

High-Frequency Trading as Microsecond Bayesian Belief Formation: From Algorithmic Priors to Autonomous AI Market Agents

Matthew Long

YonedaAI Collaboration

Magneton Research, Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

Part IV of *Markets as Distributed Bayesian Inference Engines*

March 4, 2026

Abstract

The Predictive Market Theory (PMT) establishes that markets are collective predictive mechanisms whose fundamental output is a *belief function*—a continuously updated posterior probability measure over future states of the world. In this paper we ask: *how do high-frequency trading (HFT) systems form beliefs at microsecond timescales, and how does this extend to autonomous AI agents?*

We develop a formal framework in which HFT algorithms are modeled as *rapid Bayesian updaters* operating on compressed sufficient statistics extracted from the limit order book (LOB). We show that the belief formation process in HFT decomposes into three regimes: (i) *structural beliefs* encoded in algorithm design and calibration priors, (ii) *microstructure beliefs* formed through real-time inference on order flow, and (iii) *strategic beliefs* about the behavior of other market participants modeled as a Bayesian game.

We then extend this framework to *autonomous AI market agents*—systems that learn, adapt, and form beliefs through reinforcement learning, large language models, and multi-agent interaction. We prove that the transition from HFT to AI agents constitutes a *phase transition in the belief space*: HFT beliefs are constrained to a low-dimensional sufficient-statistic manifold, while AI agents can explore the full predictive measure space $\mathcal{P}$. This expansion has profound implications for market stability, information efficiency, and the nature of price discovery itself.

We derive conditions under which AI agent ecosystems converge to efficient predictive equilibria, characterize failure modes (belief cascades, adversarial manipulation, herding), and establish an *AI Market Stability Theorem* governing the relationship between agent diversity, belief entropy, and systemic risk.

Contents

1	Introduction: Belief at the Speed of Light	4
1.1	The Central Question	4
1.2	From Human to Algorithmic to AI Beliefs	4
1.3	Outline	5
2	HFT Belief Formation: A Bayesian Decomposition	5
2.1	The HFT Agent as Bayesian Updater	5
2.2	Structural Beliefs: The Algorithmic Prior	5
2.3	Microstructure Beliefs: Order Book Inference	6
2.4	Strategic Beliefs: The Bayesian Game	7
3	Belief Extraction from the Limit Order Book	8
3.1	The LOB as a Belief Artifact	8
3.2	Information Content of Order Flow	8
3.3	The Kyle Lambda as Belief Sensitivity	8
4	Multi-Timescale Belief Structure	9
4.1	The Belief Timescale Hierarchy	9
4.2	Belief Precision vs. Latency Trade-off	9
4.3	Belief Aggregation Across HFT Agents	10
5	Extension to Autonomous AI Market Agents	10
5.1	What Changes with AI Agents	10
5.2	The AI Belief Functional	11
5.3	LLM Agents: Beliefs from Language	11
5.4	Reinforcement Learning Agents: Beliefs from Interaction	12
6	The Belief Space Phase Transition	12
6.1	Dimensionality of Belief Spaces	12
6.2	Implications of the Phase Transition	13
7	AI Market Stability Theorem	13
7.1	Belief Diversity and Systemic Risk	13
7.2	Failure Modes	15
8	Multi-Agent Belief Ecosystems	15
8.1	Ecosystem Dynamics	15
8.2	Convergence to Predictive Equilibrium	15
8.3	Specialization and Division of Belief Labor	16
8.4	Emergent Market Intelligence	16
9	Implications for Market Design and Regulation	17
9.1	Regulatory Principles	17
9.2	Market Design for AI Ecosystems	17
9.3	The Grossman-Stiglitz Paradox Revisited	17

10 Computational Framework	18
10.1 Agent Architecture	18
10.2 Training Objective	18
10.3 Multi-Agent Simulation	19
11 Empirical Predictions	19
12 Connection to the Predictive Market Theory	20
12.1 Completing the PMT Framework	20
12.2 The Complete Belief Hierarchy	20
13 Conclusion	20

1 Introduction: Belief at the Speed of Light

1.1 The Central Question

In Parts I–III of this series, we established that the stock market—or any market—is a *collective predictive mechanism* that models perceived future outcomes as a belief function. Specifically, the market produces a *predictive measure* $\mathcal{P}_t$ over future states $\omega \in \Omega$, and prices emerge from this measure through a pricing functional:

$$P_t(X) = e^{-r(T-t)} \int_{\Omega} X(\omega) \cdot R(\omega) d\mathcal{P}_t(\omega), \quad (1)$$

where $R(\omega)$ is the risk distortion kernel mapping the predictive measure to the risk-neutral pricing measure.

This raises a fundamental question that the previous papers left open: *At what timescales and through what mechanisms do individual market participants form the beliefs that collectively constitute $\mathcal{P}_t$?*

In modern markets, the dominant participants by volume are not human traders but algorithms. High-frequency trading firms account for approximately 50–60% of U.S. equity volume, executing strategies at microsecond to millisecond timescales. These algorithms form “beliefs” not through human cognition but through computational inference on market data. Understanding how these algorithmic beliefs are formed, aggregated, and how they contribute to the collective predictive measure is essential for a complete theory of markets.

1.2 From Human to Algorithmic to AI Beliefs

We identify three epochs of belief formation in financial markets:

Epoch I. Human Belief (pre-2000): Beliefs formed through fundamental analysis, technical analysis, and qualitative judgment. Timescale: hours to months. Information set: public filings, news, expert opinion.

Epoch II. Algorithmic/HFT Belief (2000–present): Beliefs formed through statistical models operating on market microstructure data. Timescale: microseconds to seconds. Information set: order book, trade flow, latency signals.

Epoch III. AI Agent Belief (emerging): Beliefs formed through learned representations over multimodal data. Timescale: variable, from microseconds to strategic horizons. Information set: market data, text, images, alternative data, other agents’ behavior, and self-generated hypotheses.

The transition between these epochs is not merely technological but represents a *phase transition in belief space*—a qualitative change in the dimensionality, adaptability, and generative capacity of the belief functions that constitute the market’s predictive measure.

1.3 Outline

Section 2 develops the formal framework for HFT belief formation, decomposing it into structural, microstructure, and strategic components. Section 3 analyzes belief extraction from the limit order book. Section 4 characterizes the multi-timescale structure of HFT beliefs. Section 5 extends the framework to autonomous AI agents. Section 6 proves the belief space phase transition theorem. Section 7 derives the AI Market Stability Theorem. Section 8 analyzes multi-agent belief ecosystems. Section 9 discusses implications for market design and regulation. Section 13 concludes.

