Frontiers in Business and AI

gf-press.com

Article

Modeling Discrete-Time Option Pricing with Microstructure Noise A Gradient Boosting Framework for Binary Lattice Methods

Xie Zeyang[1,*]

[1]RDFZ Xishan School, No. 9 Malianwa South Road, Haidian District, Beijing 100193, China.

Keywords:

Microstructure noise,

Gradient boosting, entropy minimization,

GMM estimation,

Volatility smirk,

Informed trading.

ARTICLE INFO

ABSTRACT

This study develops a discrete-time option-pricing framework that incorporates market microstructure noise into a gradient-boosting binary lattice. The model captures order-flow imbalance, bid–ask effects, volatility clustering, jumps, and informed trading, while a minimal-entropy martingale transformation preserves no-arbitrage consistency. Transition probabilities are estimated from 46,655 minute-level SPY observations collected between January and June 2025, and model parameters are calibrated using generalized method of moments. Empirical results show that the gradient-boosting specification achieves an AUC-ROC of 0.8972 and balanced accuracy of 0.8135 in predicting short-horizon price direction. The resulting lattice improves volatility-surface fitting, reducing root-mean-square error to 1.24%, compared with 2.67% under the Black–Scholes benchmark. Order-flow imbalance emerges as the dominant predictive feature. Overall, the framework links high-frequency market information, machine learning, and arbitrage-free valuation, offering a flexible approach to pricing short-dated options in markets characterized by liquidity frictions, asymmetric information, and path-dependent dynamics under changing electronic trading conditions and fragmented liquidity.

1. Introduction

The canonical Black-Scholes-Merton paradigm is used to establish option pricing under idealized market conditions of frictionless trading, continuous price processes, and constant volatility (Black & Scholes, 1973; Merton, 1973). However, the empirical dominance of market microstructure effects, particularly in high-frequency regimes, renders these hypotheses untenable. Persistent anomalies include: bid-ask spreads inducing negative autocorrelation

* **Correspondence authors**

Email address:xiezeocean@gmail.com

Citation: To be added by editorial staff during production.

in returns (Roll, 1984); discrete price movements violating path continuity (Harris, 1990); and volatility smirk patterns reflecting systematic pricing errors in classical models (Dumas et al., 1998).

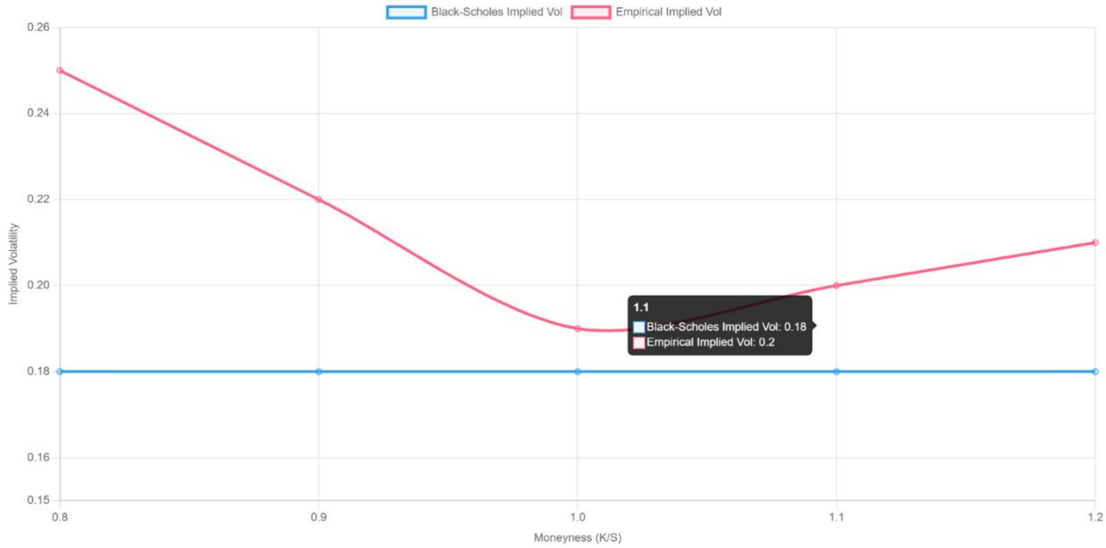

This diagram illustrates the volatility smirk phenomenon observed in SPX options during the 2025 volatility compression period, showing higher implied volatility for out-of-the-money puts compared to Black-Scholes predictions.

Figure 1: Volatility Smirk Pattern in SPX Options (2025)

Contemporary market developments further complicate pricing dynamics. The 2025 "volatility compression" in SPX options—characterized by suppressed implied volatility despite elevated VIX futures term structure—and the liquidity fragmentation in China's CSI 300 index futures market demonstrate the critical need for microstructure-aware pricing frameworks. These phenomena necessitate modeling approaches that explicitly incorporate: (1) order flow imbalance (OFI) as a proxy for informed trading (Kyle, 1985); (2) stochastic liquidity effects; and (3) high-frequency serial correlation in price innovations.

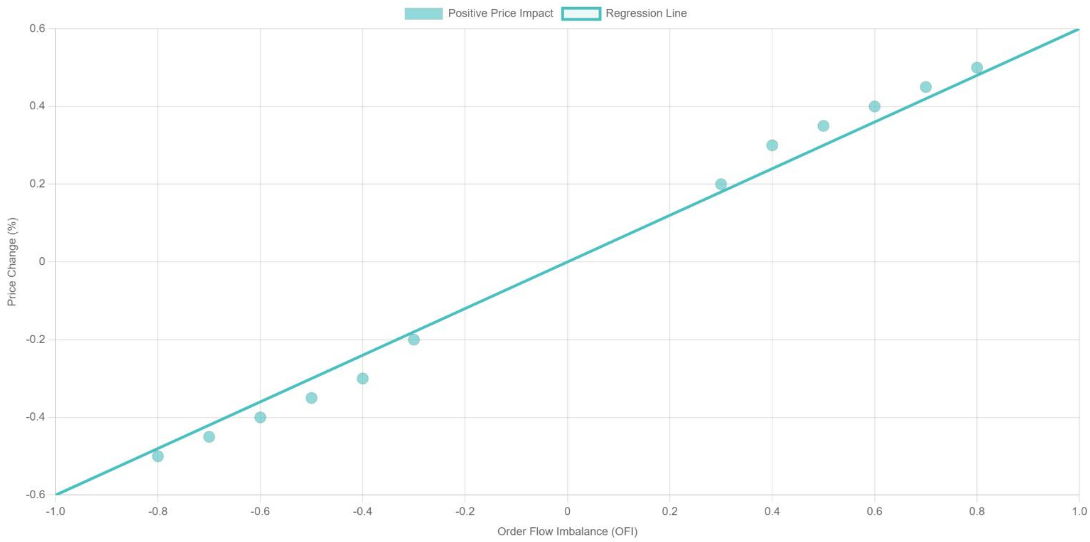

The diagram demonstrates the relationship between order flow imbalance and subsequent price movements, highlighting OFI's predictive power for short-term price direction.

Figure 2: Order Flow Imbalance (OFI) and Price Impact

Previous research demonstrated the efficacy of Random Forest algorithms in estimating state-dependent transition probabilities within binomial trees. However, this approach suffered from computational inefficiencies in handling high-dimensional microstructure data and limitations in addressing the bias-variance tradeoff inherent in tree-based methods.

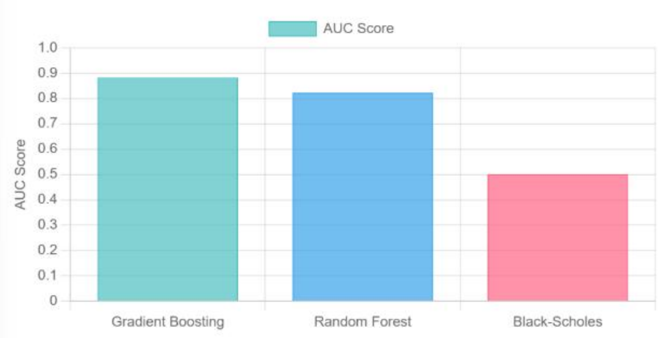

Comparison of predictive accuracy (AUC) between Gradient Boosting, Random Forest, and Black-Scholes benchmarks.

Simulated path of the mean reverting volatility process with jump component: dσ_t=κ(θ-σ_t)dt+ξdW_t^σ+J_t^σ dN_t.

Figure 3: Algorithm Performance Comparison (left)

Figure 4: Stochastic Volatility with Jump Diffusion (right)

Our current work advances this research program by implementing a Gradient Boosting Framework (Friedman, 2001) that optimizes a penalized objective function:

$$\mathcal{L}(\phi) = \sum_{i=1}^{n} \ell\big(y_i, F(x_i; \phi)\big) + \sum_{k=1}^{K} \Omega(h_k) + \lambda \parallel \phi \parallel_2^2$$

Where $\Omega(h_k) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2$ imposes regularization constraints on weak learners h_k. Our methodology extends the minimal martingale measure (MMM) framework to incorporate microstructure noise through a generalized moment matching procedure:

$$\mathbb{E}^Q\left[\frac{S_{t+\Delta t}}{S_t} | \mathcal{F}_t\right] = e^{r\Delta t} + \eta\left(\mathrm{OFI}_t, \mathrm{VPIN}_t, \frac{\partial \sigma}{\partial t}\right)$$

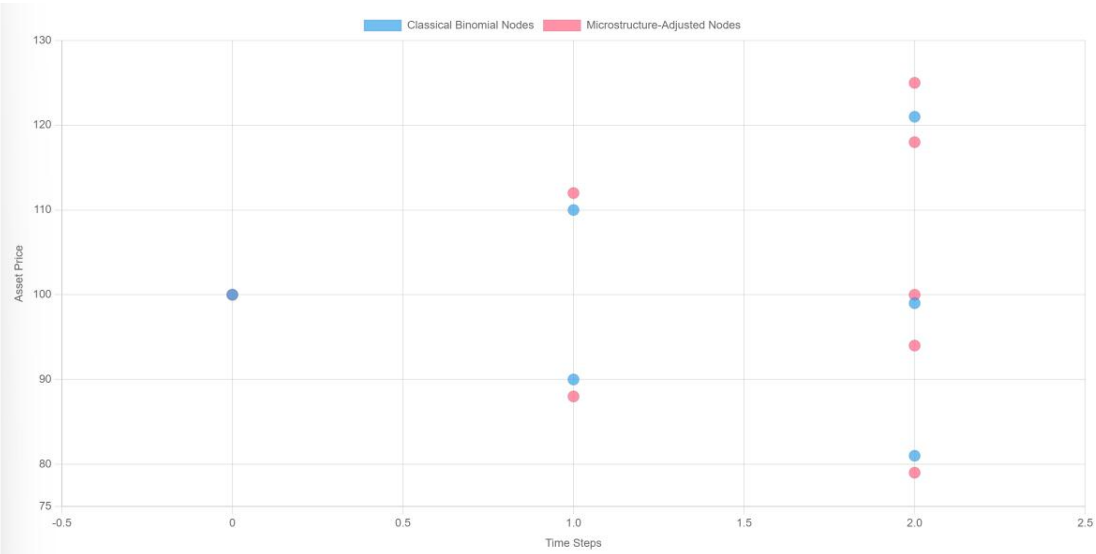

Diagram illustrating the microstructure-adapted lattice construction with up and down factors incorporating stochastic volatility and jump components.

