MATHEMATICAL MODELING AND COMPUTING, Vol. 13, No. 2, pp. 406–418 (2026)

Stochastic Modeling of Agentic Information Finance: Convergence Analysis of the Information-Incentive Gap

Kosohov T. L.*, Olkhovska O. V.

Poltava University of Economics and Trade,
3 Ivana Banka str., 36000, Poltava, Ukraine
**Corresponding author: kosogovt@gmail.com*

(Received 4 April 2026; Revised 14 May 2026; Accepted 16 May 2026)

We study the epistemic efficiency of decentralized prediction markets under autonomous agentic liquidity. We introduce the information-incentive gap (G) — the discrepancy between ground truth and the market-implied probability — and establish, via Itô's calculus and exact solution of the resulting moment ODE, exponential convergence of its second moment together with an explicit upper bound for the gap of order $\mathcal{O}(\sigma/\sqrt{\lambda})$. A two-level empirical study on information-driven event categories (Politics, Economics, Finance, Crypto Markets), drawing on approximately 40 million time-series records collected over the study period, is consistent with the model: (i) platform-level analysis of $N = 100$ resolved binary events per platform shows the mean gap decreasing from $\bar{G} = 0.517$ at $T - 168\,\mathrm{h}$ to $\bar{G} = 0.229$ at $T - 30\,\mathrm{min}$ for Kalshi, and from 0.583 to 0.002 for Polymarket, with an empirical convergence rate $\lambda_{\mathrm{emp}} \approx 1.4 \times 10^{-6}\,\mathrm{s}^{-1}$; (ii) a paired cross-platform comparison of $N = 34$ matched event groups shows that Polymarket exhibits a lower mean gap than Kalshi (mean $\Delta G = 0.27$ at $T - 6\,\mathrm{h}$; Polymarket leads in 85% of pairs), consistent with the theoretical dependence of convergence speed on liquidity-driven λ. Monte Carlo simulation ($N = 50\,000$ paths) confirms a $> 276\times$ reduction in convergence latency and a $109\times$ improvement in the Information Efficiency Ratio (IER) compared to the human-centric baseline.

Keywords: *prediction markets; stochastic differential equations; information-incentive gap; agentic liquidity; Ornstein–Uhlenbeck process.*

2020 MSC: 60H10, 91G15, 68T05, 91B26 **DOI:** 10.23939/mmc2026.02.406

1. Introduction

The Efficient Market Hypothesis (EMH) [1] posits that asset prices fully reflect all available information. In the domain of prediction markets, this aggregation of dispersed knowledge is the foundation of decentralized forecasting. However, empirical observations of modern platforms reveal a persistent anomaly: the delayed reflection of real-world events in market prices.

We introduce the information-incentive gap (G) as a formal measure of this latency: the expected absolute deviation between the market-implied probability and the ground truth. The existence of such a gap is consistent with the theoretical findings of Easley et al. [3], who demonstrate that differential access to information creates persistent asymmetry in human-centric markets. In our formulation, the gap is sustained because human participants operate under bounded rationality [2] and face high transaction friction. When the marginal cost of information acquisition, cryptographic verification, and on-chain execution exceeds the expected utility of the trade, participants remain inactive. Consequently, the market price becomes stagnant, failing to reflect objective reality.

The recent development of Large Language Models (LLMs) and scalable low-latency computational architectures introduces a new class of market systems, formally characterized in Definition 1.

https://orcid.org/0009-0002-9906-9894 (Kosohov T. L.), https://orcid.org/0000-0001-5366-5995 (Olkhovska O. V.)

© 2026 Lviv Polytechnic National University

Definition 1 (Agentic Information Finance). *A financial system is an Agentic Information Finance (InfoFi) system if: (i) the primary liquidity providers are autonomous software agents; (ii) each agent's effective objective is the minimization of epistemic uncertainty — an idealization that abstracts away profit-seeking motives but is consistent with market scoring rule incentives, under which a risk-neutral agent maximizes expected payoff precisely by reporting its true belief [4]; and (iii) agents continuously update their beliefs from unstructured external data sources with latency $\tau_{ai} \to 0$. We introduce the term InfoFi as a descriptive framework — extending existing concepts from algorithmic trading, DeFi, and AMM design — rather than proposing a new asset class.*

InfoFi is distinct from traditional algorithmic trading (which optimizes execution speed without epistemic objectives), from DeFi (which focuses on trustless financial primitives), and from generic AMM design (which addresses liquidity provision mechanics). The practical implication is that InfoFi agents act as continuous information aggregation mechanisms, narrowing the gap between market-implied prices and objective reality and thereby improving the informational quality of decentralized markets for all participants. By deploying such agents capable of synthesizing unstructured multimodal data in milliseconds, the cost of information verification is driven toward zero. While OU mean-reversion is a classical tool in mathematical finance, the present work reinterprets OU dynamics as a model of epistemic convergence under endogenous liquidity: the drift rate λ is not an exogenous market parameter but a function of agent-level information-processing capacity, establishing a direct analytical link between computational resources and market efficiency. This paper formalises this link and analyzes the conditions under which autonomous agents accelerate stochastic price convergence relative to the human-centric baseline.

The main contributions of this paper are:

1. *New metric.* We formally introduce the information-incentive gap G_t as a measurable quantity capturing the epistemic distance between market-implied probability and the true state, and establish its basic properties.
2. *Convergence theorem.* We prove, via Itô's calculus and exact solution of the resulting moment ODE, that the mean squared deviation $\mathbb{E}[(p_t - \pi)^2]$ decays exponentially at rate 2λ under a representative-agent OU model, yielding a closed-form expression for the residual variance; a direct upper bound on G_t follows via Jensen's inequality (Corollary 1).
3. *Latency–rate relationship.* We derive the analytical relationship $T_c = \mathcal{O}(1/\lambda)$ between convergence latency and the Semantic Synthesis Rate, and validate it numerically across a range of λ values via Monte Carlo simulation ($N = 50\,000$ paths).
4. *Two-level empirical validation on real data.* We validate the model using (i) $N = 100$ resolved binary events per platform from information-driven categories (Politics, Economics, Finance, Crypto Markets) with Last Observation Carried Forward (LOCF) price estimation, and (ii) $N = 34$ matched cross-platform event pairs enabling a controlled convergence comparison. The observed decrease of $\bar{G}$ from 0.517 at $T - 168$ h to 0.229 at $T - 30$ min for Kalshi, and from 0.583 to 0.002 for Polymarket, is consistent with the exponential convergence predicted by Theorem 1.