2 HFT Belief Formation: A Bayesian Decomposition

2.1 The HFT Agent as Bayesian Updater

We model an HFT agent as a constrained Bayesian updater. Let $\theta \in \Theta$ denote the latent state relevant to the agent’s trading decisions (e.g., the “true” short-term price direction, the presence of informed flow, the state of other participants’ inventories). The HFT agent maintains a belief:

$$\pi_t^{\text{HFT}}(\theta) = P(\theta \mid \mathcal{F}_t^{\text{micro}}), \quad (2)$$

where $\mathcal{F}_t^{\text{micro}}$ is the filtration generated by market microstructure observables up to time t .

Definition 2.1 (HFT Belief Decomposition). The belief of an HFT agent decomposes into three components:

$$\pi_t^{\text{HFT}}(\theta) = \underbrace{\pi_0(\theta)}_{\text{structural prior}} \cdot \underbrace{\frac{L(\mathcal{F}_t^{\text{LOB}} \mid \theta)}{Z_t^{\text{LOB}}}}_{\text{microstructure likelihood}} \cdot \underbrace{\frac{L(\mathcal{F}_t^{\text{strat}} \mid \theta)}{Z_t^{\text{strat}}}}_{\text{strategic likelihood}}, \quad (3)$$

where:

- $\pi_0(\theta)$ is the *structural prior* encoded in algorithm design and calibration,
- $L(\mathcal{F}_t^{\text{LOB}} \mid \theta)$ is the likelihood from limit order book observations,
- $L(\mathcal{F}_t^{\text{strat}} \mid \theta)$ is the likelihood from strategic inference about other agents,
- $Z_t^{\text{LOB}}, Z_t^{\text{strat}}$ are normalizing constants.

Each component operates on a different timescale and information source, creating a hierarchical belief structure.

2.2 Structural Beliefs: The Algorithmic Prior

The structural prior $\pi_0(\theta)$ is not learned in real-time but is *designed* into the algorithm through:

1. **Model specification:** The choice of model (e.g., linear regression on order imbalance, Hawkes process for order arrival, neural network for pattern recognition) constrains the hypothesis space Θ .
2. **Parameter calibration:** Historical backtesting produces point estimates or posterior distributions over model parameters.
3. **Risk constraints:** Position limits, loss limits, and inventory penalties shape the effective prior by zeroing out regions of Θ that would lead to unacceptable risk.
4. **Latency budget:** The computational time available for inference constrains the complexity of π_0 and the update mechanism.

Proposition 2.2 (Structural Prior Constraint). *Let $\mathcal{M}$ be the model class of an HFT algorithm with d free parameters and latency budget Δt . The effective dimension of the structural prior satisfies:*

$$\dim_{\text{eff}}(\pi_0) \leq \min \left(d, \left\lfloor \frac{\Delta t}{\tau_{\text{compute}}} \right\rfloor \right), \quad (4)$$

where τ_{compute} is the time required per dimension of Bayesian update.

This constraint is fundamental: HFT systems trade off belief complexity for speed. A market-making algorithm with $\Delta t = 10 \mu\text{s}$ latency budget cannot maintain a posterior over thousands of dimensions; it is restricted to a low-dimensional sufficient-statistic manifold.

2.3 Microstructure Beliefs: Order Book Inference

The microstructure component updates based on real-time observations from the limit order book:

Definition 2.3 (LOB Information Set). The LOB information set at time t is:

$$\mathcal{F}_t^{\text{LOB}} = \sigma \left(\{(p_i^b, q_i^b, p_j^a, q_j^a)\}_{i,j}, \{(\tau_k, s_k, p_k, q_k)\}_{k:\tau_k \leq t} \right), \quad (5)$$

where (p_i^b, q_i^b) are bid price-quantity pairs, (p_j^a, q_j^a) are ask price-quantity pairs, and (τ_k, s_k, p_k, q_k) are trade timestamps, sides, prices, and quantities.

The key insight is that HFT systems do not process the full LOB state but extract *sufficient statistics*:

Theorem 2.4 (LOB Sufficient Statistics). *For the class of linear-Gaussian price impact models, the following statistics are jointly sufficient for short-horizon price prediction:*

$$S_t = \left(\underbrace{\text{OI}_t}_{\text{order imbalance}}, \underbrace{\text{TI}_t}_{\text{trade imbalance}}, \underbrace{\sigma_t^{\text{micro}}}_{\text{microvolatility}}, \underbrace{\Lambda_t}_{\text{book depth}}, \underbrace{\dot{V}_t}_{\text{volume acceleration}} \right), \quad (6)$$

where:

$$\text{OI}_t = \frac{Q_t^b - Q_t^a}{Q_t^b + Q_t^a}, \quad \text{TI}_t = \frac{\sum_{k \in W_t} S_k q_k}{\sum_{k \in W_t} q_k}, \quad (7)$$

Q_t^b, Q_t^a are total bid/ask quantities within K levels, W_t is a trailing time window, σ_t^{micro} is realized volatility over W_t , Λ_t is depth-weighted book pressure, and $\dot{V}_t$ is the rate of change of volume.

Proof. Under the linear-Gaussian model $\Delta p_{t+\delta} = \alpha + \beta^\top S_t + \epsilon$, the posterior $P(\Delta p_{t+\delta} \mid \mathcal{F}_t^{\text{LOB}}) = P(\Delta p_{t+\delta} \mid S_t)$ by the factorization theorem for sufficient statistics. The Fisher information matrix $I(\beta) = \mathbb{E}[S_t S_t^\top] / \sigma_\epsilon^2$ is finite and positive definite under standard microstructure regularity conditions, confirming that S_t captures all information about the price change parameter β . $\square$

The belief update at the microstructure level is then:

$$\pi_{t+\delta}^{\text{micro}}(\theta) \propto \pi_t^{\text{micro}}(\theta) \cdot p(S_{t+\delta} \mid \theta), \quad (8)$$

which can be computed in $O(d^2)$ time for the sufficient statistic dimension $d = 5$.

2.4 Strategic Beliefs: The Bayesian Game

The strategic component models other market participants' behavior:

Definition 2.5 (HFT Bayesian Game). The HFT strategic environment is a Bayesian game $\Gamma = (N, \Theta, (A_i, u_i, \tau_i)_{i \in N})$ where:

- N is the set of market participants (HFT firms, institutional investors, retail),
- Θ is the type space (each agent's private information),
- A_i is agent i 's action space (order submission, modification, cancellation),
- $u_i : A \times \Theta \rightarrow \mathbb{R}$ is agent i 's utility,
- τ_i is agent i 's type distribution.