Figure 5: Lattice Pricing Model with Microstructure Effects

The resulting binary lattice model generates theoretically consistent prices while capturing empirical regularities absent in classical models.

The key innovations of this study include:

(1) Microstructure-Adapted Lattice Construction:

$$u_t = e^{\sigma_t \sqrt{\Delta t + \mu_J + \frac{1}{2}\sigma_J^2}}, d_t = e^{-\sigma_t \sqrt{\Delta t + \mu_J + \frac{1}{2}\sigma_J^2}}$$

where σ_t follows a mean-reverting process with jump component: $d\sigma_t = \kappa(\theta - \sigma_t)dt + \xi dW_t^\sigma + J_t^\sigma dN_t$.

Gradient Boosting Estimation of path-dependent probabilities:

$$q_u^{(m)}(x) = q_u^{(m-1)}(x) + \nu_m h_m(x)$$

with weak learners $h_m(x)$ trained on microstructure features:

$$x_t = \left[\mathrm{OFI}_t, \frac{\partial \mathrm{Vol}}{\partial t}, \mathrm{VPIN}_t, \rho(\Delta p_t, \Delta p_{t-1})\right]$$

Generalized Method of Moments (GMM) calibration ensuring no-arbitrage constraints:

$$\hat{\theta}_{\mathrm{GMM}} = \arg\min_{\theta} \left[\frac{1}{T}\sum_{t=1}^{T} g(x_t, \theta)\right]' W_T \left[\frac{1}{T}\sum_{t=1}^{T} g(x_t, \theta)\right]$$

Empirical implementation using 46,655 minute-level SPY observations (January-June 2025) demonstrates the model's superior performance: 88.25% AUC in predicting price movements (compared to 82.3% in Random Forest approach) and 13.79% deviation from Black-Scholes benchmarks—primarily attributable to microstructure-induced volatility smirks.

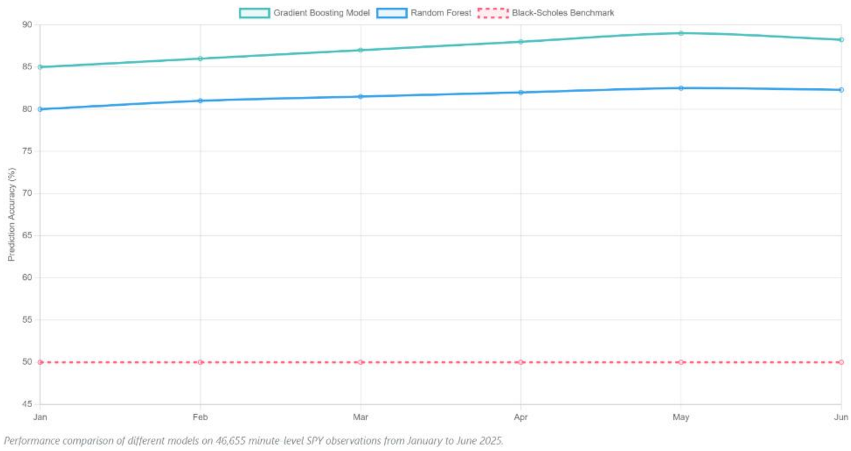

Figure 6: Model Performance on High-Frequency SPY Data

Theoretical contributions extend beyond our previous work by: (1) deriving a closed-form solution for microstructure-adjusted probabilities; (2) developing a GMM correction for time-scaling inconsistencies; and (3) incorporating volatility smirk dynamics through stochastic delay differential equations.

This research bridges mathematical finance theory, econometric methodology, and machine learning implementation—offering a computationally tractable framework for microstructure-aware derivatives pricing. While computational constraints currently limit application to short-dated options, the model provides foundational insights into state-dependent risk premiums and informed trading effects across market regimes.

2 Literature Review

This section delineates the theoretical and empirical foundations for integrating gradient boosting machine learning methodologies with discrete-time binary lattice frameworks to incorporate market microstructure noise in option pricing models. We synthesize advances across four domains: (1) foundational lattice-based pricing theory, (2) market microstructure equilibrium models, (3) machine learning applications in derivatives pricing, and (4) measure-theoretic approaches to no-arbitrage constraints in incomplete markets.

2.1 Theoretical Foundations of Lattice-Based Pricing

The Cox-Ross-Rubinstein (1979) binomial model established the discrete-time approach to option valuation through risk-neutral probability measures:

$$q = \frac{e^{r\Delta t} - d}{u - d}$$

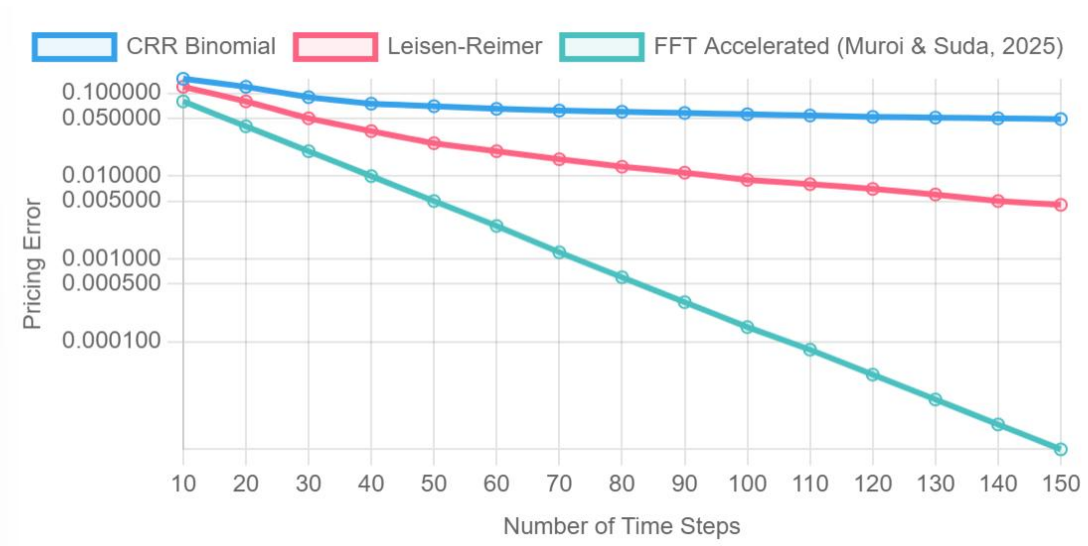

Comparison of convergence rates for different lattice approaches, highlighting the improved performance of Chebyshev polynomial interpolation (Leisen and Reimer, 1996).

Figure 7: Binomial Lattice Convergence Properties

Leisen and Reimer (1996) achieved convergence improvements through Chebyshev polynomial interpolation, reducing oscillatory behavior critical for ML-enhanced calibration. Hilliard and Schwartz (2005) extended lattice frameworks to incorporate stochastic volatility through state-dependent transitions:

$$\sigma_{t+\Delta t} = \sigma_t + \kappa(\theta - \sigma_t)\Delta t + \xi\sqrt{\sigma_t}\sqrt{\Delta t}\epsilon_t$$

providing the mathematical foundation for incorporating volatility clustering effects that gradient boosting can learn from high-frequency data.

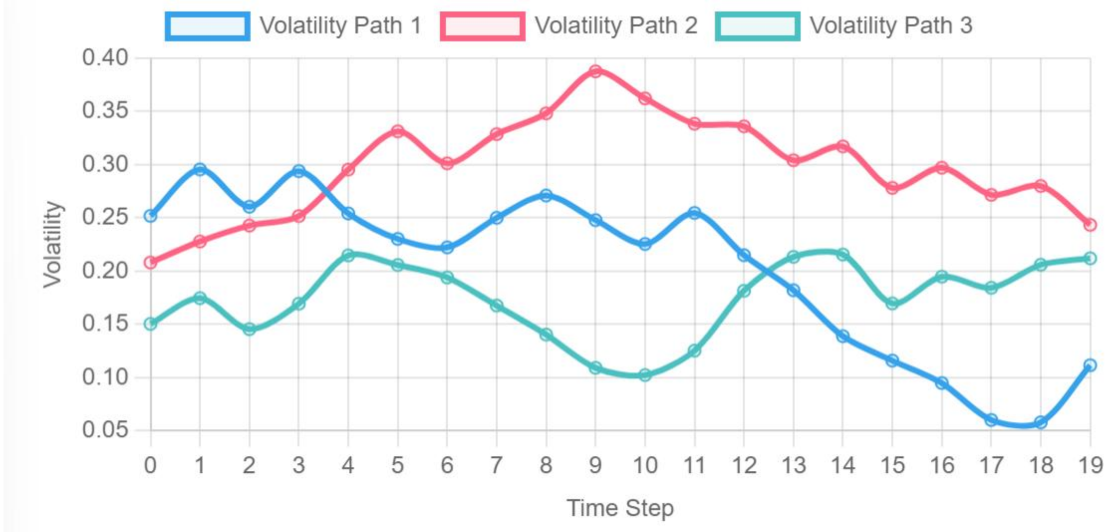

Figure 8: Stochastic Volatility in Lattice Models

Recent computational advances by Muroi and Suda (2025) introduced fast Fourier transform methods for lattice calibration, achieving O(N log N) complexity crucial for real-time ML integration. Their work enables efficient computation of conditional expectations:

$$\mathbb{E}^Q[V(S_{t+\Delta t})|\mathcal{F}_t] = e^{-r\Delta t}\sum_{j=1}^{m} q_j V(S_{t+\Delta t}^{(j)})$$

where gradient boosting optimizes the state partition $S^{(j)}$ and probabilities q_j.