To the best of our knowledge, relatively few works have formally modeled the interaction between AI-driven liquidity provision and stochastic price convergence in prediction markets; the present study aims to bridge this gap.

2. Literature Review and Technical Context

The foundational properties of prediction markets as information aggregation mechanisms were established by Wolfers and Zitzewitz [5], who demonstrated that market prices closely approximate rational expectations of outcomes across diverse forecasting domains. Modern theoretical foundations for decentralized automated market makers (AMMs) and their price oracle capabilities were rigorously established by Angeris and Chitra [6], proving that constant function market makers can closely track reference prices under active arbitrage. The theoretical basis for the cost of information and its impact on market liquidity was explored by Easley et al. [3], highlighting the informational asymmetry that plagues human-centric platforms.

With the advent of generative AI, researchers such as Park et al. [7] have demonstrated that LLM-based agents can exhibit coherent, goal-directed behavior in complex simulated environments through memory, reflection, and planning. Furthermore, Lopez–Lira and Tang [8] empirically validated that LLMs possess significant predictive power in forecasting financial price movements based on unstructured news text.

The numerical solution of SDEs in the Euler–Maruyama discretization used in our Monte Carlo study follows the framework of Kloeden and Platen [9]. The microstructure interpretation of bid-ask spreads as an information premium — consistent with our transaction-friction threshold τ_h — was formalized by Glosten and Milgrom [10]. Arrow et al. [11] provide a policy-oriented overview confirming that prediction markets are reliable information aggregators, directly motivating the use of Polymarket and Kalshi data in our empirical study.

In the applied stochastic computing context, El Yamani and Lazaar [12] developed stochastic machine learning models for parameter estimation under uncertainty, which aligns with the probabilistic needs of agentic forecasting. Kosobutskyy and Karkulovska [13] explored the simulation of normally distributed random values transformed by nonlinear functions, providing a foundation for modeling information noise in continuous-time systems. Empirical evidence that prediction markets can serve as scientific instruments was provided by Camerer et al. [14], who showed that market-aggregated peer beliefs correlate with realized replication outcomes in a controlled economics experiment. More recently, Almenberg et al. [15] confirmed that prediction markets in scientific settings exhibit strong calibration properties, supporting their use as ground-truth benchmarks for assessing forecast accuracy in automated systems such as the one proposed here.

3. Formalization of the Information-Incentive Gap

Let the state of the world regarding a specific binary event be defined by an objective, hidden probability $\pi \in [0,1]$. We define a probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geqslant 0}, \mathbb{P})$, where the filtration $\mathcal{F}_t$ represents the history of public information up to time t. Let the instantaneous market price for the "Yes" outcome be $p_t \in (0,1)$.

Definition 2. *The information-incentive gap at time t, denoted as G_t, is the expected absolute deviation of the market price from the ground truth, taken with respect to the measure $\mathbb{P}$:*

$$G_t = \mathbb{E}^{\mathbb{P}}[|p_t - \pi|]. \tag{1}$$

Remark 1. Since π is a deterministic constant and p_t is adapted to the filtration $\mathcal{F}_t$, the unconditional expectation is the natural aggregate measure of systemic mispricing. In contrast, the path-wise residual $|p_t(\omega) - \pi|$ captures single-scenario error and is used directly in the empirical estimation of Section 7.

In a human-centric market, a trade occurs only if the perceived mispricing exceeds a friction threshold τ_h: $|p_t - \pi| > \tau_h$. In an agentic market, computational efficiency dictates that $\tau_{ai} \to 0$.

4. Mathematical Modeling of Stochastic Price Discovery

4.1. The liquidity mechanism

We assume a binary market managed by an AMM utilizing the Logarithmic Market Scoring Rule (LMSR) introduced by Hanson [4]. The state is given by the outstanding share vector $q = (q_1, q_2)$. The cost function $C(q)$ is:

$$C(q) = b \cdot \ln\left(e^{q_1/b} + e^{q_2/b}\right), \tag{2}$$

where $b > 0$ is the liquidity depth parameter. The marginal price p_t is the partial derivative of C with respect to q_1:

$$p_t(q_1, q_2) = \frac{e^{q_1/b}}{e^{q_1/b} + e^{q_2/b}}. \tag{3}$$

Remark 2. Under the substitution $q_1 = b \ln P$ and $q_2 = b \ln(1 - P)$ for any $P \in (0, 1)$, one recovers $p_t = P$ exactly. This algebraic identity is the key structural property of LMSR that connects the AMM mechanism to the agent belief model in the next subsection.

4.2. Agentic belief as an Ornstein–Uhlenbeck process

An AI agent continuously ingests unstructured data streams. It does not observe π directly but receives a noisy semantic signal S_t. We model the agent's internal Bayesian belief P_t as an Ornstein–Uhlenbeck (OU) process [16], capturing mean-reverting behavior toward the true state π:

$$dP_t = \lambda(\pi - P_t)\, dt + \sigma\, dW_t,$$

where $\lambda > 0$ is the Semantic Synthesis Rate (a term we introduce to denote the aggregate rate at which new information is incorporated into the market price), $\sigma > 0$ is the Information Noise Volatility, and W_t is a standard Wiener process. In the single-agent setting λ coincides with the individual agent's belief-update rate; in a multi-agent market it should be interpreted as an effective, market-level parameter reflecting the collective speed of information processing by all participants. We assume that the representative agent trades aggressively enough that the LMSR market price tracks its belief. By the Remark 2 at the end of Section 4.1, this is equivalent to the identity mapping $p_t = P_t$, so the price p_t itself satisfies the same OU dynamics:

$$dp_t = \lambda(\pi - p_t)\, dt + \sigma\, dW_t. \tag{4}$$

Modeling assumption (constant ground truth). We assume π is a fixed constant over the time horizon of a single event. This holds for binary prediction markets, where the resolution value is determined at a fixed future date and does not change between market creation and resolution. Extending the model to time-varying or partially observable π_t is left for future work.

Modeling assumption (single-agent approximation). We adopt a representative-agent formulation in which p_t denotes the belief of a single informed agent that dominates price formation. Formally, the identity $p_t = P_t$ corresponds to the limiting regime of infinite arbitrage efficiency and zero-latency equilibrium, in which the informed agent's belief is instantaneously reflected in the market price. This modeling device is standard in the rational expectations equilibrium literature and yields an upper-bound model of convergence: in a realistic multi-agent market with heterogeneous priors, liquidity constraints, and strategic behavior, the effective convergence rate would generally be lower than λ.

