont.*

Price Prediction	Random Action	Resulting Behavior	Strategic Interpretation
FALL $\hat{r}_{i,t} < 0$	BUY	Places a bid below P_t (near $P_t(1 + \hat{r}_{i,t})$)	Dip Buying: "Buy low after the price drops."
FALL $\hat{r}_{i,t} < 0$	SELL	Places an ask near the current price P_t	Speculative Sell: "Sell now before the price drops further."

This structure allows agents with identical expectations of price to interact with the market in different ways, capturing a spectrum of real-world strategies without creating complex rules for each.

2.3.3. Order Size and Submission Frequency

All submitted orders have a constant unit size $q = 1$. At each simulation step, every agent submits exactly one order based on its current prediction. Agents are processed in a uniformly random order without replacement.

Thus, the complete order is represented as the tuple

$$o_{i,t} = (d_{i,t},\ L_{i,t},\ q_{i,t}) \tag{6}$$

where each component is specified by a simple rule set to generate heterogeneous order submissions.

2.4. Market Simulation and Matching Engine

Having defined the architecture and generation process for heterogeneous agents, we now specify the environment in which they interact. We implement a discrete-time continuous double-auction (CDA) mechanism for the matching of orders [28]. At each time step t, all submitted orders $o_{i,t} = (d_{i,t},\ L_{i,t},\ q_{i,t})$ are aggregated into a central order book. Order matching follows price–time priority: highest bid meets lowest ask, and a trade occurs if bid is higher than ask. Let B_t and A_t denote the set of buy and sell orders, respectively. Define:

$$B_t = \max\{L_{i,t} : d_{i,t} = \text{buy}\},\ A_t = \min\{L_{j,t} : d_{j,t} = \text{sell}\} \tag{7}$$

If $B_t \geq A_t$, a trade occurs and the corresponding transaction price P_k is set by the midpoint rule:

$$P_\text{k} = \frac{B_t + A_t}{2} \tag{8}$$

Matched quantities are cleared from the book. Unmatched orders remain in the book unless expired. We use a fixed lifetime $\Delta T_\text{exp} = 10$ time steps for all the orders. Any order older than ΔT_exp steps will be canceled; namely,

$$o_{i,t'} \in \mathcal{O}_t \text{ is removed if } t - t' \geq \Delta T_\text{exp} \tag{9}$$

This rule prevents infinite accumulation of orders and ensures dynamic renewal of the order book. After processing all submitted orders in a simulation step, we obtain the set of executed trades with transaction prices $\{P_k\}_{k=1}^{K}$. We define the market price for the next step as the average transaction price within the step t,

$$P_{t+1} = \begin{cases} \frac{1}{K}\sum_k^K P_{t,k} \\ P_t \end{cases} \tag{10}$$

Since all trades have a unit size $q = 1$, this is equivalent to a volume-weighted average price (VWAP) because all executed trades are identical in size. Under this mechanism, the price series is generated endogenously from the submitted orders and the matching rules, without injecting additional external price shocks.

3. Results

This section presents the simulation results of our framework. We begin by demonstrating the effectiveness of our agent generation mechanism, demonstrating it can create a population of agents that is genuinely heterogeneous in its behavior. Subsequently, we will show the central success of the framework: that these data-driven agents, in collective interaction, can spontaneously reproduce a broad spectrum of key statistical features of real markets. Finally, through an ablation study and an analysis of varying agent compositions, we will further solidify our main arguments—that all the realistic, macro-level dynamics we observe are rooted in the agents' learned, data-driven internal intelligence.

3.1. Agent Properties and Heterogeneity

Before deploying our agents into the market simulation, it is essential to first validate that our generation methodology, particularly the concepts of "personality" (via Fit Quality) and "expertise" (via Data Composition), translates into measurable behavioral differences.

3.1.1. The Effect of Fit Quality on Agent Behavior

Our analysis reveals that an agent's FQ, or its degree of in-sample fit, has a profound impact on its behavioral "personality." For this study, all agents were trained on the same dataset, which consists of 1 min level returns of stocks from the Nikkei 225's food industry and 1 min resolution intraday price series (2019–2023 sampling pool) from the randomly selected stock in the Nikkei 225's food industry. In each run, we extract a continuous 10-trading-day segment, compute 1 min log-returns, and use them as the prediction target. All agents are trained in the same segment as described in Section 2. Four types of agents are trained by setting FQ to 60%, 70%, 80% and 90%. We deployed 400 agents in each group. Behavioral heterogeneity of these agents is measured by the prediction error of return. As illustrated in Figure 2, results of the training process show clear differences in the predicting behavior at the extremes of the FQ spectrum. The distributions in this figure were generated as follows: a single historical price context $\mathrm{x_t}$ was randomly sampled from the food industry and provided as a common input to all 400 agents within each FQ group. The "prediction error" (horizontal axis) is the difference between an agent's prediction $\hat{r}_t$ and the actual historical return y_t. The "frequency" (vertical axis) represents the count of agents per error bin, thus visualizing the collective behavioral response of each group to a shared historical context.

- The left panel of Figure 2 shows that agents with moderate FQ (70% vs. 80% sign-based accuracy) exhibit very similar prediction error distributions, indicating that small variations in FQ do not produce dramatically different behaviors.
- The right panel of Figure 2, however, reveals a contrast at the extremes. The high-fit "expert" agent (90% sign-based accuracy) produces a sharply concentrated error distribution, reflecting its deep memorization of familiar patterns. Conversely, the low-fit "novice" agent (60% sign-based accuracy) generates a significantly wider, more dispersed error distribution, reflecting its less certain grasp of the training data.

This observation confirms that a higher-FQ agent, within a familiar domain, behaves with more confidence and precision—though potentially with more rigidity. Note that the in-domain personality is crucial for interpreting how these agents would later behave when confronted with unfamiliar market conditions.