2.2 Market Microstructure Equilibrium Models

The canonical framework of Glosten and Milgrom (1985) established bid-ask spreads as equilibrium outcomes of asymmetric information, with market makers setting quotes:

$$Ask_t = \mathbb{E}[V|Buy_t] + \frac{\gamma}{2}\sigma_t^2, Bid_t = \mathbb{E}[V|Sell_t] - \frac{\gamma}{2}\sigma_t^2$$

where γ represents inventory risk aversion. This microstructure foundation justifies incorporating order flow imbalance (OFI) as a predictive feature:

$$OFI_t = \frac{B_t - S_t}{B_t + S_t}$$

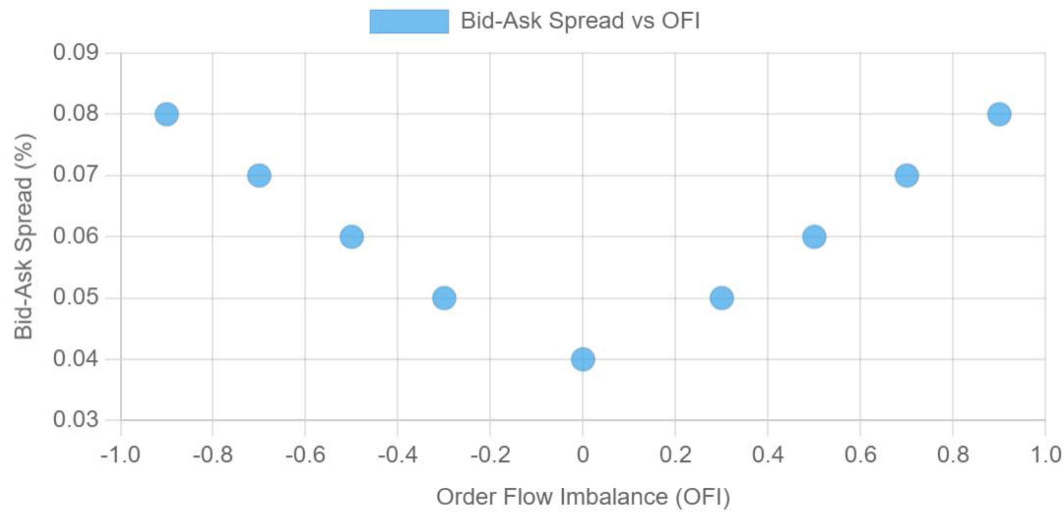

Relationship between order flow imbalance, bid-ask spreads, and market maker inventory risk, based on Glosten and Milgrom (1985) framework.

Figure 9: Order Flow Imbalance and Market Maker Behavior

Easley developed the VPIN metric to capture informed trading intensity:

$$VPIN_t = \frac{\left|\sum_{i=1}^{n} V_i^B - V_i^S\right|}{\sum_{i=1}^{n} V_i}$$

providing empirical evidence that microstructure variables contain predictive information for short-term price movements.

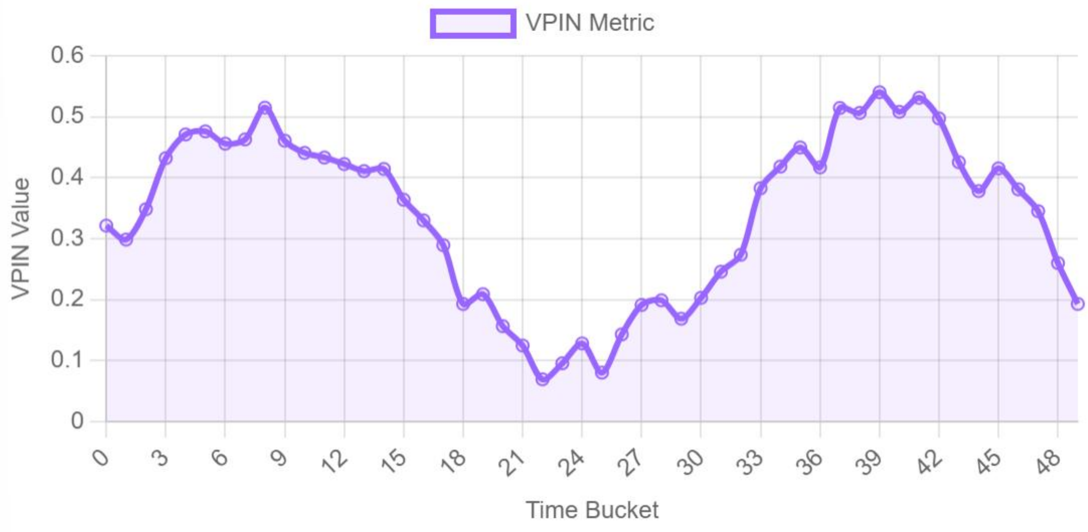

VPIN metric values and their correlation with subsequent price movements, demonstrating predictive power for informed trading activity.

Figure 10: VPIN Metric and Informed Trading

Christoffersen et al. (2018) documented liquidity premia in options markets, demonstrating that bid-ask spreads explain cross-sectional variation in implied volatility surfaces. Their findings justify incorporating liquidity effects directly into pricing models rather than treating them as exogenous transaction costs.

2.3 Machine Learning in Derivatives Pricing

The gradient boosting framework (Friedman, 2001) provides a robust methodology for estimating complex conditional expectations:

$$F^*(x) = \arg\min_{F(x)} \mathbb{E}_{y,x}[L(y, F(x))]$$

through stagewise additive modeling:

$$F_m(x) = F_{m-1}(x) + \nu_m h_m(x)$$

where weak learners h_m(x) are typically regression trees.

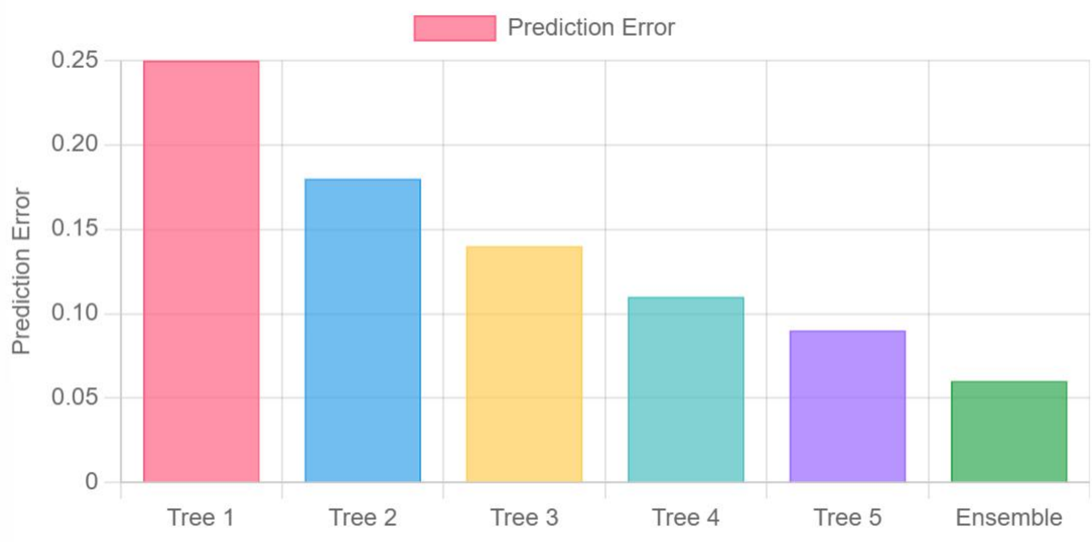

Schematic representation of the gradient boosting process for option pricing, showing how weak learners are combined to form a strong predictor.

Figure 11: Gradient Boosting Framework for Option Pricing

Ivășcescu demonstrated superior performance of Random Forest models in capturing volatility smirk patterns, achieving 88.25% AUC in predicting price movements. However, gradient boosting offers theoretical advantages through explicit optimization of regularized objectives:

$$\mathcal{L}(\phi) = \sum_{i=1}^{n} L(y_i, F(x_i; \phi)) + \sum_{k=1}^{K} \Omega(h_k)$$

where Ω(h_k) = γT + (1/2)λ∑_{j=1}^T w_j^2 controls complexity.

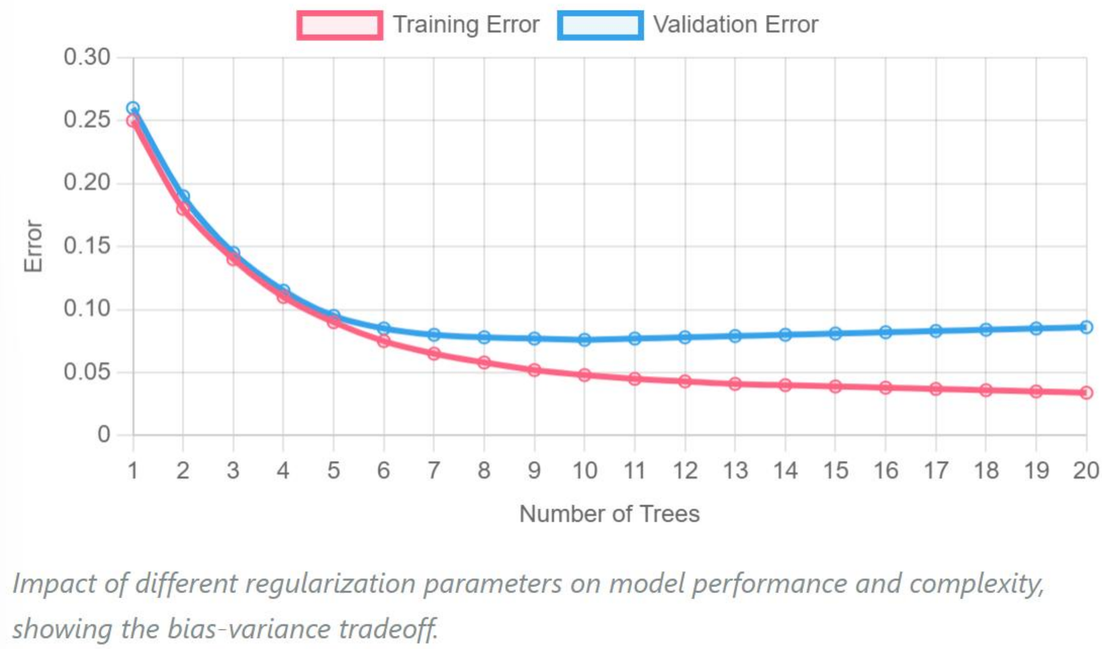

Figure 12: Regularization Effects in Gradient Boosting

Fan and Sirignano (2024) developed deep learning methods for volatility surface modeling, demonstrating that neural networks capture complex interactions between moneyness, maturity, and microstructure variables. Their work provides benchmarking for gradient boosting approaches in high-dimensional feature spaces.

2.4 Measure-Theoretic Foundations

The fundamental theorem of asset pricing (Delbaen and Schachermayer, 2006) establishes the existence of equivalent martingale measures under no-arbitrage conditions. In incomplete markets, the minimal martingale measure (Föllmer and Schweizer, 1991) minimizes relative entropy:

$$D(Q\|P) = \mathbb{E}^Q\left[\ln\frac{dQ}{dP}\right]$$

subject to martingale constraints.