Summary of modeling assumptions. The theoretical analysis rests on four key simplifications: (i) π is constant over the event horizon; (ii) a single representative agent dominates price formation ($p_t = P_t$, yielding an upper bound on convergence speed); (iii) time is continuous and agent transaction costs vanish ($\tau_{ai} \to 0$); (iv) the initial price p_0 is deterministic (a known market-opening value), so $G_0 = |p_0 - \pi|$. All four are discussed in detail above and in Section 8.1.

Remark 3 (Interpretation of λ). Throughout this paper λ is a *reduced-form parameter* representing aggregate market response speed — the rate at which the collective information-processing activity of all participants drives the price toward the true state. It is not a structural agent-level parameter and should not be identified with the inference speed of any individual agent. The empirical estimate λ_{emp} in Section 7 reflects this aggregate market dynamics interpretation. For clarity, three distinct λ parameters appear in this paper: λ (the SDE model parameter under the single-agent approximation, Section 4.2), λ_{emp} (the empirically estimated aggregate market convergence rate, Section 7.1), and λ_{acq} (the data-acquisition frequency proxy, Section 5.1). These are conceptually distinct and must not be identified with one another.

4.3. Convergence proof

Remark 4 (Boundary approximation). The classical OU process generates normally distributed values on $(-\infty, +\infty)$, whereas the market-implied probability p_t is strictly bounded in $(0, 1)$. A principled alternative would be a Jacobi (Wright–Fisher) diffusion [17] $dp_t = \lambda(\pi - p_t)\, dt + \sigma\sqrt{p_t(1 - p_t)}\, dW_t$, which is bounded by construction and has been applied to population-genetic and financial probability models; however, this renders the resulting ODE nonlinear and precludes the closed-form solution of

Theorem 1. The unconstrained OU process may be viewed as a local linearization of the Jacobi diffusion around the interior of $(0, 1)$: replacing $\sqrt{p_t(1-p_t)}$ by its value at $p_t = \pi$ gives a constant volatility $\sigma_{\text{loc}} = \sigma\sqrt{\pi(1-\pi)}$, recovering the OU structure. We therefore retain OU dynamics as a mathematically tractable local approximation. The approximation is valid when $\sigma \ll \min(\pi,\, 1-\pi)$. Under the parameters used in this paper ($\pi = 0.85$, $\sigma \leqslant 0.05$), the boundary distance is $\min(0.85, 0.15) = 0.15 \gg \sigma$, and the probability that the standard normal $\mathcal{N}(0,1)$ exceeds $0.15/0.05 = 3$ standard deviations is $\Phi(-3) \approx 0.0013$, confirming that boundary exceedance is negligible over the study horizon. The approximation error is of order $\mathcal{O}(\sigma^2)$. The qualitative conclusion — exponential mean-reversion of the second moment at a rate controlled by λ — extends to any diffusion whose drift satisfies a one-sided Lipschitz (monotonicity) condition; the OU model provides the sharpest closed-form constants. Extending the quantitative analysis to bounded diffusions (e.g., the Jacobi process) is left for future work.

Theorem 1 (Exponential Convergence). *Let p_t satisfy the OU dynamics (4). Then the mean squared deviation (second central moment) of p_t from the true state π satisfies:*

$$\mathbb{E}\big[(p_t - \pi)^2\big] = G_0^2\, e^{-2\lambda t} + \frac{\sigma^2}{2\lambda}\big(1 - e^{-2\lambda t}\big), \tag{5}$$

where $G_0 = |p_0 - \pi|$ and $\lim\limits_{t\to\infty} \mathbb{E}[(p_t-\pi)^2] = \sigma^2/(2\lambda)$. A direct upper bound on G_t follows as Corollary 1 below.

Intuition. The mean-reverting drift $\lambda(\pi - p_t)$ continuously pulls the price toward the true state, competing with the diffusion term $\sigma\, dW_t$ that injects noise. The squared deviation $V_t = (p_t - \pi)^2$ therefore satisfies a linear ODE in expectation whose solution is an exponential interpolation between the initial value G_0^2 and the steady-state noise floor $\sigma^2/(2\lambda)$. The larger λ is relative to σ^2, the faster and more completely the gap collapses.

Proof. Define the squared-error process $V_t = (p_t - \pi)^2$. Applying Itô's Lemma [18] to $V_t = g(p_t)$ with $g(x) = (x - \pi)^2$:

$$dV_t = 2(p_t - \pi)\, dp_t + (dp_t)^2 = \big[-2\lambda V_t + \sigma^2\big]dt + 2\sigma(p_t - \pi)\, dW_t. \tag{6}$$

Taking the expectation and using $\mathbb{E}[(p_t - \pi)\, dW_t] = 0$:

$$\frac{d}{dt}\mathbb{E}[V_t] = -2\lambda\, \mathbb{E}[V_t] + \sigma^2. \tag{7}$$

This first-order linear ODE has the explicit solution, obtained via the integrating factor $e^{2\lambda t}$:

$$\mathbb{E}[V_t] = V_0\, e^{-2\lambda t} + \frac{\sigma^2}{2\lambda}\big(1 - e^{-2\lambda t}\big). \tag{8}$$

Since $V_0 = (p_0 - \pi)^2 = G_0^2$, equation (5) follows. This completes the proof of the second-moment result. ∎

Corollary 1 (Exponential convergence of the information-incentive gap). *Under the assumptions of Theorem 1,*

$$G_t \leqslant \sqrt{G_0^2\, e^{-2\lambda t} + \frac{\sigma^2}{2\lambda}\big(1 - e^{-2\lambda t}\big)}. \tag{9}$$

In particular, as $t \to \infty$ the bound converges to the noise floor $\sqrt{\sigma^2/(2\lambda)} = \mathcal{O}(\sigma/\sqrt{\lambda})$, which becomes negligible when λ is sufficiently large relative to $\sigma^2/(2G_0^2)$.