The strategic belief component captures:

1. **Adverse selection:** Probability that incoming order flow is informed vs. uninformed, updated from trade arrival patterns and order size distributions.
2. **Competitor state inference:** Estimation of other HFT firms' inventory positions and likely actions from observable market behavior.
3. **Institutional flow detection:** Inference about whether large orders are being executed (e.g., VWAP, TWAP algorithms) from volume patterns.

Proposition 2.6 (Strategic Belief Hierarchy). *In the HFT Bayesian game, the strategic belief hierarchy truncates at depth $K \leq 3$ in practice:*

$$\text{Level 0: } \pi_i^{(0)}(\theta_{-i}) = \tau_{-i} \quad (\text{prior over others' types}) \quad (9)$$

$$\text{Level 1: } \pi_i^{(1)}(\theta_{-i}) = P(\theta_{-i} \mid a_{-i}^{\text{obs}}) \quad (\text{infer types from actions}) \quad (10)$$

$$\text{Level 2: } \pi_i^{(2)}(\theta_{-i}) = P(\theta_{-i} \mid a_{-i}^{\text{obs}}, \hat{\pi}_{-i}^{(1)}) \quad (\text{infer others' inference}) \quad (11)$$

The truncation at level $K = 3$ is optimal given HFT latency constraints: the marginal information gain $\Delta I_K = I(\pi^{(K)}) - I(\pi^{(K-1)})$ satisfies $\Delta I_K < \tau_{\text{compute}} / \Delta t$ for $K > 3$.

3 Belief Extraction from the Limit Order Book

3.1 The LOB as a Belief Artifact

The limit order book is not merely a queueing system; it is a *belief artifact*—a spatial representation of the aggregate beliefs of all participants who have placed resting orders.

Theorem 3.1 (LOB Belief Density). *Under the assumption that each resting limit order at price p represents a belief that the asset is worth at least p (for bids) or at most p (for asks), the LOB induces a belief density:*

$$f_t^{\text{LOB}}(v) = \frac{1}{Z_t} \left(\left. \frac{\partial Q^b}{\partial p} \right|_{p=v} - \left. \frac{\partial Q^a}{\partial p} \right|_{p=v} \right), \quad (12)$$

where $Q^b(p)$ is the cumulative bid quantity up to price p , $Q^a(p)$ is the cumulative ask quantity down to price p , and Z_t is a normalizing constant.

The spread region $[p_t^{\text{best bid}}, p_t^{\text{best ask}}]$ corresponds to the region of maximum belief uncertainty—the market’s “indifference zone.”

3.2 Information Content of Order Flow

Individual orders carry information about the beliefs of their submitters:

Definition 3.2 (Order Information Content). The information content of an order event $e = (t, \text{type}, p, q)$ is:

$$I(e) = -\log p(e \mid \mathcal{F}_t^{\text{LOB}}) = -\log p(\text{type}, p, q \mid \mathcal{H}_t), \quad (13)$$

where $\mathcal{H}_t$ is the current LOB state. High-information events (surprise orders) trigger large belief updates; low-information events (expected queue replenishment) trigger small updates.

Proposition 3.3 (Informed Flow Detection). *The posterior probability that order flow in window $[t, t + \delta]$ is informed can be estimated via:*

$$P(\text{informed} \mid \mathcal{F}_{[t, t+\delta]}) = \frac{\alpha \cdot L_I(\mathcal{F}_{[t, t+\delta]})}{\alpha \cdot L_I(\mathcal{F}_{[t, t+\delta]}) + (1 - \alpha) \cdot L_U(\mathcal{F}_{[t, t+\delta]})}, \quad (14)$$

where α is the prior probability of informed trading, L_I is the likelihood under the informed-flow model (directional, size-clustered, persistent), and L_U is the likelihood under the uninformed-flow model (balanced, Poisson arrivals). This is a real-time Bayesian classifier that forms the core of adverse selection management in HFT market-making.

3.3 The Kyle Lambda as Belief Sensitivity

Kyle’s lambda—the price impact coefficient—has a natural interpretation as the *sensitivity of the market’s belief to order flow*:

Theorem 3.4 (Lambda as Belief Gradient). *In the Kyle (1985) framework embedded within PMT, the price impact coefficient is:*

$$\lambda = \frac{\partial \mathbb{E}[\mathcal{P}_t(V > p^*) \mid \text{flow}]}{\partial \text{flow}} = \frac{\sigma_V}{2\sigma_U}, \quad (15)$$

where σ_V is the fundamental value uncertainty and σ_U is noise trading intensity. This is the gradient of the market’s belief function with respect to observed order flow.

From the UPMF perspective (Part III), this connects to the thermodynamic structure:

$$\lambda = \frac{1}{2} \frac{\partial^2 \mathcal{F}_{\text{thermo}}}{\partial \langle E \rangle^2} \Big|_{\text{eq}}, \quad (16)$$

identifying price impact as the second derivative of thermodynamic free energy—the market’s “stiffness” against belief perturbations.

4 Multi-Timescale Belief Structure

4.1 The Belief Timescale Hierarchy

HFT beliefs operate across a hierarchy of timescales, each associated with different information sources and update mechanisms:

Definition 4.1 (Timescale Decomposition). The HFT belief decomposes across timescales:

$$\pi_t^{\text{HFT}}(\theta) = \int \pi_t^{(\tau)}(\theta) d\mu(\tau), \quad (17)$$

where $\pi_t^{(\tau)}$ is the belief component at timescale τ and μ is a measure over timescales. In practice, this decomposes into discrete bands:

$$\tau_1 \sim 1\text{--}100 \mu\text{s} : \quad \text{Latency arbitrage, queue position} \quad (18)$$

$$\tau_2 \sim 100 \mu\text{s}\text{--}10 \text{ ms} : \quad \text{Order book dynamics, short-term momentum} \quad (19)$$

$$\tau_3 \sim 10 \text{ ms}\text{--}1 \text{ s} : \quad \text{Informed flow detection, inventory management} \quad (20)$$

$$\tau_4 \sim 1 \text{ s}\text{--}5 \text{ min} : \quad \text{Regime detection, volatility estimation} \quad (21)$$

$$\tau_5 \sim 5 \text{ min}\text{--}1 \text{ day} : \quad \text{Daily pattern exploitation, news integration} \quad (22)$$

Theorem 4.2 (Timescale Separation Principle). *For well-designed HFT systems, the belief components at different timescales are approximately independent:*

$$I(\pi^{(\tau_i)}; \pi^{(\tau_j)}) \leq \epsilon_{ij} \quad \text{for } |i - j| \geq 2, \quad (23)$$

where $I(\cdot; \cdot)$ denotes mutual information and $\epsilon_{ij} \rightarrow 0$ as $\tau_j/\tau_i \rightarrow \infty$. This allows parallel belief computation across timescales without information loss.