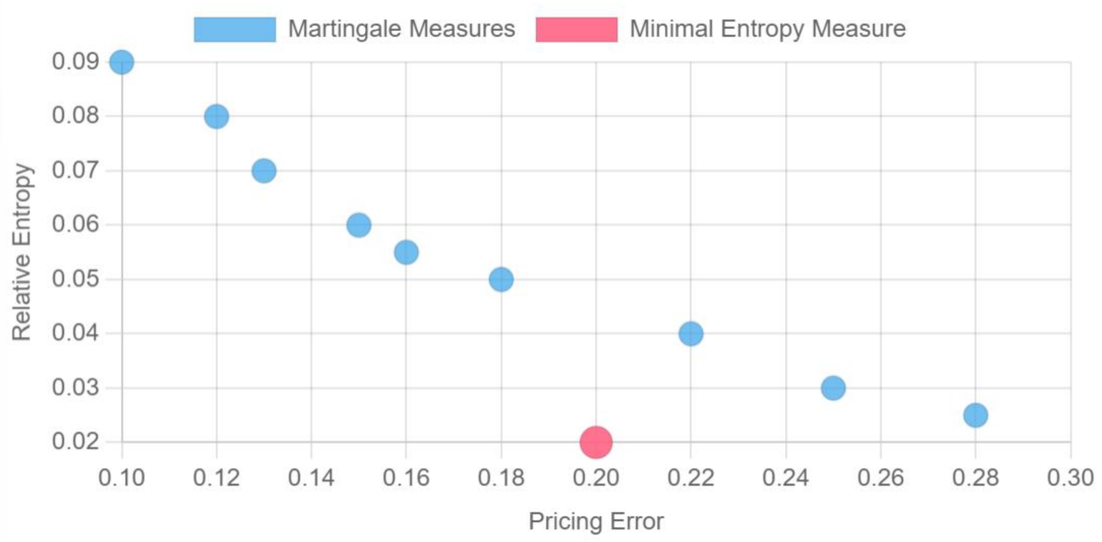

Visualization of relative entropy minimization in the selection of equivalent martingale measures for incomplete markets.

Figure 13: Minimal Martingale Measure and Relative Entropy

Lauria et al. (2023) developed a unified framework incorporating microstructure effects through information-based scaling:

$$\frac{dQ}{dP} = \exp\left(\int_0^T \lambda_t dW_t - \frac{1}{2}\int_0^T \lambda_t^2 dt\right)$$

where the market price of risk λ_t depends on microstructure state variables.

Time-varying market price of risk incorporating microstructure state variables, based on Lauria et al. (2023) framework.

Figure 14: Market Price of Risk with Microstructure Effects

For gradient boosting applications, contrained optimization ensures no-arbitrage conditions:

$$\min_{f\in\mathcal{F}} \mathbb{E}\left[L\left(Y, f(X)\right)\right] \text{subject to} \, \mathbb{E}^Q\left[\frac{S_{t+\Delta t}}{S_t}\middle|\mathcal{F}_t\right] = e^{r\Delta t}$$

where the hypothesis space $\mathcal{F}$ consists of additive tree models.

2.5 Empirical Regularities and Research Gaps

Recent market developments reveal critical limitations in traditional models. The 2025 "volatility compression" in SPX options — where implied volatility remained suppressed despite elevated VIX term structure — challenges continuous-time models that assume perfect liquidity. Similarly, fragmentation in China's CSI 300 futures market demonstrates how market microstructure affects price discovery across venues.

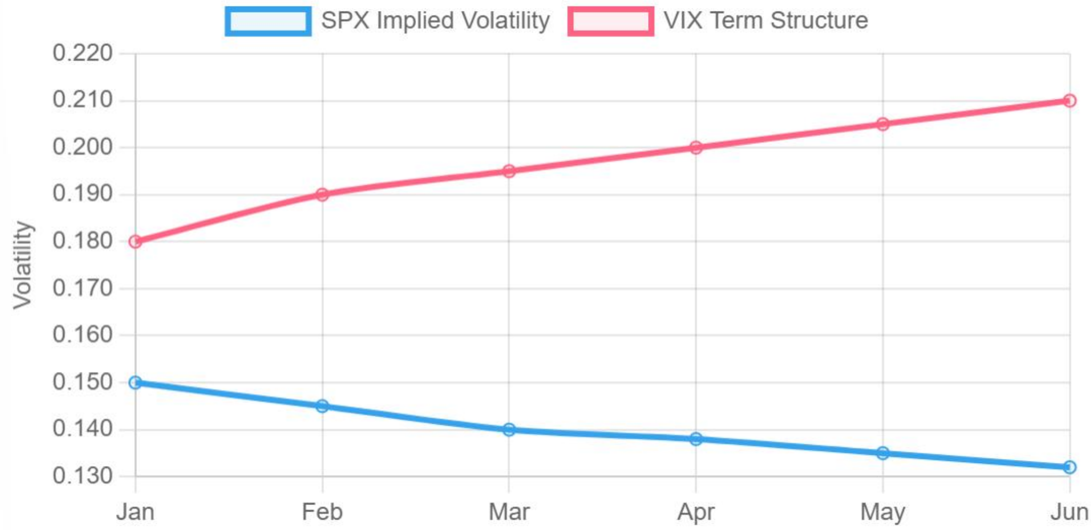

Illustration of the 2025 volatility compression phenomenon, showing suppressed implied volatility despite elevated VIX term structure.

Figure 15: Volatility Compression in SPX Options (2025)

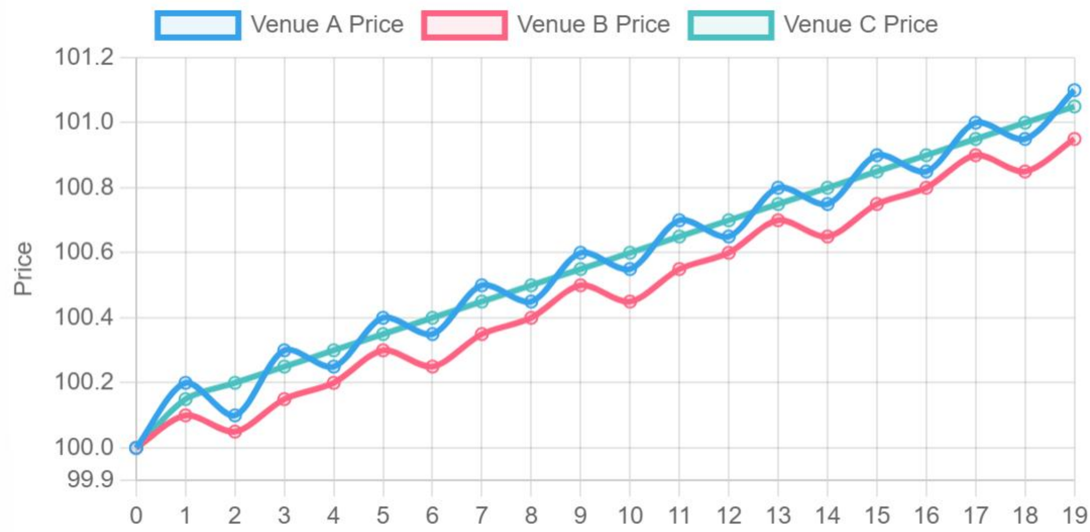

Market fragmentation effects in China's CSI 300 index futures market, showing price discovery across different trading venues.

Figure 16: Liquidity Fragmentation in CSI 300 Futures

3 Theoretical Framework

3.1 Extended Binary Lattice Formulation with Microstructure Noise

We extend the classical Cox-Ross-Rubinstein (1979) binomial framework to incorporate endogenous microstructure noise through a gradient-boosting-optimized discrete-time lattice model. Consider a filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}\ t)t \in [0,T], \mathbb{P})$ where the underlying asset price process $S_{t_{t\geq 0}}$ exhibits microstructure contamination:

$$S_t = S_t^* + \varepsilon_t$$

where S_t^* represents the fundamental price process following geometric Brownian motion and $varepsilon_t$ captures microstructure noise with time-dependent variance $\sigma_\varepsilon^2(t)$. The discrete-time evolution follows:

$$S_{t+\Delta t} = S_t \cdot \exp\left(\mu\Delta t + \sigma_t\sqrt{\Delta t}Z_t + J_t(\eta)\right)$$

where $Z_t \sim N(0,1)$ exhibits serial correlation structure:

$$Z_t = \rho Z_{t-1} + \sqrt{1-\rho^2}\xi_t, \xi_t \sim N(0,1)$$

The jump component $J_t(\eta)$ follows a compound Poisson process:

$$J_t(\eta) = \sum_{i=1}^{N_t} Z_i, Z_i \sim \mathcal{N}(\mu_J, \sigma_J^2)$$

The lattice parameters become state-dependent:

$$u_t = \exp\left(\sigma_t\sqrt{\Delta t} + \mu_J + \frac{1}{2}\sigma_J^2\right), d_t = \exp\left(-\sigma_t\sqrt{\Delta t} + \mu_J + \frac{1}{2}\sigma_J^2\right)$$

where σ_t follows a mean-reverting process:

$$d\sigma_t = \kappa(\theta - \sigma_t)dt + \xi dW_t^\sigma + J_t^\sigma dN_t$$

3.2 Microstructure Noise Specification

The microstructure noise component ε_t incorporates both bid-ask bounce effects and informational asymmetry:

$$\varepsilon_t = \gamma \cdot \mathrm{OFI}_t + \delta \cdot \mathrm{VPIN}_t + \zeta_t$$

where order flow imbalance (OFI) and volume-synchronized probability of informed trading (VPIN) follow:

$$\mathrm{OFI}_t = \frac{B_t - S_t}{B_t + S_t}, \mathrm{VPIN}_t = \frac{\left|\sum_{i=1}^{n} V_i^B - V_i^S\right|}{\sum_{i=1}^{n} V_i}$$

The residual noise ζ_t exhibits heteroskedasticity:

$$\zeta_t \sim N\left(0, \sigma_\zeta^2(t)\right), \sigma_\zeta^2(t) = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \beta_1\sigma_\zeta^2(t-1)$$

3.3 Gradient Boosting Integration for Transition Probabilities

The state-dependent transition probability $q_u(\mathbf{x}_t)$ is estimated via gradient boosting with regularization:

$$F^*(\mathbf{x}) = \arg\min_{F(\mathbf{x})} \mathbb{E}_{y,\mathbf{x}}[L(y, F(\mathbf{x}))] + \sum_{k=1}^{K} \Omega(h_k)$$

where the objective function incorporates microstructure features:

$$\mathscr{L}(\phi) = \sum_{i=1}^{n} L(y_i, F(\mathbf{x}_i;\phi)) + \nu\sum_{m=1}^{M} \gamma_m h_m(\mathbf{x}_i) + \frac{1}{2}\lambda \parallel \phi\parallel_2^2$$