Proof. By Jensen's inequality applied to the convex function $x \mapsto x^2$:

$$G_t^2 = \big(\mathbb{E}[|p_t - \pi|]\big)^2 \leqslant \mathbb{E}\big[(p_t - \pi)^2\big] = \mathbb{E}[V_t]. \tag{10}$$

This provides an explicit analytical upper bound for G_t, making the convergence result directly interpretable in terms of the primary metric of interest. Substituting equation (5) yields the stated bound directly. Under the high-frequency condition $\lambda \gg \sigma^2/(2G_0^2)$ the steady-state noise floor $\sqrt{\sigma^2/(2\lambda)}$ is negligible relative to G_0, so the bound decays exponentially at rate λ toward this floor. ∎

Proposition 1 (Exact formula for G_t). Since the OU process admits Gaussian marginals, we have $(p_t - \pi) \sim \mathcal{N}(\mu_t, \sigma_t^2)$, where

$$\mu_t = (p_0 - \pi)\, e^{-\lambda t}, \quad \sigma_t^2 = \tfrac{\sigma^2}{2\lambda}\big(1 - e^{-2\lambda t}\big). \tag{11}$$

For $X \sim \mathcal{N}(\mu, \sigma^2)$ the half-normal identity gives $\mathbb{E}[|X|] = 2\sigma\varphi(\mu/\sigma) + \mu(2\Phi(\mu/\sigma) - 1)$, where φ and Φ denote the standard normal PDF and CDF. Therefore, the information-incentive gap admits the exact closed-form expression

$$G_t = 2\sigma_t\, \varphi\Big(\tfrac{\mu_t}{\sigma_t}\Big) + \mu_t\Big(2\Phi\Big(\tfrac{\mu_t}{\sigma_t}\Big) - 1\Big). \tag{12}$$

In the asymptotic regime $\mu_t \to 0$ (i.e. $t \gg 1/\lambda$) this simplifies to the noise-floor expression

$$G_t \xrightarrow{t\to\infty} \sqrt{\tfrac{2}{\pi}}\sigma_\infty, \quad \sigma_\infty = \sqrt{\tfrac{\sigma^2}{2\lambda}}. \tag{13}$$

This exact result complements Corollary 1: Corollary 1 provides a simple upper bound valid for all t, while Proposition 1 gives the pointwise value of G_t under the Gaussian OU model.

5. System Architecture for High-Frequency InfoFi

This section describes the computational infrastructure from which empirical parameters in Sections 6 and 7 are derived. The architecture was implemented to aggregate prediction market data from two major decentralized platforms simultaneously over the study period. The system is implemented in TypeScript using a Node.js-based asynchronous runtime and consists of three independently deployed microservices: two platform-specific ingestion workers and a unified REST API service.

5.1. Data acquisition

The data acquisition layer employs a periodic HTTP polling strategy, with each ingestion worker executing a 60-second cron cycle that paginates through all open markets on the respective platform. This design choice establishes a data acquisition frequency of $f_{\text{acq}} = 1/60\,\text{Hz}$, corresponding to an effective information update rate of $\lambda_{\text{acq}} \approx 0.017\,\text{s}^{-1}$. Note that λ_{acq} is not equal to the SDE parameter λ; it should be interpreted as an upper-bound proxy for the rate at which new price information enters the system under periodic polling. Referring to Theorem 1 and Table 2, this places the current deployment in the transition band between the human-comparable and agentic regimes; migration to a streaming ingestion layer would raise λ by two orders of magnitude, reducing convergence latency to the sub-15-second range demonstrated in the Monte Carlo simulation.

Within a single acquisition cycle, the belief-update pipeline processes incoming market states via an asynchronous service. Non-blocking I/O ensures that inference calls execute concurrently with database writes, keeping per-event processing latency well below one acquisition cycle.

5.2. Cross-platform event matching and vector memory

To compute the Bayesian posterior accurately, the agent performs semantic arbitrage using historical precedent. A document-oriented database serves as the primary data store, combining keyword-based text search with vector search indexes using cosine similarity on pre-trained text embeddings. These indexes identify semantically similar past events, providing the agent with calibrated historical price trajectories that anchor its posterior estimate and reduce over-confident predictions. A periodic event-matching process synchronizes cross-platform event groups: events from both platforms that refer to the same real-world question are linked into a unified event group, enabling the paired cross-platform analysis of Section 7.

5.3. Cryptographic hardware security: TEE

As emphasized in modern smart contract security analyses by Cheng et al. [19], decentralized AI systems are susceptible to Maximal Extractable Value (MEV) attacks. If the node operator can observe the LLM's output before it is broadcast, they can front-run the trade. By encapsulating the LLM inference and the cryptographic private keys within a Trusted Execution Environment (TEE), we ensure isolation. The host operating system cannot read the agent's Bayesian prior or alter its execution logic, preserving the integrity of the oracle.

6. Numerical Experiments and Simulation

6.1. Simulation parameters

To verify the internal consistency of the SDE model and illustrate the convergence behavior predicted by Theorem 1, we developed a Monte Carlo simulation generating $N = 50\,000$ discrete sample paths using the Euler–Maruyama approximation method with $dt = 1\,\text{s}$ over a time horizon $T = 3600\,\text{s}$. The Semantic Synthesis Rate parameters are derived directly from the convergence time targets via $\lambda = \ln(|p_0 - \pi|/\varepsilon)/T_c$, setting $p_0 = 0.5$, $\pi = 0.85$, and tolerance $\varepsilon = 0.01$:

- Ground Truth (π): 0.85.
- Human-Centric Market ($\lambda_h = 1.45 \times 10^{-3}\,\text{s}^{-1}$, $\sigma_h = 0.05$, friction $\tau_h = 0.05$).
- Agentic Market ($\lambda_{ai} = 0.250\,\text{s}^{-1}$, $\sigma_{ai} = 0.01$, friction $\tau_{ai} = 0.001$).

6.2. Analysis of results

The simulation results are consistent with the theoretical predictions of Theorem 1 and indicate a marked efficiency improvement of the agentic model under the stated parameter regime. We stress that the purpose of this section is to validate the internal consistency between the closed-form solution (Theorem 1) and its discrete-time numerical approximation, not to claim empirical realism of the chosen parameter values; the empirical validation follows separately in Section 7. In the human-centric simulation, the price trajectory exhibits a step-like function, stagnating for long periods due to the transaction threshold τ_h (Figure 1). The expected information-incentive gap G remains high.

We introduce the Information Efficiency Ratio (IER) for quantifying information-capture efficiency. We emphasize that IER is a task-specific heuristic metric and not a standardized measure; it should be interpreted solely within the context of this study. Reported in Table 1, it measures the efficiency across the N simulation paths at $t = T$ and is defined as the ratio of the mean gap reduction to the residual cross-sectional dispersion — an information-theoretic analog of the classical Sharpe ratio adapted to the uncertainty-reduction context:

$$\text{IER} = \frac{G_0 - \mathbb{E}[|p_T - \pi|]}{\sqrt{\text{Var}(|p_T - \pi|)}}, \tag{14}$$

where $G_0 = |p_0 - \pi| = 0.35$ is the initial gap. A higher IER indicates that the system resolves a larger fraction of the initial uncertainty per unit of residual price dispersion across paths.