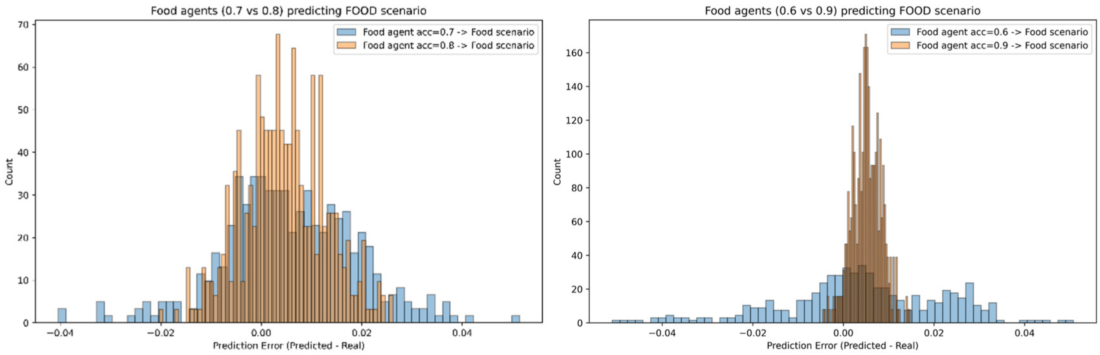

Figure 2. Behavioral impact of the in-domain Fit Quality. While agents with moderate Fit Quality (**left** panel, 70% vs. 80%) exhibit similar prediction error distributions, a stark difference emerges at the extremes (**right** panel, 60% vs. 90%), where the high-fit agents produce a much narrower (more concentrated) error distribution, indicating more precise and confident predictions, whereas the low-FQ "novice" agents' errors are far more dispersed.

3.1.2. Domain-Specific Expertise and Generalization Performance

A more complex picture emerges when these agents are confronted with unfamiliar, out-of-domain data. To investigate this, we designed an experiment utilizing four distinct agent groups, each composed of 400 agents. All groups were created using the unified architecture and formal training process described in Section 2. These groups are differentiated by two factors: "Expertise" and "Personality." Specifically, agents are trained in data from either the Food or Warehousing industry, with either high (90%) or low (60%) FQ. Figure 3 demonstrates the interplay between expertise and behavioral consistency when different groups of agents are evaluated with the prediction of stock returns in the Warehousing sector. The results are twofold:

- Agents within their domain (Specialists): The high-FQ Warehousing "expert" group (90% sign-based accuracy) produces a more concentrated distribution than the low-FQ "novice" group (60% sign-based accuracy), indicating higher behavioral consistency.
- Agents outside their domain (Non-specialists)**:** The opposite pattern emerges. The low-FQ Food "novice" group (60% sign-based accuracy) exhibits a more concentrated error distribution in predicting the Warehousing stock returns, while the high-FQ Food "expert" group (90% sign-based accuracy) becomes highly erratic with a much wider distribution.

Findings from Figure 3 are summarized in Table 2, as it is a bit more formal.

Table 2. Summary of agent behavioral patterns.

	In-Domain (Warehousing Agent Predicts Warehousing Data)	**Out-of-Domain (Food Agent Predicts Warehousing Data)**
Expert (High FQ)	Consistent & Homogeneous Behavior (Prediction errors are low and highly concentrated.)	Erratic & Heterogeneous Behavior (Prediction errors are biased and widely dispersed.)
Novice (Low FQ)	Less Consistent Behavior (Prediction errors are low but more dispersed than the expert's.)	Cautious & Agnostic Behavior (Prediction errors are biased but more concentrated than the expert's.)

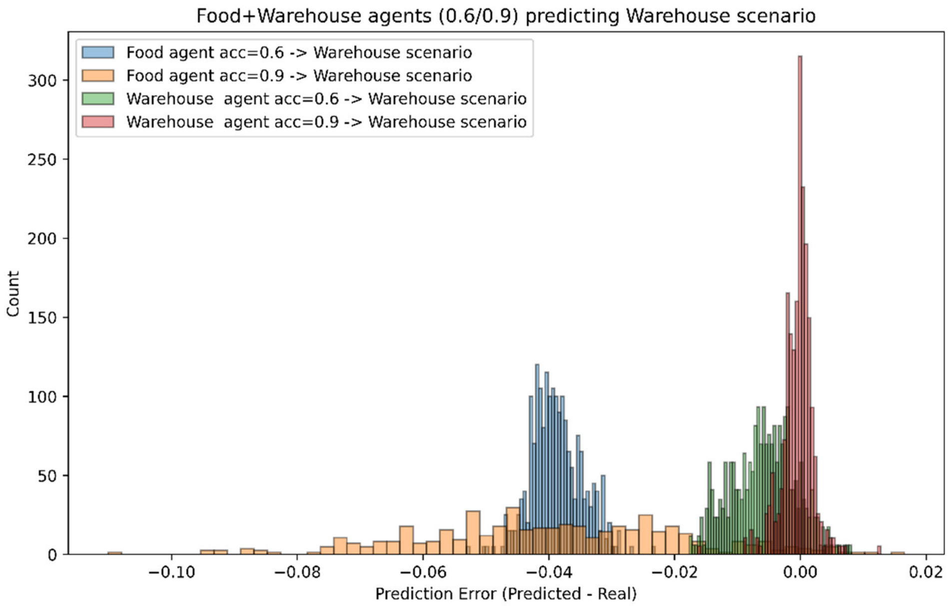

Figure 3. Behavioral consistency and domain expertise. High-FQ "expert" agents are only behaviorally consistent within their domain. When faced with an out-of-domain task, they become significantly more erratic and heterogeneous than their low-FQ "novice" counterparts.

3.1.3. Instantaneous Heterogeneity vs. Long-Term Convergence

To examine the nature of this heterogeneity in decision-making, we evaluated the same four agent groups used in the previous experiment (400 agents each, trained on Food and Warehousing data) on a new dataset unseen to all of them: historical data from the automotive industry.

Firstly, as shown in Figure 4, at any single predictive instance on an unfamiliar stock, diverse agents exhibit markedly dispersed forecasts, reflecting their unique, ingrained biases. We recognize that this instantaneous diversity is the micro-level foundation that provides liquidity and drives the (sometimes highly) fluctuating market dynamics.