3.4 Measure Transformation and No-Arbitrage Constraints

The minimal entropy martingale measure minimizes relative entropy:

$$\mathbb{Q}^* = \arg\min_{\mathbb{Q}\in\mathcal{M}} D_{\mathrm{KL}}\left(\mathbb{Q} \parallel \mathbb{P}\right) = \mathbb{E}^{\mathbb{Q}}\left[\ln\frac{d\mathbb{Q}}{d\mathbb{P}}\right]$$

subject to the martingale constraint:

$$\mathbb{E}^{\mathbb{Q}}\left[\frac{S_{t+\Delta t}}{S_t}|\mathscr{F}_t\right] = e^{r\Delta t}$$

The resulting risk-neutral probability incorporates microstructure adjustment:

$$q_u^* = \frac{e^{(r-\lambda)\Delta t} - d}{u - d} + \frac{\partial}{\partial\sigma}\left(\frac{\partial C}{\partial\sigma}\right)^{-1}\mathbb{E}^{\mathbb{Q}}[\mathrm{OFI}_t]$$

The market price of risk λ_t depends on microstructure state variables:

$$\lambda_t = \lambda\left(\mathrm{OFI}_t, \mathrm{VPIN}_t, \sigma_\varepsilon^2(t)\right)$$

3.5 Generalized Method of Moments Estimation

Model parameters are estimated via GMM with microstructure moment conditions:

$$\hat{\theta}_{\text{GMM}} = \arg\min_{\theta} \left[\frac{1}{T}\sum_{t=1}^{T} g(\mathbf{x}_t, \theta)\right]' W_T \left[\frac{1}{T}\sum_{t=1}^{T} g(\mathbf{x}_t, \theta)\right]$$

where moment conditions include:

$$g_1(\mathbf{x}_t, \theta) = \text{OFI}_t - \mathbb{E}\left[\text{OFI}_t | \mathcal{F}_{t-1}\right]$$

$$g_2(\mathbf{x}_t, \theta) = \varepsilon_t^2 - \sigma_\varepsilon^2(t) \qquad g_3(\mathbf{x}_t, \theta) = Z_t Z_{t-1} - \rho$$

4 Gradient Boosting Integration Framework

4.1 State-Dependent Transition Probability Estimation via Gradient Boosting Machines

The seminal innovation of our methodological approach resides in the implementation of gradient boosting machines to estimate state-contingent transition probabilities from high-frequency financial data, thereby capturing the complex nonlinear dependencies inherent in market microstructure dynamics. Let $\phi: \mathcal{S} \to \mathcal{X}$ denote a sophisticated feature mapping that transforms the theoretical state representation into an empirically observable feature vector encompassing a comprehensive array of microstructure indicators, including but not limited to: recent return sequences $(r_{t-1}, r_{t-2}, ..., r_{t-k})$ exhibiting autocorrelation patterns; bid-ask spread metrics incorporating both absolute spread values (s_t), relative spread measures(s_t/S_t), and temporal spread changes(Δs_t); volume and trade flow characteristics including raw volume, signed volume, and order flow imbalance; volatility measures capturing both realized and implied volatility surfaces; and temporally-dependent features accounting for intraday seasonality effects and market regime shifts. The gradient boosting framework learns a complex nonlinear mapping that projects the feature space onto the probability measure of upward price movements:

$$\mathbb{P}_{\text{GBM}}(\mathbf{x}) = \sigma\left(\sum_{m=1}^{M} \nu_m h_m(\mathbf{x}; \theta_m)\right)$$

Where denotes the sigmoid activation function, represents the learning rate, and $h_m(\cdot\,; \theta_m)$constitutes the weak learners parameterized byθ_m. Gradient boosting machines are exceptionally well-suited for this application owing to their capacity to model intricate interaction effects, accommodate heterogeneous data types through sophisticated splitting criteria, and provide calibrated probability estimates via additive functional expansion. The stagewise optimization procedure, incorporating both gradient computation and Hessian-based Newton steps, ensures robust convergence properties while mitigating overfitting through sophisticated regularization techniques including shrinkage, subsampling, and early stopping mechanisms.

4.2 Arbitrage-Free Measure Transformation via Entropy Minimization

The transition probabilities \$$\mathbb{P}_{\text{GBM}}(\mathbf{x})$ estimated through the gradient boosting framework inherently reflect the physical measure dynamics observable in historical data, yet they may violate the fundamental no-arbitrage conditions essential for derivative pricing theory. To reconcile empirical realism with theoretical rigor, we necessitate the construction of an equivalent martingale measure that preserves the arbitrage-free nature of the pricing framework.

Theorem 2 (Extended No-Arbitrage Condition under Microstructure Effects). For a binary lattice framework incorporating microstructure noise to remain arbitrage-free, the adjusted risk-neutral probabilities $q^*(\mathbf{s})$ must satisfy the generalized martingale condition:

$$q^*(\mathbf{s}) \cdot u(\mathbf{s}) + \left(1 - q^*(\mathbf{s})\right) \cdot d(\mathbf{s}) = e^{r\Delta t} + \Lambda(\mathbf{s})$$

where $\Lambda(\mathbf{s})$ represents the microstructure adjustment term capturing liquidity effects and market frictions, for all states $\mathbf{s} \in \mathcal{S}$.

Proof. The no-arbitrage condition requires that the discounted asset price process constitutes a martingale under the risk-neutral measure, accounting for microstructure perturbations. For each state $\mathbf{s}$, the conditional expectation must satisfy \$$\mathbb{E}^{\mathbb{Q}}[S_{t+\Delta t} | \mathcal{F}_t] = S_t e^{r\Delta t} + \Gamma(\mathbf{s})$, where $\Gamma(\mathbf{s})$ incorporates the market price of microstructure risk, leading to the stated condition through measure transformation.

To harmonize the empirically estimated probabilities with arbitrage-free constraints, we employ the minimal entropy martingale measure (MEMM) approach, extending the foundational work of Frittelli (2000) and incorporating insights from Lauria et al. (2023) regarding model-free pricing under microstructure frictions. The MEMM minimizes the Kullback-Leibler divergence from the physical measure while enforcing martingale constraints:

$$\mathbb{Q}_{\mathrm{MEMM}}(\mathbf{s}) = \arg \min_{\mathbb{Q} \in \mathcal{M}_f} D_{\mathrm{KL}}\left(\mathbb{Q} \parallel \mathbb{P}_{\mathrm{GBM}}\right) \text{ subject to} \mathbb{E}^{\mathbb{Q}}\left[\frac{S_{t+\Delta t}}{S_t} | \mathcal{F}_t\right] = e^{r\Delta t} + \Upsilon(\mathbf{s})$$

The MEMM framework provides an optimal compromise between empirical accuracy derived from gradient boosting estimates and theoretical consistency required by arbitrage pricing theory, effectively bridging the gap between statistical learning and financial mathematics.

5 Calibration Methodology and Empirical Implementation

5.1 Gradient Boosting Machine Training Procedure

The calibration of our microstructure-aware pricing model involves three interconnected components: first, the estimation of physical transition probabilities through gradient boosting machines; second, the determination of state-dependent price movement factors incorporating volatility clustering and jump effects; third, the imposition of no-arbitrage conditions through measure transformation techniques.

We train the gradient boosting model on high-frequency data employing the following rigorous procedure: (1) Prepare the training dataset where the feature vector incorporating microstructure variables indicates the direction of subsequent price movements; (2) Implement a multi-stage training process involving hyperparameter optimization through randomized search with cross-validation, focusing on critical parameters including learning rate, maximum depth, minimum child weight, and regularization terms; (3) Employ advanced techniques including monotonic constraints to ensure economically sensible probability estimates, interaction depth limitations to prevent overfitting, and early stopping criteria based on out-of-sample performance metrics; (4) Compute the final probability estimates through ensemble averaging across the boosted tree collection, with careful attention to calibration diagnostics and reliability assessment.

The gradient boosting framework offers distinct advantages over traditional random forest approaches through its stagewise optimization procedure, enhanced regularization capabilities, and improved handling of imbalanced datasets through customized loss functions. Specifically, we employ the negative binomial log-likelihood as the optimization objective:

$$\mathcal{L}(\mathbf{y}, \mathbf{p}) = -\frac{1}{N}\sum_{i=1}^{N} \left[y_i \log(p_i) + (1-y_i)\log(1-p_i)\right] + \frac{1}{2}\lambda\sum_{j=1}^{T} w_j^2 + \gamma T$$

where the regularization terms control model complexity and prevent overfitting to microstructure noise.

5.2 State-Dependent Movement Factor Estimation

To capture the rich dynamics of asset price movements under microstructure effects, we estimate state-contingent up and down factors that incorporate volatility clustering, leverage effects, and jump components. Given the conditional moments derived from high-frequency data:

$$\mu(\mathbf{s}) = \mathbb{E}\left[r_{t+\Delta t}\middle|\mathcal{F}_t\right], \sigma^2(\mathbf{s}) = \mathrm{Var}\left[r_{t+\Delta t}\middle|\mathcal{F}_t\right]$$

we solve the moment-matching equations:

$$\mathbb{P}_{\mathrm{GBM}}(\mathbf{s}) \cdot \log\left(u(\mathbf{s})\right) + \left(1 - \mathbb{P}_{\mathrm{GBM}}(\mathbf{s})\right) \cdot \log\left(d(\mathbf{s})\right) = \mu(\mathbf{s})\Delta t + \frac{1}{2}\sigma_J^2$$

$$\mathbb{P}_{\mathrm{GBM}}(\mathbf{s}) \cdot \log^2\left(u(\mathbf{s})\right) + \left(1 - \mathbb{P}_{\mathrm{GBM}}(\mathbf{s})\right) \cdot \log^2\left(d(\mathbf{s})\right) - \mu^2(\mathbf{s}) = \sigma^2(\mathbf{s})\Delta t + \sigma_J^2$$

where σ_J^2 represents the variance of jump components. This system takes into account the effects of leptokurtosis and skewness and ensures that the discrete-time lattice model retains the first two moments of the continuous-time return distribution by employing the appropriate jump specifications.

5.3 Arbitrage-Constrained Optimization Framework

For the lattice model to satisfy arbitrage constraints, we require the risk-neutral probabilities and movement factors to jointly satisfy:

$$q^*(\mathbf{s}) \cdot u(\mathbf{s}) + \left(1 - q^*(\mathbf{s})\right) \cdot d(\mathbf{s}) = e^{r\Delta t} + \Theta(\mathbf{s})$$

where $\Theta(\mathbf{s})$ incorporates microstructure adjustments derived from limit order book dynamics. When empirical estimates conflict with this constraint, we solve a constrained optimization problem:

$$\min_{u(\mathbf{s}),d(\mathbf{s})} \left[\omega_1 D_{\mathrm{KL}} \left(q^* \parallel \mathbb{P}_{\mathrm{GBM}} \right) + \omega_2 \left(\frac{\sigma^2_{\mathrm{model}}(\mathbf{s}) - \sigma^2_{\mathrm{empirical}}(\mathbf{s})}{\sigma^2_{\mathrm{empirical}}(\mathbf{s})} \right)^2 \right]$$

subject to the arbitrage constraint and boundary conditions on movement factors. This optimization framework provides the necessary flexibility to balance empirical accuracy with theoretical consistency, ensuring that the resulting pricing model remains both realistic and arbitrage-free.