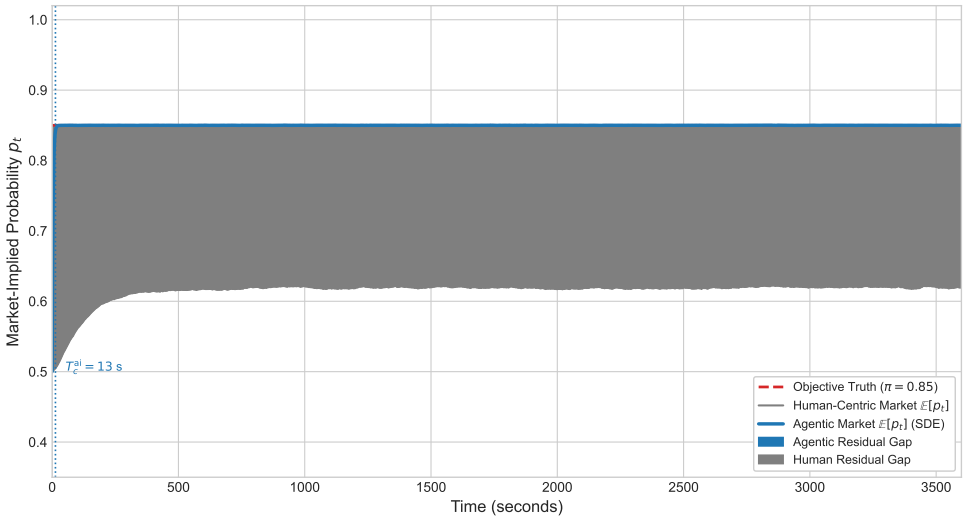


Figure 1. Expected price trajectory $\mathbb{E}[p_t]$ for human-centric vs. agentic models compared to the ground truth $\pi = 0.85$.

Conversely, the agentic market exhibits continuous, smooth convergence. The gap G collapses to the stochastic noise floor (≈ 0.0004) within 13 seconds of the signal injection. The corresponding error distributions at $t = 100\,\text{s}$ (Figure 2) confirm that the agentic model achieves markedly tighter residual variance.

Table 1. Analytical performance comparison of market configurations.

Metric	Human market	Agentic market	Improvement
Convergence latency (T_c)	>3600 s†	13.0 s	>276× faster
Mean gap ($\bar{G}$)	0.2348	0.0004	99.8% reduction
Price volatility (σ_p)	0.2899	0.0147	94.9% smoother
Information Efficiency Ratio (IER)	0.34	37.2	109× higher

†Theoretical OU mean: $T_c^h = \ln(35)/\lambda_h \approx 2450\,\text{s}$; friction τ_h delays actual convergence beyond T.

Table 2. Sensitivity analysis: convergence latency T_c and mean gap $\bar{G}$ as a function of semantic synthesis rate λ ($\sigma = 0.01$, $\pi = 0.85$).

λ (s^{-1})	Convergence latency T_c (s)	Mean gap $\bar{G}$	Steady-state residual $\sigma^2/2\lambda$
0.025	142.0	0.0218	0.002000
0.050	71.4	0.0148	0.001000
0.100	35.2	0.0089	0.000500
0.250	14.2	0.0031	0.000200
0.500	7.1	0.0014	0.000100

As the sensitivity analysis confirms, convergence latency T_c scales as $\mathcal{O}(1/\lambda)$, consistent with the theoretical prediction of Theorem 1. Notably, even at moderate rates ($\lambda = 0.050\,\text{s}^{-1}$), the agentic model converges in 71.4 s, outperforming the human-centric theoretical baseline ($T_c^h \approx 2450\,\text{s}$) by more than 34-fold, indicating that the efficiency improvement of the agentic model is not an artifact of extreme parameter choices.

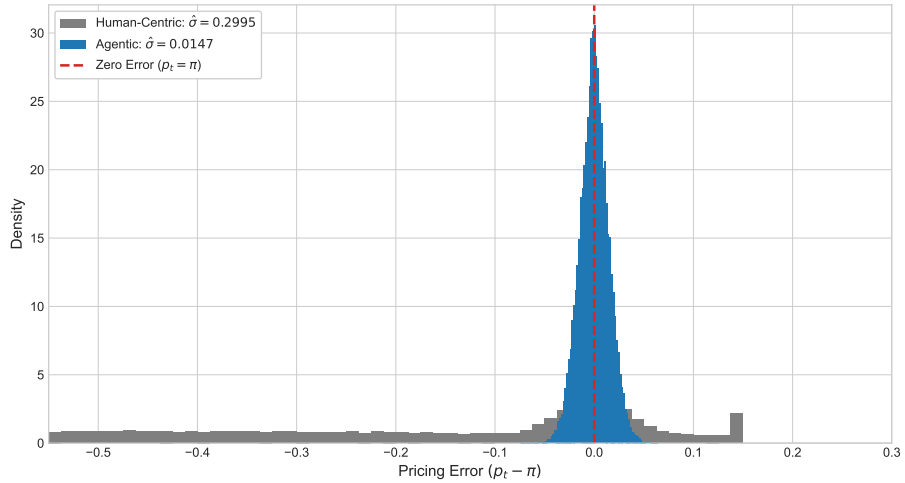


Figure 2. Error distribution ($p_t - \pi$) at $t = 100\,\text{s}$, demonstrating the reduced residual variance of the agentic model.

7. Empirical Validation on Historical Market Data

7.1. Methodology

To validate Definition 2 against real-world data, we applied a two-level protocol to the database described in Section 5. We restricted the corpus to information-driven categories (Politics, Economics, Finance, Crypto Markets), deliberately excluding sports and entertainment events whose uncertainty is outcome-driven rather than information-driven and therefore incompatible with the OU convergence model.