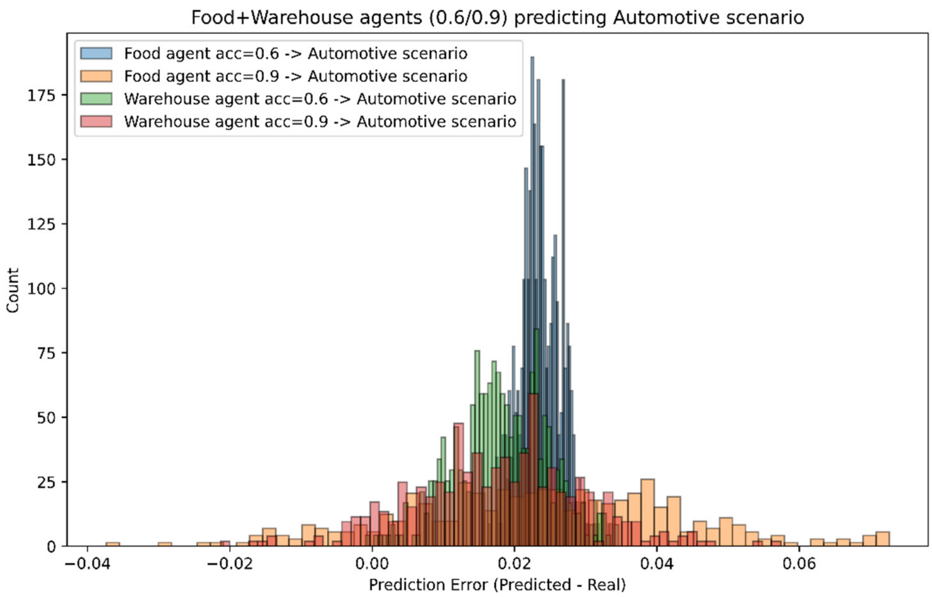

Figure 4. Instantaneous predictive heterogeneity. At any single moment, the agent population exhibits markedly different forecasts, reflecting the diverse biases that serve as the micro-level foundation for market dynamics.

However, from a long-term statistical perspective, these individual biases tend to average out. In our tests, convergence of the aggregate distribution was observed after roughly 1000–1500 steps. Figure 5 shows that if the experiment is iterated over numerous time steps, the aggregate distribution of predictions converges toward a more centralized, symmetric shape. The difference in error distributions due to different expertise becomes less pronounced. Agents with the same FQ begin to behave in an increasingly similar (consistent) manner. This long-term statistical convergence indicates that, although heterogeneity drives micro-level variation, the aggregate system exhibits macroscopic stability and efficient-like behavior at the collective level.

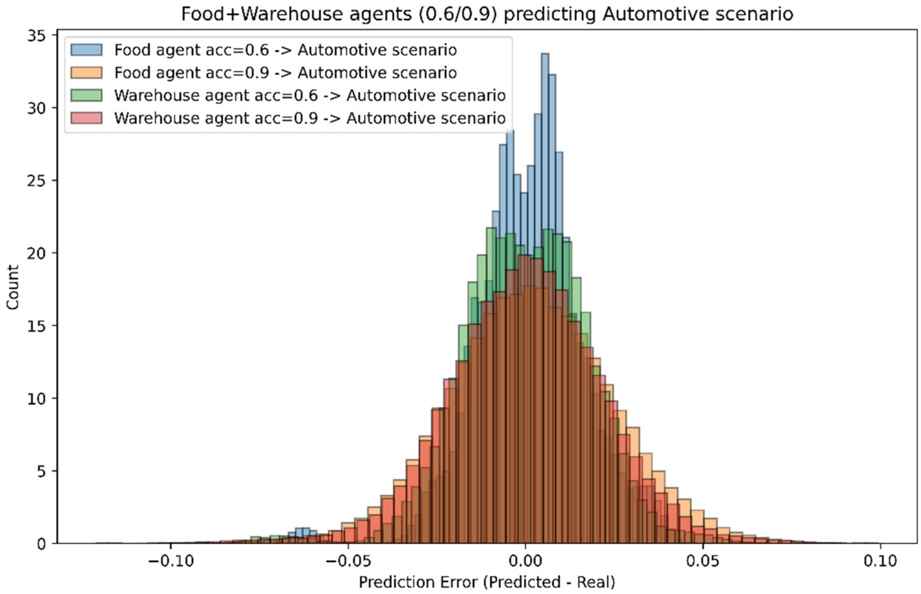

Figure 5. Long-Term convergence of aggregate predictions. In contrast to the instantaneous view (Figure 4), the long-term aggregate distribution converges to a symmetric shape, illustrating a statistical averaging effect.

3.2. Market Simulation

After verifying the micro-level behaviors of our heterogeneous agents, this section analyzes the emergent macro-level dynamics of the market simulation. First, we demonstrate the model's ability to reproduce the key stylized facts of financial markets. Second, we use an ablation study to isolate the causal source of these emergent properties. Finally, we investigate how the market's dynamics can be modulated by altering the composition of the agent population.

3.2.1. Emergence of Stylized Facts

A crucial benchmark for any agent-based financial model is its ability to spontaneously reproduce the "stylized facts" observed in real markets. Before presenting the representative results in Figure 6, we first clarify the experimental setup. This specific configuration consists of 400 agents trained on 1 min intraday price series from the JPX market (2019–2023) [26]. Concretely, for each training episode we randomly select stock within the target industry and extract a continuous 10-trading-day price sequence. The population contains a mixture of agents with different Fit Qualities (FQs): 10% of agents with 60% FQ (sign-based accuracy), 20% with 70% FQ, 50% with 80% FQ and 20% with 90% FQ. Notice that this specific composition is not a prerequisite for the emergence of stylized facts. Indeed, we found that the proposed framework is highly robust: the stylized facts emerge consistently across simulations with varying FQ compositions, provided that the total population of agents is sufficiently large. The threshold of the population, although empirically found to be N ≈ 400 in our experiment, is an important research topic to be investigated in the future. However, we have observed its complex dependency on multiple factors, such as the dataset used in training and the relative proportions of agents with different FQs.