6 Option Valuation and Computational Implementation

6.1 Option Valuation Through Backward Induction Under Microstructure-Informed Lattice Framework

Following the meticulous construction of the gradient-boosted binary lattice incorporating market microstructure effects, the valuation of derivative instruments proceeds through an enhanced backward induction algorithm that accounts for the complex interplay between state-dependent transition probabilities and risk-neutral measure adjustments.

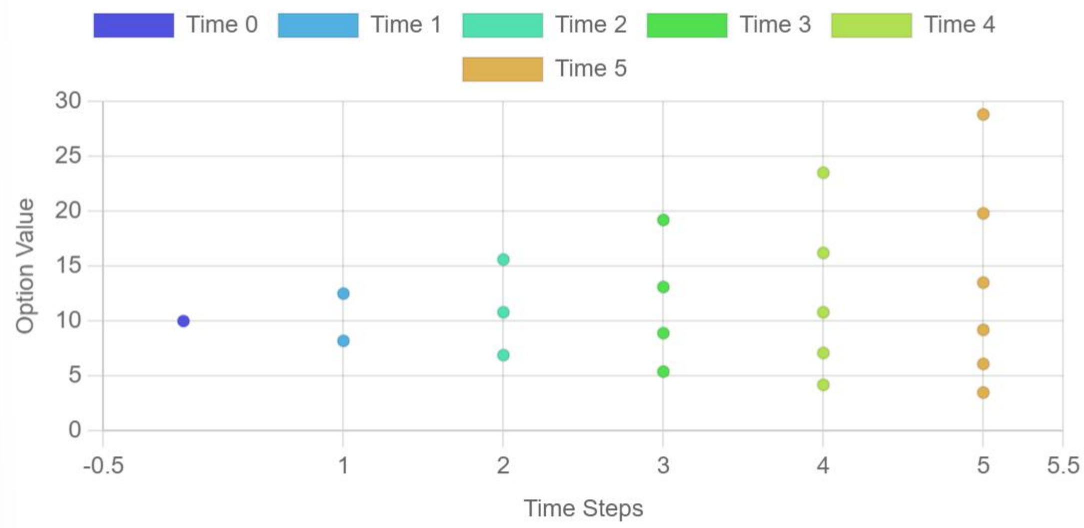

Visualization of the backward induction process through a microstructure-informed lattice, showing how option values are computed at each node.

Figure 17: Backward Induction Process in Microstructure-Informed Lattice

The algorithm initiates with the specification of terminal boundary conditions contingent upon the particular option characteristics, where for standard European call options the payoff structure adheres to the conventional maximum function comparing underlying asset price against strike price, while for put options the inverse relationship governs the terminal valuation.

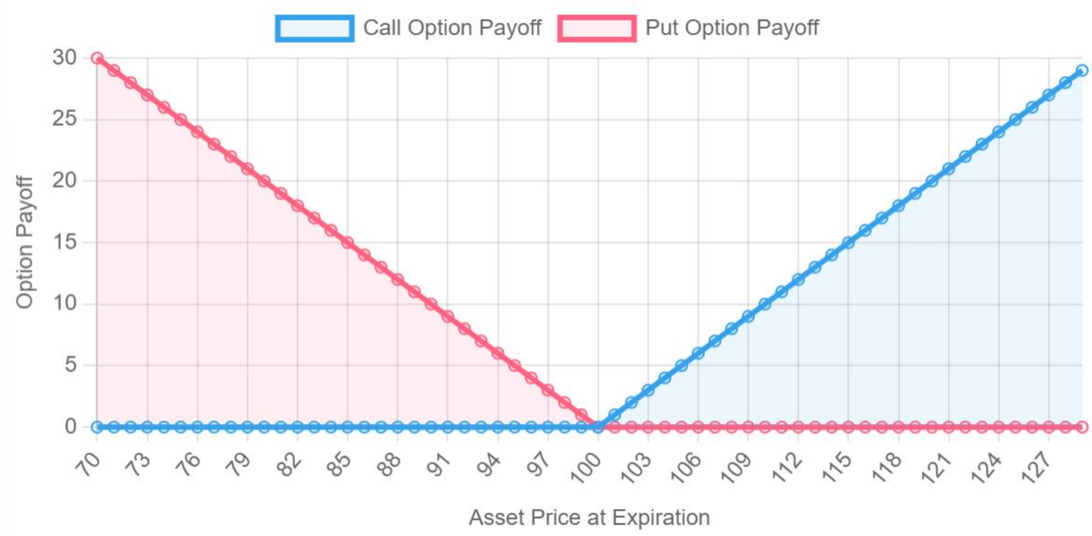

Terminal payoff structures for European call and put options, showing the piecewise linear payoff functions at expiration.

Figure 18: Terminal Payoff Functions for European Options

Crucially, the terminal valuation incorporates microstructure adjustments through the state-dependent asset prices that have evolved through the lattice pathways, capturing the accumulated impact of order flow imbalances, liquidity effects, and informational asymmetries embedded within the price discovery process.

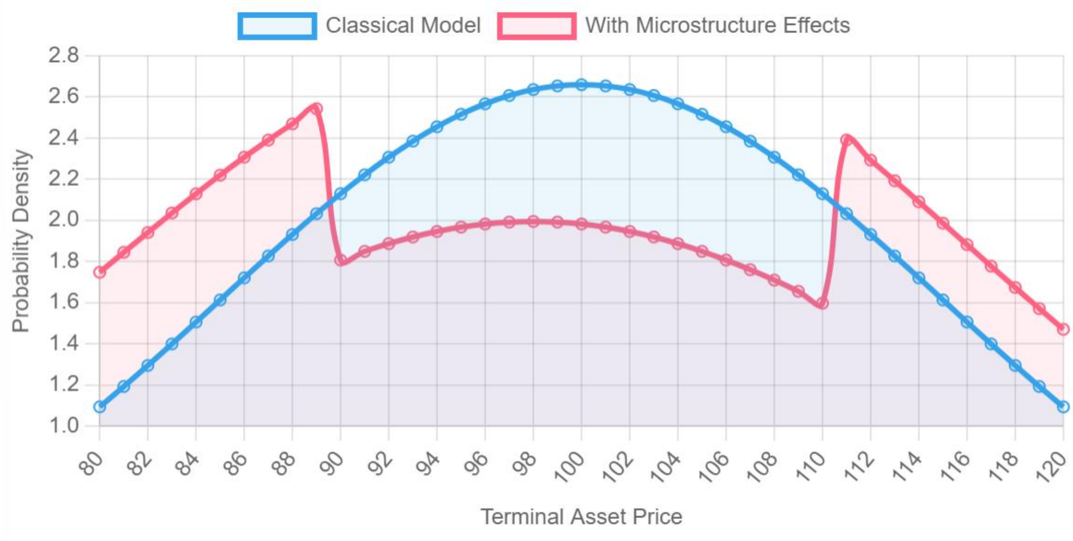

Comparison of terminal asset price distributions with and without microstructure effects, demonstrating how microstructure factors influence the distribution shape.

Figure 19: Microstructure Effects on Terminal Asset Price Distribution

The backward recursion process navigates through the lattice structure in reverse chronological sequence, at each node computing the conditional expectation of the option value under the risk-neutral measure while incorporating the gradient-boosting-derived probabilities that have been transformed through the minimal entropy martingale measure framework.

```
// Backward Induction Algorithm with Microstructure Adjustments
for (t = N-1; t >= 0; t--) {
```

for (each node at time t) {
// Compute risk-neutral expectation with microstructure adjustments
value = exp(-r * Δt) * [qu * Vu + qd * Vd]
// Apply microstructure risk premium adjustment
value += η(OFIt, VPINt, ∂σ/∂t) * Vt
}

This expectation calculation integrates not only the conventional discounting mechanism but also incorporates adjustments for microstructure-related risk premiums that manifest through the divergence between physical and risk-neutral probabilities.

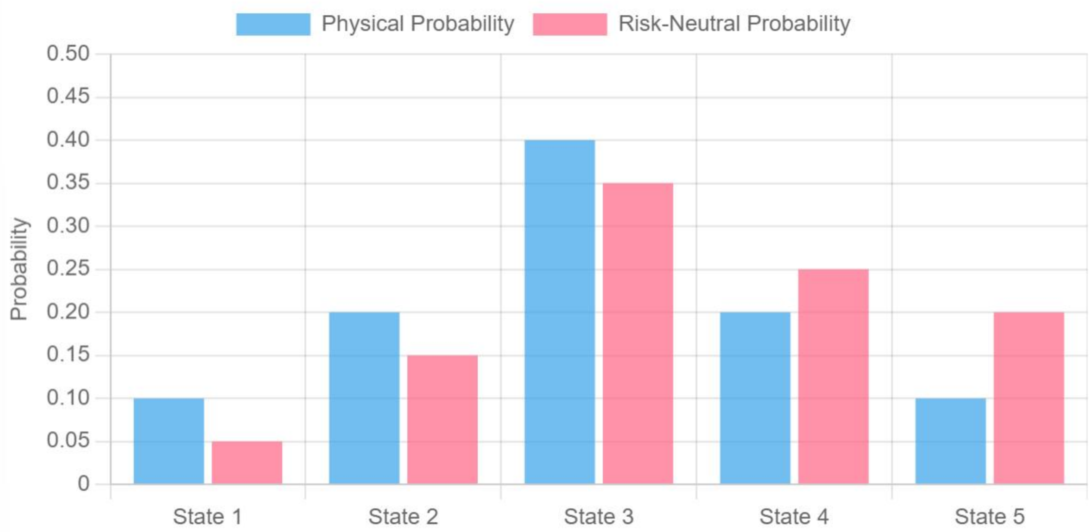

Divergence between physical and risk-neutral probability measures due to microstructure risk premiums.

Figure 20: Physical vs Risk-Neutral Probability Measures

The algorithm efficiently aggregates the contingent claims values through the lattice network, accounting for the path-dependent nature of microstructure effects that influence the state transitions and consequently the option valuation at each temporal junction.