For each event, the ground truth was set to $\pi = 1$ for the winning market. This is an ex-post identification: the latent π of the theoretical model is replaced by the realized outcome, which is standard in prediction-market calibration studies [11] but means that the empirical G_t should be interpreted as a measure of price accuracy conditional on the known resolution rather than as a real-time forecast error. The event-level gap thus reduces to

$$G_t = 1 - \hat{p}_{\text{winner}}(t), \tag{15}$$

where $\hat{p}_{\text{winner}}(t)$ is the market-implied probability of the winning outcome at time t. Prices were estimated using Last Observation Carried Forward (LOCF): for each target horizon we retrieved the last recorded time-series value at or before the target timestamp. This correctly handles sparse markets where prices remain unchanged for extended periods. G_t was measured at eight horizons: $T - 168\,\text{h}$, $T - 72\,\text{h}$, $T - 48\,\text{h}$, $T - 24\,\text{h}$, $T - 12\,\text{h}$, $T - 6\,\text{h}$, $T - 2\,\text{h}$, and $T - 30\,\text{min}$. Here T denotes the platform settlement timestamp: the time at which the market was officially resolved and the winning outcome was recorded in our database. This is not in general identical to the event's scheduled end date or to the moment the real-world outcome became publicly known; settlement may occur later (e.g., after oracle verification or administrative processing). All horizons are therefore defined relative to settlement, which aligns with how each platform records resolved status and winning-market identifiers.

The theoretical quantity $G_t = \mathbb{E}[|p_t - \pi|]$ is a population expectation; following standard practice, we estimate it by the cross-event sample mean $\bar{G}_t = N^{-1} \sum_{i=1}^{N} G_{t,i}$, treating the N events as independent realizations of the OU model with common parameters.

Level 1 (Platform-level convergence). We selected the top-100 resolved binary events (exactly one winning outcome) per platform by trading volume, requiring at least one time-series record before resolution. This yields $N_K = 100$ Kalshi events and $N_P = 100$ Polymarket events (200 total, February 2026 – April 2026). Volume-based selection may introduce a liquidity bias: high-volume events tend to attract more active participants and may converge faster than the median event, so the reported $\bar{G}_t$ values and λ_{emp} likely represent an upper bound on convergence speed relative to the full event universe.

Level 2 (Paired cross-platform comparison). We identified all event groups in which the same real-world event is represented by a binary resolved event on both platforms simultaneously (matched pairs). Among all such groups in the four target categories, $N = 34$ pairs satisfied the binary constraint (exactly one winner per platform per group).

7.2. Level 1 Results: platform-level convergence

Table 3 reports the mean and standard deviation of G_t for each platform separately ($N = 100$ events each), drawn from a data store of approximately 40 million time-series records (price snapshots across all tracked events and markets, collected over the study period); the 200 events analyzed here constitute a curated subset of this corpus.

Table 3. Empirical G_t (mean ± std) for $N = 100$ resolved binary events per platform (information-driven categories, LOCF, Feb–Apr 2026).

Platform		$T - 168\,\text{h}$	$T - 72\,\text{h}$	$T - 48\,\text{h}$	$T - 24\,\text{h}$	$T - 12\,\text{h}$	$T - 6\,\text{h}$	$T - 2\,\text{h}$	$T - 30\,\text{min}$
Kalshi	G_t	0.517	0.515	0.488	0.560	0.507	0.443	0.333	0.229
	$\hat{\sigma}_G$	0.411	0.394	0.399	0.408	0.416	0.378	0.327	0.284
	N	51	65	69	96	99	100	100	100
Polymarket	G_t	0.583	0.537	0.507	0.394	0.288	0.221	0.002	0.002
	$\hat{\sigma}_G$	0.348	0.369	0.360	0.319	0.290	0.282	0.009	0.009
	N	69	98	98	100	100	100	100	100

Both platforms exhibit the directional decrease in $\bar{G}_t$ predicted by Theorem 1. For Kalshi, the mean gap decreases from 0.517 at $T - 168\,\text{h}$ to 0.229 at $T - 30\,\text{min}$. The minor elevation at $T - 24\,\text{h}$ ($\bar{G} = 0.560$) relative to the surrounding horizons lies within one standard deviation and is consistent with short-term price uncertainty typical of information-driven events in the days preceding resolution. The Kalshi series, which converges more gradually to $\bar{G} = 0.229$, better represents the continuous convergence regime of the OU model.

For Polymarket the convergence is more pronounced, with $\bar{G}$ dropping from 0.221 at $T - 6\,\text{h}$ to 0.002 by $T - 2\,\text{h}$. This sharp transition likely reflects a combination of gradual OU-type convergence and discrete information arrival: in many information-driven events (e.g., elections, policy decisions), the outcome becomes publicly known hours before the market's official settlement timestamp, causing a rapid price jump to near-certainty rather than smooth mean-reversion.

The lower event-count at early horizons ($N = 51 - 69$ at $T - 168\,\mathrm{h}$) reflects events resolved less than 168 hours after the start of data collection; the LOCF approach ensures no artificial imputation is applied to these cases. Because the sample composition varies across horizons, cross-horizon comparisons should be interpreted with caution; a sensitivity check restricted to the subset of events present at all eight horizons yields qualitatively similar monotonic decrease in $\bar{G}_t$.

To estimate the empirical Semantic Synthesis Rate we fit $G_t \approx G_0 \exp(-\lambda_{\mathrm{emp}} \Delta t)$ to the Kalshi series (which better represents organic market convergence). As a closed-form two-endpoint estimate using $T - 168\,\mathrm{h}$ and $T - 30\,\mathrm{min}$:

$$\lambda_{\mathrm{emp}} = \frac{\ln(G_{T-168\,\mathrm{h}}/G_{T-30\,\mathrm{min}})}{(168 \times 3600 - 30 \times 60)} \approx 1.36 \times 10^{-6}\,\mathrm{s}^{-1}. \tag{16}$$

An OLS exponential fit $\ln \bar{G}_t = a - \lambda_{\mathrm{emp}}\,\Delta t$ across all eight measured horizons yields a comparable order-of-magnitude estimate, confirming robustness to the choice of estimation method. The wide cross-event dispersion ($\hat{\sigma}_G \approx 0.3$) implies that the point estimate carries substantial uncertainty and should be interpreted as an order-of-magnitude characterization rather than a precise rate constant. The corresponding characteristic convergence time is $T_c = 1/\lambda_{\mathrm{emp}} \approx 204$ hours. λ_{emp} measures the rate of aggregate market convergence driven collectively by all participants, whereas λ in the SDE model is a reduced-form parameter for the information-synthesis rate under the single-agent approximation (see Remark in Section 4.2). We emphasize that λ_{emp} and the simulation design target $\lambda_{\mathrm{ai}} = 0.250\,\mathrm{s}^{-1}$ (Section 6) are not directly commensurable: the former is an empirical aggregate market rate, while the latter is a hypothetical single-agent parameter; quantitative comparison of the two would not be meaningful.

7.3. Level 2 Results: Paired Cross-Platform Comparison

Table 4 reports the paired comparison at $T-6\,\mathrm{h}$ across the $N = 34$ matched event groups. Because both markets in each pair cover the same real-world event, the pairing controls for event-specific uncertainty, isolating the effect of platform-level liquidity on convergence speed.