4.2 Belief Precision vs. Latency Trade-off

Theorem 4.3 (Belief Uncertainty Relation). *For an HFT agent with latency budget Δt and computational capacity C (operations per second), the belief precision $\tau_\pi = 1/\text{Var}(\pi)$ satisfies:*

$$\tau_\pi \cdot \Delta t \leq \frac{C \cdot \Delta t^2}{d_{\text{eff}}}, \quad (24)$$

where d_{eff} is the effective belief dimension. This is analogous to the Heisenberg uncertainty relation: one cannot simultaneously achieve high belief precision and low latency.

The optimal trade-off satisfies:

$$\tau_{\pi}^* = \frac{C \cdot \Delta t}{d_{\text{eff}}}, \quad \Delta t^* = \sqrt{\frac{d_{\text{eff}} \cdot \tau_{\pi}}{C}}. \quad (25)$$

This theorem has a direct market structure implication: the “arms race” in HFT latency is actually an arms race in *belief formation speed*—the ability to form more precise beliefs faster.

4.3 Belief Aggregation Across HFT Agents

When multiple HFT agents operate simultaneously, their beliefs aggregate into a collective micro-timescale predictive measure:

Proposition 4.4 (HFT Belief Aggregation). *The collective HFT belief at the micro-timescale is:*

$$\mathcal{P}_t^{\text{HFT}} = \frac{\sum_{i \in \mathcal{I}_{\text{HFT}}} w_i \cdot \tau_{\pi,i} \cdot \pi_t^i}{\sum_{i \in \mathcal{I}_{\text{HFT}}} w_i \cdot \tau_{\pi,i}}, \quad (26)$$

where w_i is agent i ’s capital weight and $\tau_{\pi,i}$ is its belief precision. This is the precision-weighted consensus from Part I, now operating at microsecond timescales. The key insight is that HFT “price discovery” is Bayesian belief aggregation in fast-forward.

5 Extension to Autonomous AI Market Agents

5.1 What Changes with AI Agents

The transition from HFT to AI market agents is not merely quantitative (faster, more data) but qualitative. We identify five fundamental differences:

Definition 5.1 (AI Market Agent). An *AI market agent* is a market participant characterized by:

1. **Learned prior:** π_0 is not hand-designed but learned from data through neural network training, encoding implicit beliefs about market structure.
2. **Multimodal information processing:** The information set extends beyond market microstructure:

$$\mathcal{F}_t^{\text{AI}} = \mathcal{F}_t^{\text{LOB}} \vee \mathcal{F}_t^{\text{text}} \vee \mathcal{F}_t^{\text{alt}} \vee \mathcal{F}_t^{\text{social}} \vee \mathcal{F}_t^{\text{agent}}, \quad (27)$$

where $\mathcal{F}_t^{\text{text}}$ includes news, filings, and LLM-processed text; $\mathcal{F}_t^{\text{alt}}$ includes alternative data (satellite imagery, credit card transactions, web traffic); $\mathcal{F}_t^{\text{social}}$ includes social media sentiment; and $\mathcal{F}_t^{\text{agent}}$ includes observations of other AI agents’ behavior.

3. **Adaptive model:** The model class $\mathcal{M}$ itself evolves through online learning, not just parameters.

4. **Generative capacity:** The agent can form beliefs about states it has never observed, through learned generative models.
5. **Strategic depth:** The agent can model other agents’ learning processes, not just their current actions.

5.2 The AI Belief Functional

Definition 5.2 (AI Belief Functional). An AI market agent’s belief is a functional:

$$\pi_t^{\text{AI}} = \Phi_\theta(\mathcal{F}_t^{\text{AI}}), \quad (28)$$

where $\Phi_\theta : \mathcal{F} \rightarrow \Delta(\Theta)$ is a *learned mapping* from information to beliefs, parameterized by neural network weights θ . Unlike HFT beliefs, which are analytical functions of sufficient statistics, AI beliefs are *amortized inference*—the network has learned to approximate the Bayesian posterior without explicit computation.

Theorem 5.3 (Universal Belief Approximation). *For any measurable belief function $\pi^* : \mathcal{F} \rightarrow \Delta(\Theta)$ and any $\epsilon > 0$, there exists a neural network Φ_θ such that:*

$$\sup_t \text{KL}(\pi_t^* \parallel \Phi_\theta(\mathcal{F}_t)) < \epsilon. \quad (29)$$

That is, AI agents can approximate any belief formation process, including optimal Bayesian inference, given sufficient capacity.

Proof. This follows from the universal approximation theorem for neural networks applied to the space of probability measures. The key observation is that $\Delta(\Theta)$ can be embedded in $L^1(\Theta)$, and neural networks are universal approximators of measurable functions into Banach spaces (Hornik, 1991; Kidger and Lyons, 2020, for the modern continuous-depth variant). $\square$

5.3 LLM Agents: Beliefs from Language

Large language model (LLM) based market agents form beliefs through a fundamentally different channel: *natural language understanding of market narratives*.

Definition 5.4 (LLM Market Belief). An LLM market agent forms beliefs via:

$$\pi_t^{\text{LLM}}(\theta) \propto P_{\text{LLM}}(\theta \mid \text{prompt}_t), \quad (30)$$

where prompt_t encodes the current market context, recent news, analyst reports, and the agent’s reasoning chain. The LLM’s pre-trained weights encode an implicit prior over market dynamics learned from the entire textual history of financial markets.

The power of LLM beliefs lies in their ability to process *narratives*—causal stories about why markets move—that cannot be captured by statistical models of microstructure data alone.

Proposition 5.5 (Narrative Information Content). *Let I^{quant} be the information content of quantitative microstructure signals and I^{narr} be the information content of market narratives. For prediction horizons h :*

$$\frac{I^{\text{narr}}(h)}{I^{\text{quant}}(h)} \sim h^\gamma, \quad \gamma > 0, \quad (31)$$

i.e., narrative information becomes increasingly dominant at longer horizons. Quantitative microstructure dominates at $h < 1$ s; narratives dominate at $h > 1$ day.