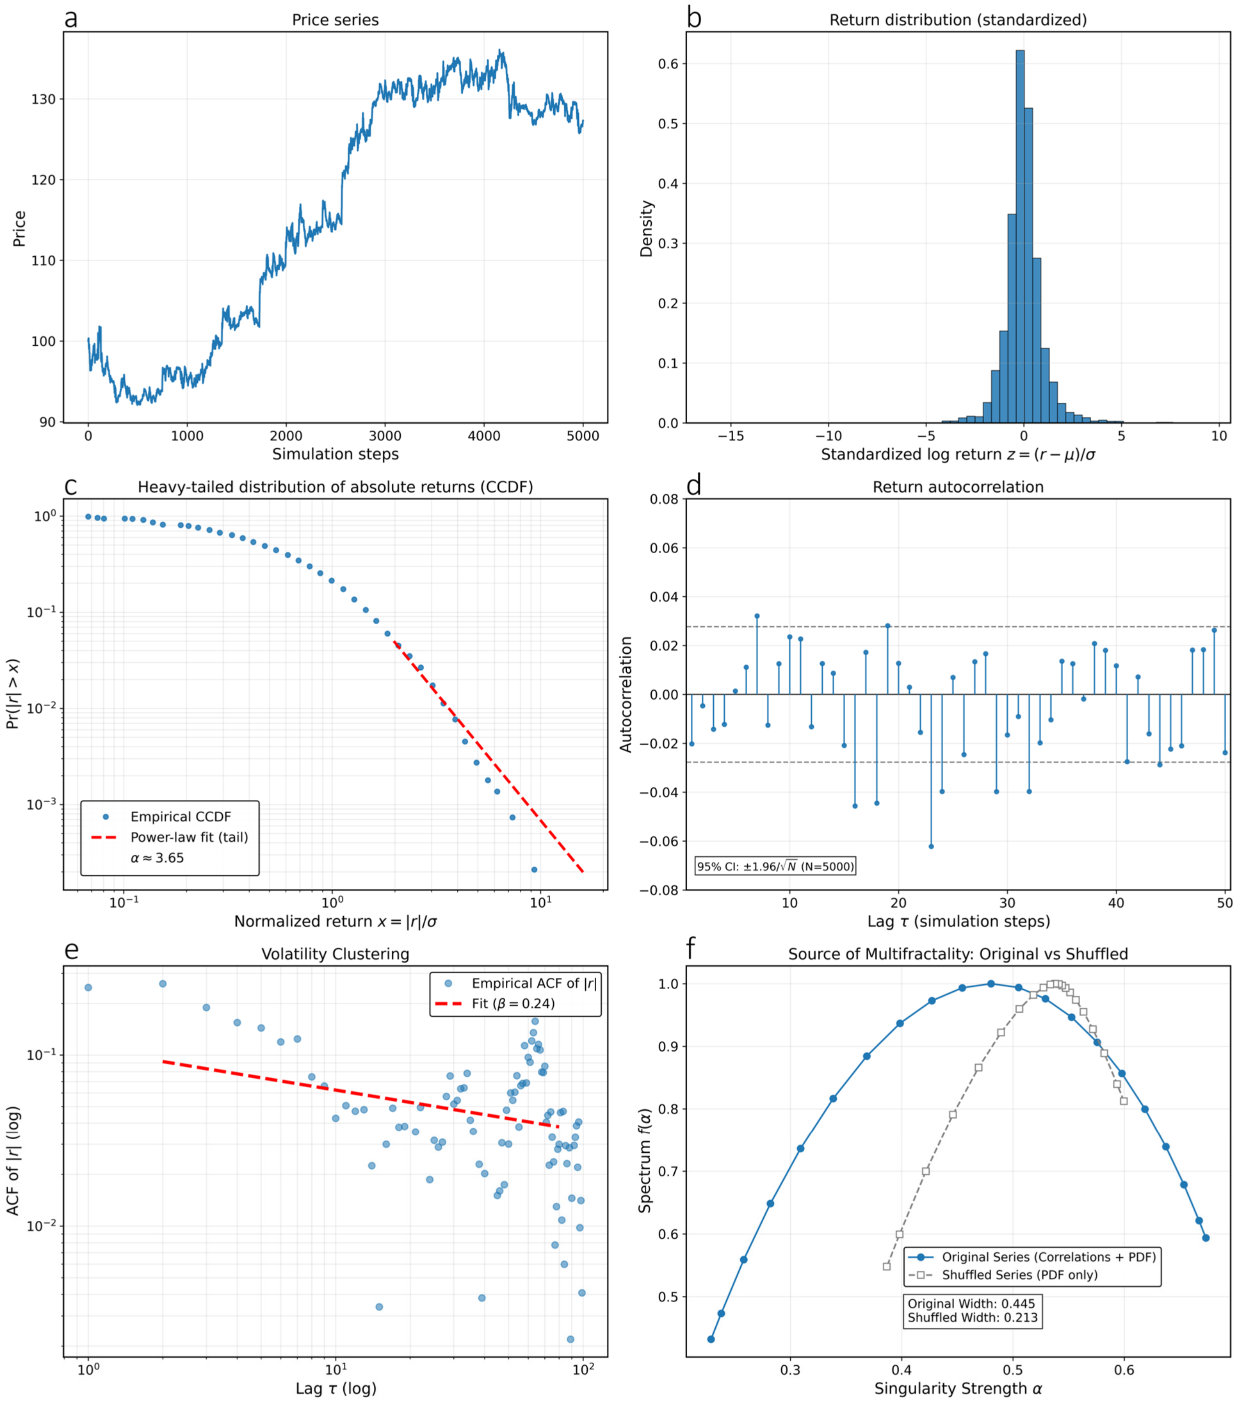

Figure 6. Reproduction of market stylized facts: (**a**) price series, (**b**) standardized return distribution, (**c**) heavy-tailed absolute returns shown by the CCDF with a power-law tail fit (red dashed line), (**d**) return autocorrelation with a 95% confidence band (gray dotted lines indicate the 95% confidence band), (**e**) volatility clustering quantified by the ACF of $|r|$ and a power-law decay fit (red dashed line), and (**f**) multifractality assessed by the MF-DFA singularity spectrum (original vs. shuffled).