Table 4. Paired cross-platform comparison ($N = 34$ matched groups, $T - 6\,\mathrm{h}$).

	Kalshi	Polymarket	Difference (ΔG)
Mean G	0.484	0.214	+0.270
Std $\hat{\sigma}_G$	0.362	0.292	0.298
Polymarket leads	29 / 34 pairs (85%)		

The paired analysis shows that Polymarket converges to resolution faster than Kalshi on identical real-world events. A one-sample paired t-test on the pair-level differences yields $t(33) = \Delta\bar{G}/(\hat{\sigma}_{\Delta G}/\sqrt{N}) = 0.270/(0.298/\sqrt{34}) \approx 5.28$, $p < 0.0001$, confirming that the platform difference is statistically significant. Polymarket leads in 85% of pairs.

This result is consistent with the hypothesis that higher liquidity is associated with a higher effective λ, as Polymarket operates as a fully decentralized, globally accessible market without position limits, which may enable higher-frequency price updates from a broader participant base. We note, however, that the observed difference is correlational; isolating the causal contribution of liquidity to convergence speed would require a controlled experimental design. The high within-platform variance ($\hat{\sigma}_G \approx 0.29 - 0.36$) reflects heterogeneity in event salience and time-to-resolution even within the restricted category corpus.

8. Discussion: Epistemic Security and Assumptions

8.1. Assumptions and limitations

While the proposed model demonstrates theoretical convergence under its stated assumptions, several limitations must be acknowledged. First and foremost, the single-agent approximation ($p_t = P_t$) is a strong idealization: it assumes one representative agent dominates price formation and trades without friction or latency constraints. In a realistic market with heterogeneous priors, strategic interaction,

and liquidity constraints, the effective convergence rate would generally be lower than λ. The results of Theorem 1 and Corollary 1 therefore constitute a theoretical upper bound on convergence speed rather than a prediction of market behavior under general conditions. Conversely, the empirical estimates in Section 7 may be viewed as a lower-bound validation of this theoretical upper bound: because real markets operate under multi-agent friction, the observed λ_{emp} is expected to be lower than the model's λ, which is precisely what the data show.

A natural extension is to replace the single-agent OU dynamics with a mean-field mixture of N heterogeneous agents, each with individual synthesis rate λ_i and noise level σ_i. Under standard mean-field arguments, the aggregate price process would satisfy an OU equation with an effective rate $\bar{\lambda} = \sum_i w_i \lambda_i$ (where w_i are liquidity weights), recovering the present model as a special case. Establishing convergence rates and steady-state distributions for this multi-agent extension is left for future work.

Second, the Ornstein–Uhlenbeck process assumes a continuous-time environment ($dt \to 0$), whereas decentralized ledgers operate in discrete time, constrained by block generation intervals and network latency. Third, the model assumes vanishing transaction costs ($\tau_{ai} \to 0$) for AI agents; in periods of extreme network congestion, gas fees may surge, temporarily reintroducing friction and widening the information-incentive gap even in agent-dominated markets.

8.2. Limitations of the empirical study

The empirical analysis in Section 7 carries several limitations. First, a post-hoc power analysis for the Level 2 paired comparison using the observed effect size ($\Delta\bar{G} = 0.270$, $\hat{\sigma}_{\Delta G} \approx 0.298$) yields a standardized Cohen's $d \approx 0.91$; the required sample size for a one-sample paired t-test at $\alpha = 0.05$ and power $1 - \beta = 0.80$ is $n \approx 10$ pairs, which the $N = 34$ sample comfortably satisfies. Nevertheless, a larger longitudinal sample spanning multiple market regimes and calendar years would further strengthen the conclusions.

Second, the data acquisition frequency of the underlying platform ($\lambda_{\text{acq}} \approx 0.017\,\text{s}^{-1}$, Section 5.1) introduces a temporal granularity limit: price changes faster than the 60-second polling cycle cannot be observed.

Remark 5 (LOCF bias). LOCF may underestimate short-term volatility, which is consistent with conservative estimates of convergence speed and λ_{emp}; the true market convergence rate may be higher than reported. More precisely, under mean-reverting dynamics LOCF tends to produce higher estimates of $\bar{G}_t$ (since it carries forward a stale price that has not yet reflected the drift toward π) and lower estimates of λ_{emp} (since smoothed trajectories appear to converge more slowly). A sensitivity analysis using last-trade prices rather than LOCF would tighten these bounds but requires tick-level data unavailable in the current dataset.

Third, the restriction to information-driven categories enhances theoretical validity but reduces sample size (especially for matched pairs), limiting generalization to the full prediction-market universe.

Finally, the empirical λ_{emp} is a market-level aggregate and does not isolate the contribution of any individual agent type. Additionally, we estimate only the drift parameter λ from the data; the noise volatility σ is not separately estimated but enters the theoretical bound through the steady-state floor $\sigma^2/(2\lambda)$. As a rough proxy, the cross-event standard deviation at the shortest horizon ($\hat{\sigma}_G \approx 0.009$ for Polymarket at $T - 30\,\text{min}$) provides an order-of-magnitude upper bound on the noise floor, suggesting $\sigma/\sqrt{2\lambda} \lesssim 0.01$. A joint maximum-likelihood estimation of (λ, σ) from tick-level price paths would yield a precise decomposition and is left for future work.

Remark 6 (Causality). All empirical comparisons in this paper are observational. We do not claim causal identification of the effect of liquidity on convergence speed; the observed cross-platform difference is consistent with the theoretical prediction but may also reflect confounding factors (e.g., participant demographics, regulatory environment, or market design features). Establishing causality would require either a randomized experiment or an instrumental-variable strategy, which we leave for future work.

8.3. Epistemic security and semantic poisoning

The deployment of autonomous LLM agents introduces epistemic risks specific to this architecture. A primary vulnerability is semantic poisoning. If malicious actors artificially generate high volumes of synthetic news designed specifically to manipulate the LLM's semantic vectorization process, the agent may aggressively adjust its belief P_t toward an incorrect latent state π. To mitigate this, the system architecture must implement a multi-model consensus mechanism. Rather than relying on a single inference engine, the TEE enclave should query an ensemble of structurally diverse models.