5.4 Reinforcement Learning Agents: Beliefs from Interaction

RL-based market agents form beliefs through *interaction with the market environment*:

Definition 5.6 (RL Market Belief). An RL agent’s belief is implicitly encoded in its value function:

$$V_\theta(s_t) = \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k} \mid s_t, \pi_\theta \right], \quad (32)$$

where s_t is the market state, r_t is the reward (P&L), γ is the discount factor, and π_θ is the learned policy. The belief about future states is encoded in the gradient:

$$\pi_t^{\text{RL}}(\theta) \propto \exp \left(-\frac{1}{T_{\text{agent}}} \cdot (-V_\theta(\cdot)) \right), \quad (33)$$

where T_{agent} is the agent’s “temperature” (exploration parameter).

Remark 5.7. The RL agent’s temperature T_{agent} is distinct from the market temperature T in Parts I–III. However, they interact: when $T_{\text{agent}} \rightarrow 0$ (greedy exploitation), the agent’s beliefs become degenerate (point mass), reducing the entropy contribution to the market’s belief diversity.

6 The Belief Space Phase Transition

6.1 Dimensionality of Belief Spaces

The fundamental difference between HFT and AI agents can be characterized by the dimensionality of their belief spaces.

Definition 6.1 (Belief Space Dimension). The *effective belief dimension* of an agent class $\mathcal{A}$ is:

$$d_{\text{eff}}(\mathcal{A}) = \inf \{ d : \exists S : \mathcal{F} \rightarrow \mathbb{R}^d \text{ such that } \pi(\theta \mid \mathcal{F}) = \pi(\theta \mid S(\mathcal{F})) \text{ a.s.} \}, \quad (34)$$

i.e., the minimal dimension of a sufficient statistic for belief formation.

Theorem 6.2 (Belief Space Phase Transition). *Let $\mathcal{A}_{\text{HFT}}$ and $\mathcal{A}_{\text{AI}}$ denote the classes of HFT and AI agents respectively. Then:*

1. **HFT dimensionality:** $d_{\text{eff}}(\mathcal{A}_{\text{HFT}}) = O(\log |\Theta|)$, constrained by latency.
2. **AI dimensionality:** $d_{\text{eff}}(\mathcal{A}_{\text{AI}}) = O(|\Theta|)$ in principle, bounded only by model capacity.
3. **Phase transition:** As the fraction ϕ of AI agents in the market increases past a critical threshold ϕ_c , the effective dimension of the market’s collective belief undergoes a phase transition:

$$d_{\text{eff}}(\mathcal{P}_t) \sim \begin{cases} O(\log |\Theta|) & \text{if } \phi < \phi_c, \\ O(|\Theta|^\alpha) & \text{if } \phi > \phi_c, \end{cases} \quad (35)$$

where $0 < \alpha \leq 1$ depends on the diversity of AI architectures.

Proof. For Part 1: HFT agents are constrained to $d_{\text{eff}} \leq C \cdot \Delta t / \tau_{\text{compute}}$ by Theorem 4.3. For typical HFT parameters ($\Delta t \sim 10 \mu\text{s}$, $\tau_{\text{compute}} \sim 1 \mu\text{s}$, $C \sim 10^9 \text{ ops/s}$), this gives $d_{\text{eff}} \leq 10$, which is $O(\log |\Theta|)$ for $|\Theta| \sim 10^3$.

For Part 2: AI agents using neural networks with N parameters can represent belief distributions over Θ with $O(N)$ effective dimensions, and N can scale with $|\Theta|$ without latency constraints (since AI agents need not operate at HFT speeds).

For Part 3: The collective belief dimension of a mixture of agent types is:

$$d_{\text{eff}}(\mathcal{P}_t) = d_{\text{eff}}\left(\sum_i w_i \pi_i\right) \leq \sum_i d_{\text{eff}}(\pi_i). \quad (36)$$

When $\phi < \phi_c$, the HFT agents dominate and the collective belief is low-dimensional. When $\phi > \phi_c$, the AI agents' high-dimensional beliefs cannot be projected onto the HFT sufficient-statistic manifold without information loss, forcing the collective belief into higher dimensions. The transition is sharp because the rank of the belief covariance matrix has a threshold behavior as a function of ϕ . $\square$

6.2 Implications of the Phase Transition

Corollary 6.3 (Information Efficiency Jump). *At the phase transition $\phi = \phi_c$, the market's information efficiency (as measured by the mutual information between the predictive measure and the true state) undergoes a discontinuous increase:*

$$I(\mathcal{P}_t; \omega_{\text{true}}) \sim \begin{cases} I_{\text{low}} + O(\log \phi) & \text{if } \phi < \phi_c, \\ I_{\text{high}} + O(\phi^\alpha) & \text{if } \phi > \phi_c, \end{cases} \quad (37)$$

where $I_{\text{high}} - I_{\text{low}} > 0$ is the “efficiency gap” between the HFT-dominated and AI-dominated regimes.

Corollary 6.4 (Volatility Regime Change). *The phase transition alters volatility structure:*

$$\sigma_t^2 \sim \begin{cases} \sigma_{\text{micro}}^2 + \frac{T}{d_{\text{eff,HFT}}} & \text{if } \phi < \phi_c, \\ \sigma_{\text{micro}}^2 + \frac{T}{d_{\text{eff,AI}}} & \text{if } \phi > \phi_c. \end{cases} \quad (38)$$

Since $d_{\text{eff,AI}} \gg d_{\text{eff,HFT}}$, the AI-dominated regime has lower fundamental volatility but potentially higher tail risk from belief correlation.

7 AI Market Stability Theorem

7.1 Belief Diversity and Systemic Risk

The central stability question is: *under what conditions does an AI agent ecosystem converge to a stable, efficient predictive equilibrium?*

Definition 7.1 (Belief Diversity). The *belief diversity* of an agent ecosystem $\{\pi_i\}_{i=1}^N$ is:

$$\mathcal{D}_t = \frac{1}{N^2} \sum_{i,j} \text{KL}(\pi_i^t \| \pi_j^t), \quad (39)$$

measuring the average divergence between agents' beliefs.

Definition 7.2 (Belief Correlation). The *belief correlation* is:

$$\rho_t^{\text{belief}} = 1 - \frac{\mathcal{D}_t}{\mathcal{D}_{\max}}, \quad (40)$$

where $\mathcal{D}_{\max}$ is the maximum achievable diversity for the given agent architectures.