As shown in Figure 6, this price-data-trained agent-based model for financial markets can capture a broad suite of well-known statistical properties:

- Price and Return: Figure 6a shows a representative simulated price series, and Figure 6b shows the corresponding distribution of standardized log-returns.
- Heavy-Tailed Returns: The complementary cumulative distribution function (CCDF) of normalized absolute returns $x = |r|/\sigma$ exhibits an approximately linear regime on log–log axes (Figure 6c), indicating a power-law tail. Following the standard practice in our pipeline, we fit the tail $x \geq x_q$, where x_q is the q-quantile of x (we use $q = 0.95$), via ordinary least squares in log–log space, and report the tail exponent $\alpha \approx 3.65$ (PDF form $p(x) \propto x^{-\alpha}$). The estimated α falls within the empirically accepted range reported for real markets [29,30].
- Absence of Linear Autocorrelations: Figure 6d shows the autocorrelation function (ACF) of raw returns up to 50 lags, together with the 95% confidence band $\pm 1.96/\sqrt{N}$. The correlations remain statistically negligible for essentially all lags, supporting linear unpredictability (efficient behavior) at the return level [31].
- Volatility Clustering: The ACF of absolute returns, $\rho_{|r|}(\tau)$, decays slowly over a wide range of lags (Figure 6e), indicating persistent volatility clustering. To quantify the long-memory decay, we fit $\rho_{|r|}(\tau) \propto \tau^{-\beta}$ on log–log axes over an intermediate-lag window $\tau \in [2, 80]$, excluding very small lags (microstructure-dominated) and very large lags (noise floor). The reported $\beta = 0.24$ is obtained by linear regression in log–log space using only positive ACF values in the chosen window [32].
- Multifractality: Figure 6f compares the singularity spectra $f(\alpha)$ obtained by multifractal detrended fluctuation analysis (MF-DFA) for the original return series and its temporally shuffled counterpart. Shuffling preserves the marginal distribution but removes temporal correlations; therefore, the larger spectral width in the original series indicates that multifractality is not solely driven by a heavy-tailed PDF but is substantially amplified by temporal dependence [33,34].

3.2.2. The Source of Complexity: An Ablation Study

To assess whether the emergence of stylized facts is driven by the agents' learned behaviors rather than being an artifact of the order-mapping and matching mechanics, we conducted an ablation study by constructing a "null" market. In the null market, we replaced the data-trained agents with "random-noise" traders: each agent's neural-network prediction was substituted by an i.i.d. random value drawn from a zero-mean Normal distribution, while keeping all other components (order mapping, matching engine, and market rules) identical. The variance of the random predictions was set to a small, fixed value so that the generated quotes remain within a comparable price scale to the trained-agent simulations.

Figure 7 compares the trained agent simulation and the null market. For heavy-tailed behavior, we plot the CCDF of standardized absolute log-returns, with normalization $x = |r|/\sigma$. Following the same pipeline as in Section 3.2.1, we fit the tail using a fixed quantile threshold and estimate the corresponding power-law exponent in the tail region. In this particular example, the fitted exponent is relatively large ($\alpha \approx 6.8$), as shown in Figure 7a, and is outside the range typically reported for empirical financial returns [29,30]. For linear return dependence, we compute the autocorrelation function (ACF) of raw returns together with the standard 95% confidence bounds. As shown in Figure 7c, the return ACF remains close to zero across lags in both the trained-agent simulation and the null market, with only occasional fluctuations. This is consistent with the standard expectation that raw returns are approximately linearly unpredictable at these horizons.

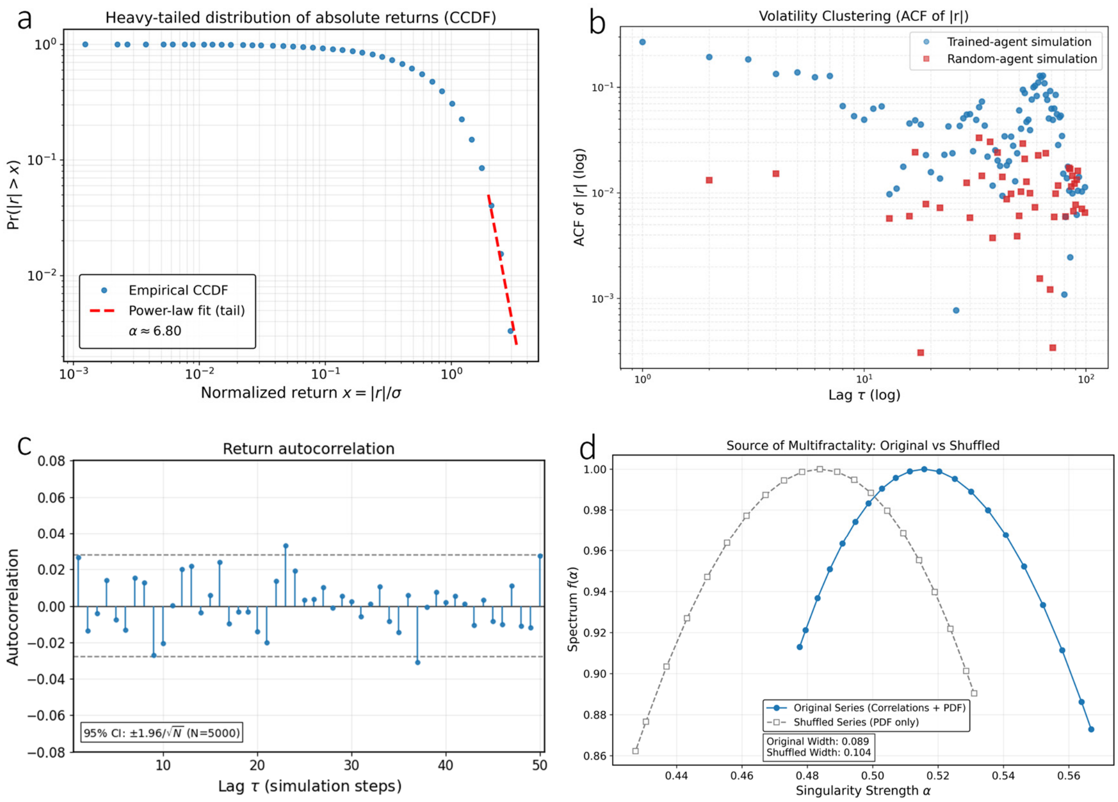

Figure 7. Ablation study based on stylized facts. (**a**) CCDF of normalized absolute log-returns with a power-law tail fit estimated over the fixed upper-tail region (red dashed line indicates the power-law tail fit). (**b**) Volatility clustering measured by the ACF of $|r|$ (log–log axes), comparing the trained-agent market to the random-noise ("null") market under identical order-mapping and matching mechanics. (**c**) Return autocorrelation function (ACF) with 95% confidence bands (gray dotted lines indicate the 95% confidence band). (**d**) Multifractal singularity spectrum $f(\alpha)$ from MF-DFA for the original return series and a shuffled surrogate (distribution preserved, temporal dependence removed); the spectrum width is reported for reference.