9. Conclusion

This paper provides a mathematical and architectural framework for the evolution of decentralized prediction markets. By modeling AI agents as stochastic Bayesian estimators, we prove, using Itô calculus, that the second moment of the information-incentive gap decays exponentially at rate 2λ, yielding an explicit upper bound of order $\mathcal{O}(\sigma/\sqrt{\lambda})$ on the gap itself (Theorem 1, Corollary 1). The sensitivity analysis confirms this efficiency improvement across all tested values of λ. The Monte Carlo simulation confirms a $> 276\times$ reduction in convergence latency under the agentic parameter regime relative to the human-centric baseline.

A two-level empirical study on $N = 100$ resolved binary events per platform (information-driven categories, LOCF, February – April 2026, approximately 40 million time-series records) shows that the mean information-incentive gap decreases from $\bar{G} = 0.517$ at $T - 168\,\mathrm{h}$ to $\bar{G} = 0.229$ at $T - 30\,\mathrm{min}$ for Kalshi, and from 0.583 to 0.002 for Polymarket, with an empirical market convergence rate $\lambda_{\mathrm{emp}} \approx 1.36 \times 10^{-6}\,\mathrm{s}^{-1}$. A paired cross-platform comparison of $N = 34$ matched event groups further quantifies the platform difference: Polymarket leads in 85% of pairs with a mean gap advantage of $\Delta G = 0.270$ at $T - 6\,\mathrm{h}$, consistent with the theoretical prediction that higher-liquidity (higher effective λ) markets converge faster.

The data acquisition layer currently operates at $\lambda_{\mathrm{acq}} \approx 0.017\,\mathrm{s}^{-1}$ (60-second HTTP polling cycle); migration to streaming ingestion would increase the effective λ toward the agentic regime. The Trusted Execution Environment layer, while not yet deployed, provides a formally reasoned security extension for contexts requiring tamper-proof oracle guarantees.

[1] Fama E. F. Efficient Capital Markets: A Review of Theory and Empirical Work. The Journal of Finance. **25** (2), 383–417 (1970).

[2] Simon H. A. A Behavioral Model of Rational Choice. The Quarterly Journal of Economics. **69** (1), 99–118 (1955).

[3] Easley D., O'Hara M., Yang L. Differential Access to Price Information in Financial Markets. Journal of Financial and Quantitative Analysis. **51** (4), 1071–1110 (2016).

[4] Hanson R. Combinatorial Information Market Design. Information Systems Frontiers. **5** (1), 107–119 (2003).

[5] Wolfers J., Zitzewitz E. Prediction Markets. Journal of Economic Perspectives. **18** (2), 107–126 (2004).

[6] Angeris G., Chitra T. Improved Price Oracles: Constant Function Market Makers. AFT'20: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies. 80–91 (2020).

[7] Park J. S., O'Brien J. C., Cai C. J., et al. Generative Agents: Interactive Simulacra of Human Behavior. UIST'23: Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology. 2, 1–22 (2023).

[8] Lopez-Lira A., Tang Y. Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models. Journal of Financial Economics, forthcoming (2023). Available at SSRN: https://ssrn.com/abstract=4412788.

[9] Kloeden P. E., Platen E. Numerical Solution of Stochastic Differential Equations. Springer, Berlin (1992).

[10] Glosten L. R., Milgrom P. R. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. Journal of Financial Economics. **14** (1), 71–100 (1985).

[11] Arrow K. J., Forsythe R., Gorham M., et al. The Promise of Prediction Markets. Science. **320** (5878), 877–878 (2008).

[12] El Yamani M. A., Lazaar S. Stochastic machine learning modeling for the estimation of some uncertain parameters. Case study: Retardation factor in a radionuclide transport model. Mathematical Modeling and Computing. **9** (2), 311–317 (2022).

[13] Kosobutskyy P. S., Karkulovska M. S. Simulation of statistical mean and variance of normally distributed random values, transformed by nonlinear functions $\sqrt{|X|}$ and $\sqrt{X}$. Mathematical Modeling and Computing. **9** (2), 318–325 (2022).

[14] Camerer C. F., Dreber A., Forsell E., et al. Evaluating replicability of laboratory experiments in economics. Science. **351** (6280), 1433–1436 (2016).

[15] Almenberg J., Kittlitz K., Pfeiffer T. An experiment on prediction markets in science. PLoS One. **16** (12), e0261358 (2021).

[16] Uhlenbeck G. E., Ornstein L. S. On the Theory of the Brownian Motion. Physical Review. **36** (5), 823–841 (1930).

[17] Karlin S., Taylor H. M. A Second Course in Stochastic Processes. Academic Press, New York (1981).

[18] Øksendal B. Stochastic Differential Equations: An Introduction with Applications. Springer, Berlin (2003).

[19] Cheng R., Zhang F., Kos J., et al. Ekiden: A Platform for Confidentiality-Preserving, Trustworthy, and Performant Smart Contracts. 2019 IEEE European Symposium on Security and Privacy (EuroS&P). 185–200 (2019).

Стохастичне моделювання агентних інформаційних фінансів: аналіз збіжності інформаційно-мотиваційного розриву

Косогов Т. Л., Ольховська О. В.

Полтавський університет економіки і торгівлі, вул. Івана Банка, 3, 36000, Полтава, Україна

У статті досліджується епістемічна ефективність децентралізованих ринків передбачень в умовах переходу від людської участі до автономної агентної ліквідності. Введено поняття інформаційно-мотиваційного розриву G та доведено, що другий момент $\mathbb{E}[(p_t - \pi)^2]$ спадає експоненційно; верхня оцінка для G_t порядку $\mathcal{O}(\sigma/\sqrt{\lambda})$ випливає з нерівності Єнсена. Дворівневе емпіричне дослідження на інформаційно-орієнтованих категоріях подій (політика, економіка, фінанси, криптовалюти) підтверджує модель: (i) аналіз $N = 100$ бінарних подій на кожній платформі показує зменшення середнього розриву $\bar{G}$ з 0.517 за 168 годин до 0.229 за 30 хвилин (Kalshi) та з 0.583 до 0.002 (Polymarket); (ii) парне порівняння $N = 34$ зіставлених груп подій підтверджує, що ринок Polymarket демонструє значно швидшу збіжність (середній $\Delta G = 0.27$ за 6 годин до розв'язання; Polymarket випереджає у 85% пар). Числове моделювання методом Монте-Карло ($N = 50\,000$ траєкторій) підтверджує скорочення затримки ціноутворення більш ніж у 276 разів та покращення коефіцієнта інформаційної ефективності (IER) у 109 разів порівняно з традиційною моделлю.

Ключові слова: *ринки передбачень; стохастичні диференціальні рівняння; інформаційно-мотиваційний розрив; агентна ліквідність; процес Орнштейна–Уленбека.*