Theorem 7.3 (AI Market Stability Theorem). *An AI market agent ecosystem is stable if and only if the following conditions hold:*

1. **Diversity condition:** *The belief diversity satisfies:*

$$\mathcal{D}_t > \mathcal{D}_{\text{crit}} = \frac{T_c}{T} \cdot \log \left(\frac{V_{\text{total}}}{\Lambda_{\min}} \right), \quad (41)$$

where T_c is the critical temperature from Part I, V_{total} is total market volume, and $\Lambda_{\min}$ is minimum liquidity.

2. **Update rate bound:** *The maximum belief update rate satisfies:*

$$\sup_i \left\| \frac{d\pi_i^t}{dt} \right\|_{\text{TV}} < \frac{\Lambda_t}{N \cdot \bar{q}}, \quad (42)$$

where $\|\cdot\|_{\text{TV}}$ is total variation norm, Λ_t is instantaneous liquidity, and $\bar{q}$ is average order size.

3. **Ergodicity:** *The Markov chain on the collective belief state is ergodic:*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mathcal{P}_t dt = \mathcal{P}_{\text{eq}}^* \quad a.s. \quad (43)$$

When any condition is violated:

- *Violation of (1):* Belief cascade—agents’ beliefs converge, collapsing diversity below critical threshold, triggering a first-order phase transition (crash or bubble).
- *Violation of (2):* Liquidity crisis—belief updates exceed the market’s ability to absorb order flow, causing flash crash dynamics.
- *Violation of (3):* Belief lock-in—the market becomes trapped in a non-equilibrium attractor, failing to incorporate new information.

Proof sketch. Condition (1) follows from the connection between belief diversity and entropic liquidity (Part I, Section 8): $\Lambda = \exp(H[\text{beliefs}])$, and the market is stable iff $\Lambda > \Lambda_{\min}$, which requires sufficient belief entropy, which requires sufficient diversity.

Condition (2) follows from the market’s finite absorption capacity. If all agents update beliefs simultaneously, the resulting order flow is $\Delta Q \sim N \cdot \bar{q} \cdot \|d\pi/dt\|_{\text{TV}}$. For the market to absorb this without price discontinuity, we need $\Delta Q < \Lambda_t$.

Condition (3) follows from the requirement that the collective belief converges to the PMT predictive measure $\mathcal{P}^*$, which requires ergodicity of the underlying Markov dynamics (the system must eventually explore all relevant belief states). $\square$

7.2 Failure Modes

Proposition 7.4 (Belief Cascade Dynamics). *When belief diversity drops below $\mathcal{D}_{\text{crit}}$, the time to cascade (crash) follows:*

$$\tau_{\text{cascade}} \sim \frac{1}{\mathcal{D}_{\text{crit}} - \mathcal{D}_t} \cdot \frac{\Lambda_t}{V_t}, \quad (44)$$

which diverges at the critical point and goes to zero as diversity collapses—consistent with the “calm before the storm” empirically observed before market crashes.

Proposition 7.5 (Adversarial Belief Manipulation). *An adversarial agent can manipulate the collective belief by injecting signals that exploit the training distribution of AI agents. The vulnerability is:*

$$\Delta \mathcal{P} = \sum_i w_i \cdot \nabla_{\mathcal{F}} \Phi_{\theta_i}(\mathcal{F}) \cdot \delta \mathcal{F}_{\text{adv}}, \quad (45)$$

where $\delta \mathcal{F}_{\text{adv}}$ is an adversarial perturbation. The attack surface scales with the homogeneity of agent architectures: if all agents use similar neural networks, a single adversarial signal can shift the collective belief.

8 Multi-Agent Belief Ecosystems

8.1 Ecosystem Dynamics

Definition 8.1 (Market Belief Ecosystem). *A market belief ecosystem is a tuple $(\mathcal{A}, \mathcal{E}, \mathcal{I}, \mathcal{R})$ where:*

- $\mathcal{A} = \{a_1, \dots, a_N\}$ is the set of AI agents,
- $\mathcal{E} : \mathcal{A} \times \mathcal{A} \rightarrow [0, 1]$ is the *interaction graph* (information exchange intensity),
- $\mathcal{I} = \{\mathcal{F}_i\}_{i=1}^N$ is the information partition,
- $\mathcal{R} = \{r_i\}_{i=1}^N$ are reward functions (market P&L plus any additional objectives).

The ecosystem evolves through coupled belief dynamics:

$$\frac{d\pi_i^t}{dt} = -\eta_i \nabla_{\pi_i} \mathcal{F}[\pi_i^t; \mathcal{F}_i^t] + \sum_{j \neq i} \mathcal{E}_{ij} \cdot (\pi_j^t - \pi_i^t) + \xi_i^t, \quad (46)$$

where the first term is free energy minimization (each agent improving its belief), the second term is belief exchange through the interaction graph, and ξ_i^t is stochastic innovation from new information.

8.2 Convergence to Predictive Equilibrium

Theorem 8.2 (Ecosystem Convergence). *If the interaction graph $\mathcal{E}$ is connected and the information partition $\mathcal{I}$ is such that $\bigvee_i \mathcal{F}_i = \mathcal{F}_{\text{market}}$ (collective completeness), then the ecosystem converges to the PMT predictive measure:*

$$\pi_i^t \xrightarrow{t \rightarrow \infty} \mathcal{P}^* \quad \forall i, \quad (47)$$

at rate:

$$\text{KL}(\pi_i^t || \mathcal{P}^*) \leq C \cdot e^{-\lambda_2(\mathcal{E}) \cdot t}, \quad (48)$$

where $\lambda_2(\mathcal{E})$ is the spectral gap of the interaction graph Laplacian.

Proof. The dynamics (46) combine local free energy descent with consensus-seeking on the interaction graph. The free energy term drives each agent toward the locally optimal belief given its information; the consensus term drives agreement. Under connectedness of $\mathcal{E}$ and collective completeness of $\mathcal{I}$, the unique fixed point is $\pi_i = \mathcal{P}^*$ for all i , and convergence follows from standard results on consensus dynamics on graphs (Olfati-Saber and Murray, 2004) combined with the contraction property of free energy minimization. $\square$

8.3 Specialization and Division of Belief Labor

Proposition 8.3 (Belief Specialization). *In equilibrium, AI agents specialize in belief formation along information dimensions:*

$$\pi_i^* \approx \mathcal{P}^* \text{ restricted to } \mathcal{F}_i, \quad (49)$$

i.e., each agent becomes an “expert” in its information domain. The market achieves collective intelligence through division of belief labor: no single agent needs to process all information, but the ecosystem as a whole covers the full information space.