In contrast, a clearer difference between the two settings is observed in volatility clustering. We compute the ACF of absolute returns $|r|$ and visualize it on log–log axes to examine slow decay over lags. In Figure 7b, the trained-agent simulation exhibits systematically stronger persistence in ACF($|r|$) than the null market over a broad range of lags, whereas the null market produces weaker long-memory effect. This comparison suggests that volatility persistence is not solely a by-product of the matching mechanics and order-mapping rules, and that the trained agents' learned behaviors may contribute to generating clustered volatility under the same market infrastructure. Finally, we examine multifractality via MF-DFA. In Figure 7d, the spectrum widths for the original and shuffled series are both small, and the shuffled width is not smaller than the original. This suggests that any multifractality observed in the trained-agent market was largely due to the distribution of returns rather than dynamic correlations—and since both trained and null showed similarly small spectrum widths, the learned behaviors did not introduce significant multifractal scaling.

3.2.3. Influence of Agent Composition on Market Dynamics

A primary limitation of traditional ABMs is their fragility. They are often calibrated to reproduce a single market environment, and any significant parameter adjustment risks breaking the model's ability to generate stylized facts. Consequently, exploring a spectrum of different market styles requires a laborious, often intractable process of system-wide recalibration [6,35,36]. Here we conducted a series of experiments in which the influence of agent composition on the market risk profile was investigated, with all other components held fixed, to examine whether our framework can mitigate this limitation.

Each simulation consisted of a fixed population of 400 agents, all trained on the same Manufacturing-industry dataset to isolate the effect of FQ composition. We considered nine market compositions by gradually replacing low-FQ "novice" agents (FQ = 60%) with high-FQ "expert" agents (FQ = 90%): Group 1 contains 400 "novices" (400 × 60%); Group 2 contains 350 novices and 50 experts (350 × 60% + 50 × 90%); . . .; and Group 9 contains 400 "experts" (400 × 90%). For each group, we ran 100 independent simulations, computed the stylized-fact diagnostics for each run, and then summarized the results across runs (Figure 8).

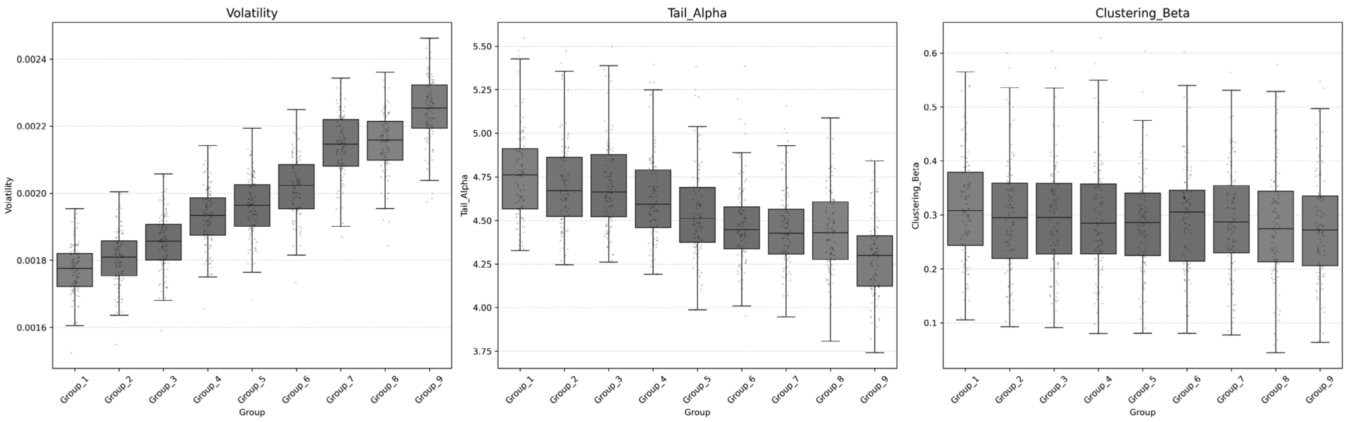

Figure 8. Market dynamics as a function of agent composition. Across nine compositions (400 agents total; the share of high-FQ agents increases from Group 1 to Group 9; 100 independent runs each), boxplots summarize the across-run distributions of event-scale volatility $\sigma(\Delta \log P)$ (left), tail index α of standardized $|r|$ estimated on the tail $x = |r|/\sigma \geq x_{0.95}$ (middle; PDF form $p(x) \propto x^{-\alpha}$), and clustering exponent β obtained from a log–log fit of ACF($|r|$) $\propto \tau^{-\beta}$ over $\tau \in [2, 80]$ using positive ACF values (right). Boxes indicate the interquartile range with median lines; whiskers show the full range across runs; dots represent individual runs. Across these settings, the volatility summaries shift upward and the estimated α tends to decrease as the high-FQ share increases, whereas β exhibits substantial overlap across groups.

Across this spectrum of compositions, the stylized-fact diagnostics discussed in Section 3.2.1 remain qualitatively present without requiring any system-wide recalibration. But traditional ABMs often need complete recalibration for different market conditions [6,35,36]; here we test whether simply changing the agent makeup can modulate the market statistics in our framework without breaking them. As summarized in Figure 8 (left), the distribution of the average volatility shifts upward as the proportion of high-FQ "expert" agents increases. This volatility metric is computed at the event scale consistent with the simulation's time step, and thus its numerical scale is not intended to be directly compared with conventional daily volatility. Nevertheless, the monotonic increase across Groups 1–9 indicates that composition changes can systematically modulate short-horizon price variability under otherwise identical market rules.