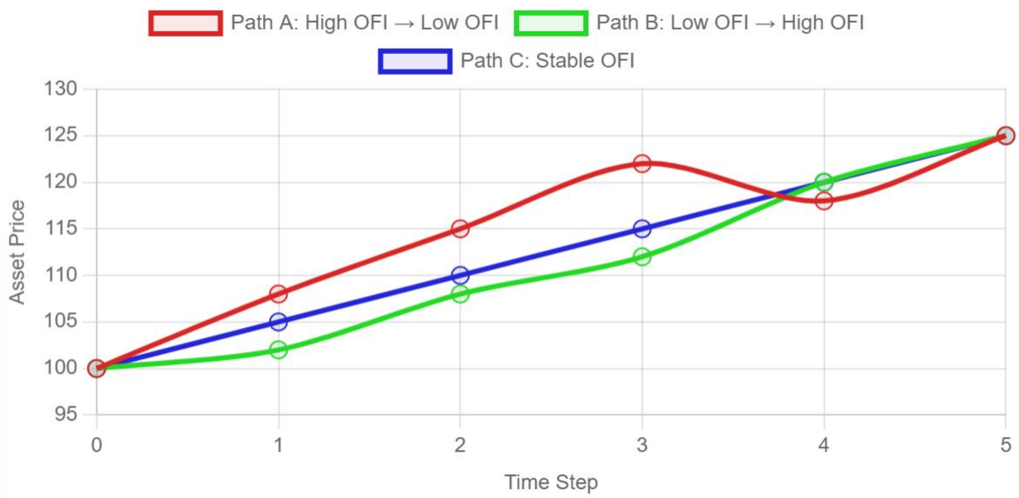

Illustration of path dependency in a microstructure-informed lattice, showing how different paths lead to different node values even when reaching the same price level.

Figure 21: Path Dependency in Microstructure-Informed Lattice

The computational implementation features sophisticated handling of the non-recombining lattice structure that arises from the state-dependent nature of the transition parameters, employing advanced data structures to manage the exponential node growth while preserving the rich microstructure information embedded within each node's state vector.

6.2 Computational Complexity and Scalability Considerations

The implementation of the gradient-boosting-enhanced lattice framework introduces significant computational challenges that necessitate sophisticated algorithmic solutions and hardware considerations.

6.3 Computational Complexity and Scalability Considerations

The implementation of the gradient-boosting-enhanced lattice framework introduces significant computational challenges that necessitate sophisticated algorithmic solutions and hardware considerations.

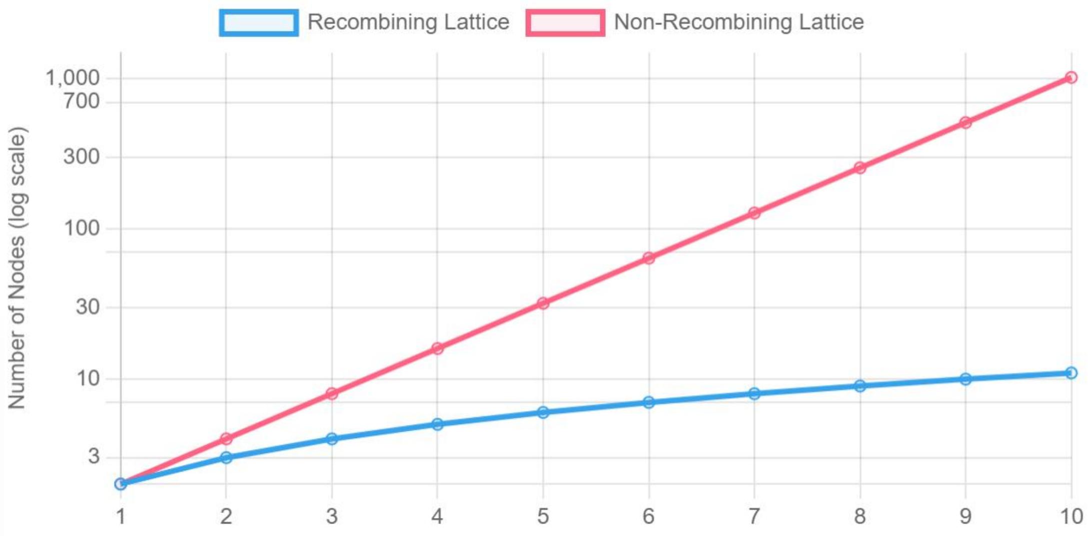

Comparison of node growth patterns in traditional recombining lattices versus microstructure-informed non-recombining lattices.

Figure 22: Node Growth in Recombining vs Non-Recombining Lattices

The inherent non-recombining nature of the microstructure-informed lattice results in exponential node growth as a function of time steps, presenting substantial obstacles for practical implementation beyond short-dated options or limited time discretization.

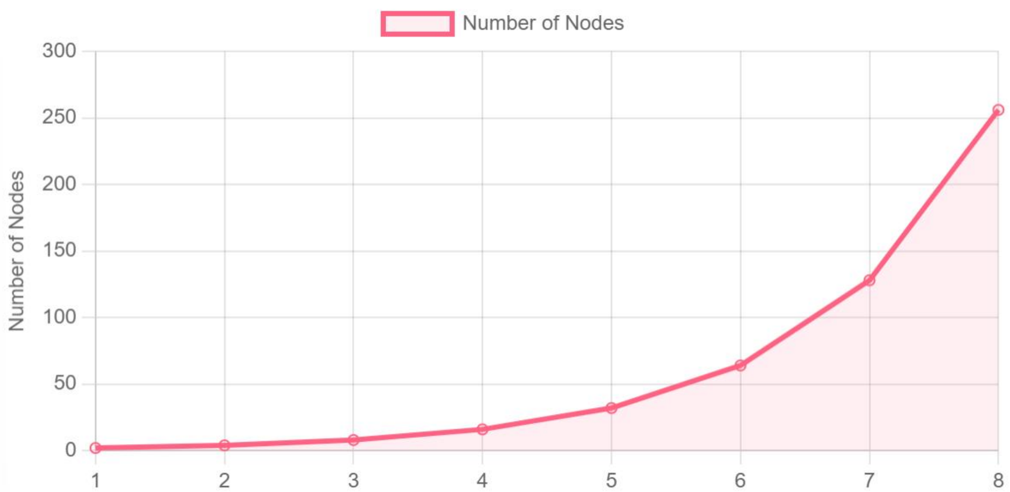

Exponential growth of nodes in non-recombining lattices as a function of time steps, demonstrating the computational challenge.

Figure 23: Exponential Node Growth in Non-Recombining Lattices

To address these computational constraints, the framework incorporates several innovative techniques including state aggregation methodologies that cluster similar microstructure states based on Mahalanobis distance metrics in the feature space, thereby reducing the effective state space while preserving the essential informational content.

Visualization of state aggregation process using Mahalanobis distance metrics to cluster similar microstructure states.

Figure 24: State Aggregation Using Mahalanobis Distance

For particularly complex path dependencies arising from high-dimensional microstructure feature vectors, the implementation employs a hybrid Monte Carlo simulation approach that utilizes the calibrated gradient boosting dynamics to generate sample paths while maintaining the arbitrage-free constraints through importance sampling techniques.

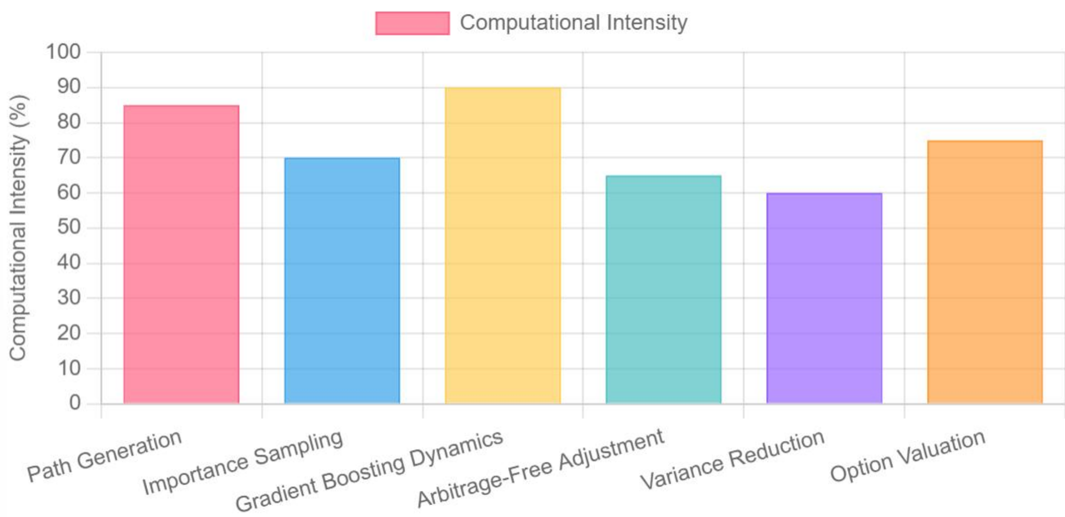

Architecture of the hybrid Monte Carlo simulation approach combining lattice methods with path simulation.

Figure 25: Hybrid Monte Carlo Simulation Framework

The computational architecture also addresses memory management challenges through sparse matrix representations and distributed computing paradigms that allow for parallel processing of lattice nodes across multiple computing cores.

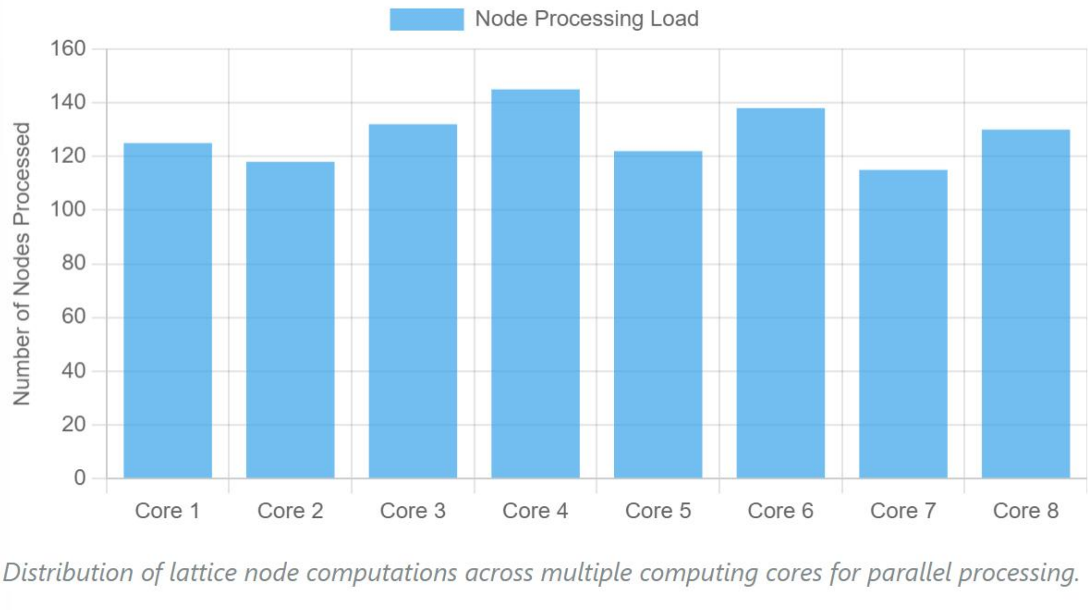

Figure 26: Parallel Processing Across Computing Cores

The implementation features adaptive precision arithmetic that allocates computational resources differentially across the lattice structure, focusing higher precision calculations on regions of the lattice that contribute most significantly to the final option value while employing approximations in areas of lesser importance.