The efficiency of belief specialization is:

$$\eta_{\text{spec}} = \frac{I(\mathcal{P}^*; \omega_{\text{true}})}{N \cdot \bar{I}_{\text{agent}}} \geq 1, \quad (50)$$

where $\bar{I}_{\text{agent}}$ is the average information per agent. The inequality $\eta_{\text{spec}} \geq 1$ means specialization is always at least as efficient as uniform processing, with strict inequality when information domains are separable.

8.4 Emergent Market Intelligence

Definition 8.4 (Collective Intelligence Measure). *The collective intelligence of the market is:*

$$\mathcal{CI}_t = I(\mathcal{P}_t; \omega_{\text{true}}) - \max_i I(\pi_i^t; \omega_{\text{true}}), \quad (51)$$

the excess information in the collective belief over the best individual agent.

Theorem 8.5 (Superadditivity of Market Belief). *Under the conditions of Theorem 8.2, the collective intelligence is strictly positive:*

$$\mathcal{CI}_t > 0 \quad \text{for all } t > t_0, \quad (52)$$

i.e., the market “knows more” than any individual AI agent. Moreover:

$$\mathcal{CI}_t \geq H[\text{agent specializations}] = - \sum_i p_i \log p_i, \quad (53)$$

where p_i is the fraction of total information contributed by agent i . Collective intelligence is maximized when agents are maximally diverse in their information sources.

9 Implications for Market Design and Regulation

9.1 Regulatory Principles

The framework yields concrete regulatory principles:

1. **Diversity mandates:** Regulators should monitor and promote AI agent diversity. A market dominated by agents using the same architecture (e.g., all using GPT-based systems) is vulnerable to correlated belief failures. Minimum diversity requirements analogous to biodiversity preservation are warranted.
2. **Belief update speed limits:** The update rate bound (Theorem 7.3, condition 2) provides a principled basis for speed limits—not arbitrary latency floors, but limits tied to the market’s absorption capacity.
3. **Adversarial robustness requirements:** AI agents should be tested against adversarial belief manipulation (Theorem 7.5) before deployment, analogous to stress testing for banks.
4. **Transparency of belief formation:** Regulators should require AI agents to make their belief formation process interpretable—not necessarily revealing proprietary strategies, but providing sufficient information to assess systemic risk contribution.

9.2 Market Design for AI Ecosystems

Proposition 9.1 (Optimal Market Microstructure for AI). *The optimal market microstructure for AI agent ecosystems differs from the current HFT-optimized design:*

1. **Batch auctions at AI-natural frequencies:** *Rather than continuous limit order books, batch auctions at intervals matching the AI belief update cycle (~ 1 s to 1 min) would improve belief aggregation quality.*
2. **Rich signal spaces:** *Markets should provide structured data channels beyond price and quantity—belief reports, confidence intervals, reasoning traces—to improve inter-agent information exchange.*
3. **Belief-weighted voting:** *Order priority should account for belief precision, not just time priority. More informed beliefs (higher τ_π) should have greater price impact.*

9.3 The Grossman-Stiglitz Paradox Revisited

Theorem 9.2 (AI Resolution of Grossman-Stiglitz). *In AI agent markets, the Grossman-Stiglitz paradox is resolved differently than in human markets:*

$$\text{Cost of belief formation} = \underbrace{C_{\text{compute}}}_{\text{near-zero marginal}} + \underbrace{C_{\text{data}}}_{\text{diminishing}} + \underbrace{C_{\text{model}}}_{\text{amortized}}, \quad (54)$$

where $C_{\text{compute}} \rightarrow 0$ with Moore's law, C_{data} is subject to diminishing returns, and C_{model} is a one-time training cost amortized over the agent's lifetime. This means the equilibrium level of market efficiency increases monotonically with AI adoption:

$$\eta_{\text{eff}}(\phi) = 1 - \frac{C(\phi)}{\bar{r}(\phi)}, \quad (55)$$

where $C(\phi)$ is the aggregate cost of belief formation at AI adoption level ϕ and $\bar{r}(\phi)$ is the average return to informed trading. Since $C(\phi) \rightarrow 0$, we get $\eta_{\text{eff}} \rightarrow 1$ —near-perfect efficiency.

10 Computational Framework

10.1 Agent Architecture

We provide the computational architecture for the AI market agent belief system, extending the UPMF implementation from Part III.

Definition 10.1 (Computational Belief Agent). The computational agent implements:

1. **Belief Encoder:** Neural network $f_\theta : \mathcal{F}_t \rightarrow \mathbb{R}^d$ mapping observations to a latent belief representation.
2. **Belief Decoder:** Network $g_\phi : \mathbb{R}^d \rightarrow \Delta(\Theta)$ mapping latent representation to a probability distribution over states.
3. **Update Rule:** $z_{t+1} = z_t + \alpha \cdot h_\psi(z_t, \delta\mathcal{F}_t)$, where h_ψ is a learned update function and $\delta\mathcal{F}_t$ is new information.
4. **Action Policy:** $a_t = \pi_\omega(z_t)$, mapping belief state to market actions.

The full pipeline is:

$$\mathcal{F}_t \xrightarrow{f_\theta} z_t \xrightarrow{g_\phi} \pi_t(\theta) \xrightarrow{\text{pricing}} P_t(X) \xrightarrow{\pi_\omega} a_t \xrightarrow{\text{market}} \mathcal{F}_{t+1}. \quad (56)$$

10.2 Training Objective

Proposition 10.2 (Optimal Training Loss). *The optimal training objective for an AI market agent combines predictive accuracy with free energy minimization:*

$$\mathcal{L}(\theta, \phi, \psi, \omega) = \underbrace{\mathbb{E}_t[-\log p(x_{t+h} | g_\phi(f_\theta(\mathcal{F}_t)))]}_{\text{predictive loss}} + \underbrace{\beta \text{KL}(g_\phi(f_\theta(\mathcal{F}_t)) || \pi_0)}_{\text{complexity penalty}} - \underbrace{\gamma \mathbb{E}_t[r_t(\pi_\omega(z_t))]}_{\text{reward}}, \quad (57)$$

where β controls belief complexity (analogous to market temperature) and γ controls the trade-off between prediction and profit.

10.3 Multi-Agent Simulation

The ecosystem simulation extends the MarketSimulator from Part III:

Definition 10.3 (AI Ecosystem Simulator). The simulator evolves N agents on a shared LOB:

1. Initialize N agents with diverse architectures and information feeds.
2. At each timestep:
 - (a) Generate market events (news, order flow, fundamental signals).
 - (b) Each agent updates its belief: $\pi_i^{t+1} = \Phi_{\theta_i}(\mathcal{F}_i^{t+1})$.
 - (c) Each agent submits orders: $a_i^t = \pi_{\omega_i}(z_i^t)$.
 - (d) LOB processes orders, generates trades, updates prices.
 - (e) Agents receive P&L and update through online learning.
3. Measure: collective intelligence $\mathcal{CI}_t$, belief diversity $\mathcal{D}_t$, market efficiency η_t .