Figure 8 further reports two stylized-fact summary exponents estimated from each run: the tail index α (Figure 8 middle) and the volatility-clustering exponent β (Figure 8 right). The median α decreases as the market becomes more "expert"-dominated, suggesting progressively heavier tails under the standard power-law convention. In contrast, the β estimates show substantial overlap across groups, indicating that the long-memory proxy captured by $ACF(|r|)$ is comparatively less sensitive to composition within the present fitting window and horizon. Importantly, the resulting α and β values remain of the same order of magnitude as those reported in empirical studies (with the exact numerical range depending on sampling frequency and estimation protocol), supporting that the simulated markets stay within a broadly realistic regime [29–32].

A plausible interpretation is that, because all data-trained agents are (intentionally) overfitted to the same class of historical intraday series, their predictions tend to be more aligned than those generated by the random-noise baseline. Such alignment can increase the frequency with which many agents react in the same direction at similar times, which in turn raises the likelihood of large, short-horizon price moves and intermittency—consistent with the increase in volatility and the downward shift in the estimated tail index α in Figure 8. By contrast, the decay pattern of $ACF(|r|)$ (summarized by β) appears less systematically altered by composition alone in our setting, as the distributions of β largely overlap across groups. For transparency, we provide the run-level outputs underlying Figure 8 (all 100 runs per group) in the Supplementary Material, together with additional parameter cases to verify that these trends are not driven by a single hand-picked realization.

4. Conclusions and Future Work

In this work, motivated by the difficulty of inferring micro-level trading rules from aggregated price data, we introduced and evaluated a data-driven agent-based market in which neural agents trained on historical price series interact through a fixed continuous double-auction limit order book. Within this setting, we focused on how to represent and organize heterogeneity when one does not prespecify a small set of behavioral archetypes. The core contributions of this study are threefold.

1. A Shift in Trader Taxonomy via Fit Quality (FQ): Instead of assigning agents to a small set of ex ante labels such as "fundamentalist" or "trend-follower," we instantiate each trader as an individual neural mapping from recent price histories to order decisions, and organize heterogeneity along an effect-defined coordinate, Fit Quality. FQ summarizes how strongly a learned decision rule is anchored to in-sample historical patterns and is used purely as a neutral ex post descriptive and experimental axis, rather than as a definitive taxonomy of real-world strategies or a normative ranking of forecasting skill.
2. Reproduction of Key Stylized Facts under a Unified Agent Design: We examined the time series generated when populations of such agents interact under a simple order-submission and matching protocol. Within this unified agent and market design, and without hand-crafting distinct behavioral types or tuning parameters to target particular statistics, the simulated index-level returns exhibit several widely discussed stylized regularities of equity markets, such as heavy-tailed return distributions, weak linear autocorrelation of raw returns, and slowly decaying dependence in volatility proxies. In addition, an ablation experiment—replacing learned prediction signals with random-noise signals while keeping the remaining simulation components unchanged—leads to noticeably weaker stylized-fact signatures in the diagnostics. This comparison supports that the learned decision signals can contribute to the emergence of these statistics under the same market environment.
3. Controlled Modulation of Market Styles through FQ Composition: Finally, we used FQ as a control variable to study how the composition of the agent population along this axis shapes aggregate market statistics. By varying the mixture of lower-FQ and higher-FQ agents, we observe systematic but moderate shifts in summary exponents that characterize tail heaviness and volatility persistence, within ranges that remain broadly consistent with empirical benchmarks. In this sense, the distribution of "history-anchoring strength" across agents acts as a coarse control knob on market style, while preserving the presence of core stylized-fact patterns. We view this as a

way to probe the contribution of different segments of a data-trained agent population to aggregate behavior, rather than as a complete explanation of empirical markets.

Several extensions follow directly from the present design. On the modeling side, it would be informative to compare FQ with alternative effect-based coordinates—for example, measures of out-of-sample stability or cross-domain generalization—and to test whether the qualitative relationships reported here persist under different network architectures or training protocols. At the microstructure level, incorporating additional constraints such as richer order-type menus, inventory considerations, or more detailed cancellation rules could help clarify how far a simple matching mechanism can be pushed before such refinements become essential. Finally, broader sensitivity analyses—covering alternative agent architectures, training protocols, and market-simulation design choices—would help delineate which stylized-fact outcomes are robust across modeling degrees of freedom and which are regime-dependent.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/complexities2010004/s1.

Author Contributions: Q.Z. and Y.C. conceived the study; Q.Z. designed and performed the research; Q.Z. and Y.C. discussed the results and contributed to the text of the manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: The APC was funded by the Department of Human and Engineered Environment, The University of Tokyo, and supported by the International Graduate Program of Innovation for Intelligent World (IIW) and SPRING GX program.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The market data used in this study were obtained from the historical data service of Japan Exchange Group (JPX) [26]. As these data are provided under JPX's paid historical data offering and are subject to contractual and licensing restrictions, the raw data cannot be publicly shared by the authors. Interested readers may obtain equivalent data directly from JPX under a separate agreement.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Axtell, R.L.; Farmer, J.D. Agent-based modeling in economics and finance: Past, present, and future. *J. Econ. Lit.* **2025**, *63*, 197–287. [CrossRef]
2. Paulin, J.; Calinescu, A.; Wooldridge, M. Agent-based modeling for complex financial systems. *IEEE Intell. Syst.* **2018**, *33*, 74–82. [CrossRef]
3. Wang, L.; Ahn, K.; Kim, C.; Ha, C. Agent-based models in financial market studies. *J. Phys. Conf. Series.* **2018**, *1039*, 012022. [CrossRef]
4. Windrum, P.; Fagiolo, G.; Moneta, A. Empirical validation of agent-based models: Alternatives and prospects. *J. Artif. Soc. Soc. Simul.* **2007**, *10*, 8.
5. Fagiolo, G.; Guerini, M.; Lamperti, F.; Moneta, A.; Roventini, A. Validation of agent-based models in economics and finance. In *Computer Simulation Validation: Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*; Springer International Publishing: Cham, Switzerland, 2019; pp. 763–787.
6. Kukacka, J.; Kristoufek, L. Does parameterization affect the complexity of agent-based models? *J. Econ. Behav. Organ.* **2021**, *192*, 324–356. [CrossRef]
7. Alfarano, S.; Lux, T.; Wagner, F. Estimation of agent-based models: The case of an asymmetric herding model. *Comput. Econ.* **2005**, *26*, 19–49. [CrossRef]
8. Gontis, V.; Kononovicius, A. Consentaneous agent-based and stochastic model of the financial markets. *PLoS ONE* **2014**, *9*, e102201. [CrossRef] [PubMed]

9. Alfi, V.; Cristelli, M.; Pietronero, L.; Zaccaria, A. Minimal agent based model for financial markets I: Origin and self-organization of stylized facts. *Eur. Phys. J. B* **2009**, *67*, 385–397. [CrossRef]
10. Kononovicius, A.; Ruseckas, J. Order book model with herd behavior exhibiting long-range memory. *Phys. A Stat. Mech. Its Appl.* **2019**, *525*, 171–191. [CrossRef]
11. Lux, T. Estimation of agent-based models using sequential Monte Carlo methods. *J. Econ. Dyn. Control* **2018**, *91*, 391–408. [CrossRef]
12. Turgut, Y.; Bozdag, C.E. A framework proposal for machine learning-driven agent-based models through a case study analysis. *Simul. Model. Pract. Theory* **2023**, *123*, 102707. [CrossRef]
13. Ionescu, Ș.; Delcea, C.; Chiriță, N.; Nica, I. Exploring the use of artificial intelligence in agent-based modeling applications: A bibliometric study. *Algorithms* **2024**, *17*, 21. [CrossRef]
14. Jäger, G. Replacing rules by neural networks a framework for agent-based modelling. *Big Data Cogn. Comput.* **2019**, *3*, 51. [CrossRef]
15. Jäger, G. Using neural networks for a universal framework for agent-based models. *Math. Comput. Model. Dyn. Syst.* **2021**, *27*, 162–178. [CrossRef]
16. Jäger, G.; Reisinger, D. Can we replicate real human behaviour using artificial neural networks? *Math. Comput. Model. Dyn. Syst.* **2022**, *28*, 95–109. [CrossRef]
17. Shi, Z.; Cartlidge, J. Neural stochastic agent-based limit order book simulation: A hybrid methodology. *arXiv* **2023**, arXiv:2303.00080. [CrossRef]
18. Zhang, W.; Valencia, A.; Chang, N.-B. Synergistic integration between machine learning and agent-based modeling: A multidisciplinary review. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *34*, 2170–2190. [CrossRef] [PubMed]
19. Angione, C.; Silverman, E.; Yaneske, E. Using machine learning as a surrogate model for agent-based simulations. *PLoS ONE* **2022**, *17*, e0263150. [CrossRef] [PubMed]
20. Lussange, J.; Lazarevich, I.; Bourgeois-Gironde, S.; Palminteri, S.; Gutkin, B. Modelling stock markets by multi-agent reinforcement learning. *Comput. Econ.* **2021**, *57*, 113–147. [CrossRef]
21. Maeda, I.; Degraw, D.; Kitano, M.; Matsushima, H.; Sakaji, H.; Izumi, K.; Kato, A. Deep reinforcement learning in agent based financial market simulation. *J. Risk Financ. Manag.* **2020**, *13*, 71. [CrossRef]
22. Shavandi, A.; Khedmati, M. A multi-agent deep reinforcement learning framework for algorithmic trading in financial markets. *Expert Syst. Appl.* **2022**, *208*, 118124. [CrossRef]
23. Steinbacher, M.; Raddant, M.; Karimi, F.; Cuena, E.C.; Alfarano, S.; Iori, G.; Lux, T. Advances in the agent-based modeling of economic and social behavior. *SN Bus. Econ.* **2021**, *1*, 99. [CrossRef]
24. Engle, R.F. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **1982**, *50*, 987–1007. [CrossRef]
25. Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *J. Econom.* **1986**, *31*, 307–327. [CrossRef]
26. Japan Exchange Group (JPX). Historical Data (Paid Information Services (Equities)). Available online: https://www.jpx.co.jp/english/markets/paid-info-equities/historical/index.html (accessed on 6 February 2026).
27. Nikkei Indexes. Components: Nikkei Stock Average (Nikkei 225). Available online: https://indexes.nikkei.co.jp/en/nkave/index/component?idx=nk225 (accessed on 6 February 2026).
28. Gould, M.D.; Porter, M.A.; Williams, S.; McDonald, M.; Fenn, D.J.; Howison, S.D. Limit order books. *Quant. Financ.* **2013**, *13*, 1709–1742. [CrossRef]
29. Liu, Y.; Gopikrishnan, P.; Cizeau, P.; Meyer, M.; Peng, C.K.; Stanley, H.E. Statistical properties of the volatility of price fluctuations. *Phys. Rev. E* **1999**, *60*, 1390. [CrossRef]
30. Stanley, H.E.; Plerou, V.; Gabaix, X. A statistical physics view of financial fluctuations: Evidence for scaling and universality. *Phys. A Stat. Mech. Its Appl.* **2008**, *387*, 3967–3981. [CrossRef]
31. Chakraborti, A.; Toke, I.M.; Patriarca, M.; Abergel, F. Econophysics review: I. Empirical facts. *Quant. Financ.* **2011**, *11*, 991–1012. [CrossRef]
32. Cont, R. Volatility clustering in financial markets: Empirical facts and agent-based models. In *Long Memory in Economics*; Springer: Berlin/Heidelberg, Germany, 2007; pp. 289–309.
33. Kantelhardt, J.W.; Zschiegner, S.A.; Koscielny-Bunde, E.; Havlin, S.; Bunde, A.; Stanley, H. Multifractal detrended fluctuation analysis of nonstationary time series. *Phys. A Stat. Mech. Its Appl.* **2002**, *316*, 87–114. [CrossRef]
34. Thompson, J.R.; Wilson, J.R. Multifractal detrended fluctuation analysis: Practical applications to financial time series. *Math. Comput. Simul.* **2016**, *126*, 63–88. [CrossRef]

35. Srikrishnan, V.; Keller, K. Small increases in agent-based model complexity can result in large increases in required calibration data. *Environ. Model. Softw.* **2021**, *138*, 104978. [CrossRef]
36. Miserez, N.B. An Agent-Based Model of High-Frequency Trading. Master's Thesis, University of Zurich, Department of Business Administration, Zurich, Switzerland, 2017.