11 Empirical Predictions


The framework generates specific, testable predictions:

1. **HFT belief dimension:** The effective dimension of HFT belief formation should be ≤ 10 , measurable through the rank of the cross-correlation matrix of HFT order flow features.
2. **Phase transition signature:** As AI agent market share increases, there should be a discontinuous increase in market information efficiency at $\phi \approx \phi_c$, detectable as a structural break in return predictability regressions.
3. **Diversity-stability correlation:** Markets with more diverse AI agent architectures should exhibit lower crash frequency and smaller tail risk, measurable through cross-sectional comparison of markets with different AI adoption profiles.
4. **Narrative vs. quantitative information ratio:** The relative importance of textual vs. quantitative signals should follow $I^{\text{narr}}/I^{\text{quant}} \sim h^\gamma$ with $\gamma \approx 0.5$, testable through horizon-dependent return prediction comparisons.
5. **Belief cascade precursors:** Declining belief diversity $\mathcal{D}_t$ should predict market crashes with lead time $\tau_{\text{cascade}} \propto (\mathcal{D}_{\text{crit}} - \mathcal{D}_t)^{-1}$, testable through real-time monitoring of algorithmic trading diversity metrics.
6. **Adversarial vulnerability:** Markets dominated by homogeneous AI architectures should be more susceptible to manipulation, testable through experimental market simulations.

12 Connection to the Predictive Market Theory

12.1 Completing the PMT Framework

This paper completes the PMT framework by filling in the *microstructure of belief formation*:



- Part I established *what* the market computes (Bayesian posterior).
- Part II established *why* (predictive measure minimizes free energy).
- Part III established *how* (UPMF dynamics and computational implementation).
- Part IV (this paper) establishes *who* computes and *at what timescales* (HFT and AI agents as belief-forming participants).

12.2 The Complete Belief Hierarchy

The full belief hierarchy of the market is now:

$$\mathcal{P}_t = \underbrace{\sum_{i \in \mathcal{I}_{\text{human}}} w_i \pi_i^{\text{human}}}_{\text{human beliefs (slow)}} + \underbrace{\sum_{j \in \mathcal{I}_{\text{HFT}}} w_j \pi_j^{\text{HFT}}}_{\text{HFT beliefs (fast)}} + \underbrace{\sum_{k \in \mathcal{I}_{\text{AI}}} w_k \pi_k^{\text{AI}}}_{\text{AI beliefs (adaptive)}}, \quad (58)$$

where the weights w_i, w_j, w_k are proportional to capital deployed and belief precision. The market's predictive power comes from the *diversity* of these belief sources across timescales, information domains, and computational architectures.

13 Conclusion

We have developed a formal framework for understanding how high-frequency trading systems and autonomous AI agents form the beliefs that collectively constitute the market's predictive measure. The key results are:

1. HFT belief formation decomposes into structural (designed), microstructure (real-time LOB inference), and strategic (game-theoretic) components, each operating at different timescales and constrained by the latency-precision trade-off.

2. AI agents represent a qualitative expansion of the belief space, processing multimodal information through learned representations that can approximate any belief formation process.
3. The transition from HFT-dominated to AI-dominated markets constitutes a *phase transition in belief space*, with implications for market efficiency, volatility structure, and systemic risk.
4. Market stability requires sufficient belief diversity, bounded update rates, and ergodic dynamics—conditions formalized in the AI Market Stability Theorem.
5. The market’s collective intelligence exceeds any individual agent’s, with the surplus proportional to the diversity of agent specializations.

The central insight connecting all four parts of this series is now complete: **the market is a distributed belief computation engine, and understanding its behavior requires understanding how each participant—human, algorithmic, and AI—forms, updates, and aggregates beliefs about the future.** The predictive measure $\mathcal{P}_t$ is not an abstraction but the real-time output of this vast, heterogeneous, multi-timescale belief formation process.

As AI agents become increasingly dominant market participants, the character of this belief formation process will change fundamentally. The framework developed here provides the theoretical tools to understand, predict, and design for this transition.

References

- [1] Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica*, 53(6), pp.1315–1335.
- [2] Glosten, L.R. and Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), pp.71–100.
- [3] Easley, D., Kiefer, N.M., O’Hara, M. and Paperman, J.B., 1996. Liquidity, information, and infrequently traded stocks. *Journal of Finance*, 51(4), pp.1405–1436.
- [4] Hasbrouck, J., 2007. *Empirical Market Microstructure*. Oxford University Press.
- [5] Bouchaud, J.P., Bonart, J., Donier, J. and Gould, M., 2018. *Trades, Quotes and Prices: Financial Markets Under the Microscope*. Cambridge University Press.
- [6] Cartea, Á., Jaimungal, S. and Penalva, J., 2015. *Algorithmic and High-Frequency Trading*. Cambridge University Press.
- [7] Grossman, S.J. and Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3), pp.393–408.
- [8] Hayek, F.A., 1945. The use of knowledge in society. *American Economic Review*, 35(4), pp.519–530.

- [9] Hornik, K., 1991. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2), pp.251–257.
- [10] Kidger, P. and Lyons, T., 2020. Universal approximation with deep narrow networks. *Proceedings of the 33rd COLT*, pp.2306–2327.
- [11] Olfati-Saber, R. and Murray, R.M., 2004. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9), pp.1520–1533.
- [12] Silver, D., Singh, S., Precup, D. and Sutton, R.S., 2021. Reward is enough. *Artificial Intelligence*, 299, p.103535.
- [13] Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), pp.223–236.
- [14] Bouchaud, J.P. and Potters, M., 2003. *Theory of Financial Risk and Derivative Pricing*. Cambridge University Press.
- [15] Amari, S., 2016. *Information Geometry and Its Applications*. Springer.
- [16] Friston, K., 2010. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), pp.127–138.
- [17] Long, M., 2026. Markets as distributed Bayesian inference engines. *GrokRxiv*, 2603.04.markets-bayesian-inference.
- [18] Long, M., 2026. Predictive market theory: a Bayesian-probabilistic foundation. *GrokRxiv*, 2603.04.predictive-market-theory.
- [19] Long, M., 2026. Unified predictive market framework: synthesis and implementation. *GrokRxiv*, 2603.04.unified-predictive-markets.