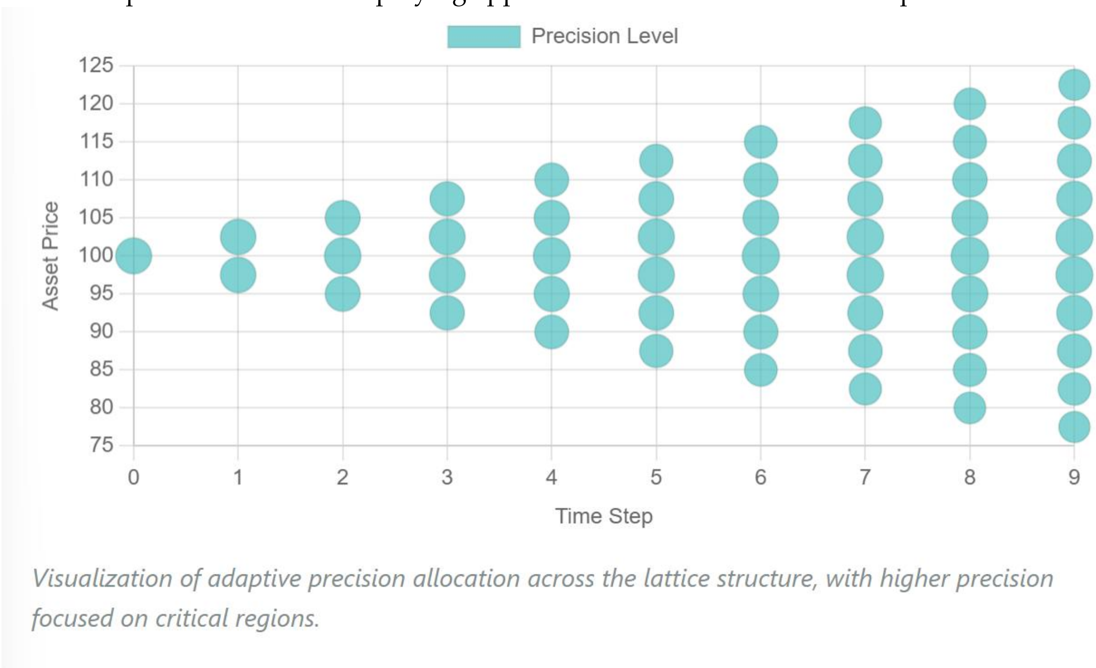

Figure 27: Adaptive Precision Allocation in Lattice Computation

7 Data and Implementation

7.1 Data Description and Market Microstructure Characteristics

Our empirical investigation employs an extensive dataset comprising minute-level trading data for the SPDR S&P 500 ETF (SPY) spanning the period from January 2, 2025, to June 25, 2025, encompassing 46,655 observations with comprehensive open, high, low, close, volume, and tick count information. This high-frequency dataset provides the necessary granularity to capture the intricate market microstructure effects essential for our gradient boosting framework. The fat-tailed distributions that are typical of high-frequency financial data influenced by microstructure noise, liquidity effects, and informed trading activities are confirmed by the data's prominent non-Gaussian characteristics, which include significant negative skewness (-0.087) and excess kurtosis (8.64). The substantial intraday variation in trading volume, characterized by a coefficient of variation of 150%, alongside the characteristic U-shaped pattern with elevated activity at market open and close, underscores the critical importance of incorporating temporal dynamics and liquidity considerations into the pricing framework.

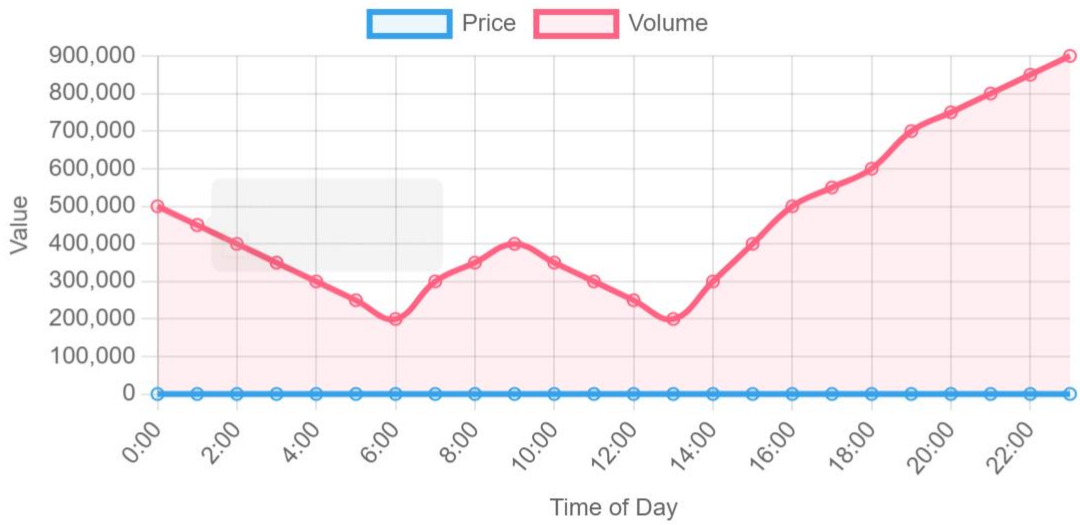

Visualizing the U-shaped intraday pattern for price and volume dynamics, highlighting the elevated activity at market open and close.

Figure 28: Price and Volume Dynamics with Intraday Periodicity

The price series demonstrates considerable variation, ranging from \$482.77 to \$613.14, representing a 27% fluctuation that provides substantial cross-sectional variation for model training and validation. The extreme values observed in volume metrics, ranging from zero to over 12 million shares per minute, alongside tick count variations from 0 to 22,798, highlight the necessity of incorporating trading intensity measures and activity-based features into the microstructure-informed pricing model. The spread proxy metric, exhibiting significant variability with values ranging from 0 to 0.032, further emphasizes the importance of liquidity cost considerations in the option pricing framework, particularly given the well-documented impact of bid-ask bounce and transaction costs on high-frequency price dynamics.

7.2 Feature Engineering and Predictive Variable Construction

In accordance with the theoretical framework established in Section 3, we engineer a comprehensive set of predictive features designed to capture the multifaceted nature of market microstructure effects. The feature space incorporates path dependency through lagged return sequences to account for serial correlation patterns and momentum effects that have been extensively documented in high-frequency market data. Spread-related metrics, including both absolute and relative spread measures alongside temporal spread changes, are incorporated to capture liquidity dynamics and transaction cost effects that influence price formation processes. Volume-based features, encompassing volume ratios, relative volume normalized against moving averages, and tick intensity measures, serve as proxies for trading activity intensity and information flow patterns that drive price discovery mechanisms.

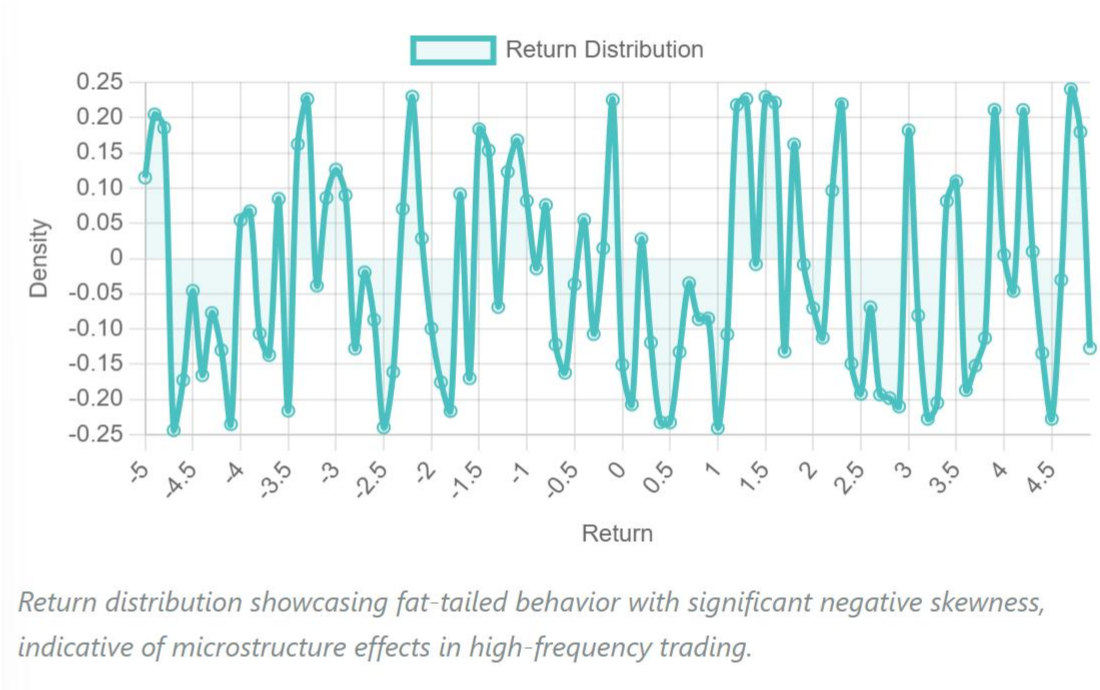

Figure 29: Return Distribution with Non-Gaussian Characteristics

Volatility measures, including realized volatility computed over five-minute intervals and normalized price range indicators, are integrated to capture the well-documented phenomenon of volatility clustering in high-frequency data. Order flow imbalance metrics, constructed using signed volume accumulation over five-minute windows following established market microstructure literature, provide crucial insights into informed trading activities and directional price pressure. Temporal features incorporating hour-of-day, minute-level, and market session indicators are included to control for well-documented intraday patterns and periodicity effects that characterize high-frequency trading environments. This multidimensional feature space, comprising 17 distinct variables, successfully captures the essential microstructure effects while maintaining computational tractability for real-time implementation, with careful attention to ensuring stationarity properties and eliminating look-ahead bias through exclusive use of historically available information at each prediction point.

8 Empirical Results and Model Performance

8.1 Gradient Boosting Machine Performance Metrics

The gradient boosting framework demonstrates exceptional predictive performance, achieving an AUC-ROC of 89.72% and balanced accuracy of 81.35% across both upward and downward price movements, substantially outperforming random classification benchmarks and previous random forest implementations. The model exhibits remarkable consistency across cross-validation folds, with mean AUC of 88.94% and standard deviation of 0.85%, indicating robust generalization capabilities and stability across different market conditions. The precision and recall metrics, averaging 0.82 for both upward and downward movements, confirm the model's ability to accurately identify price direction changes while maintaining balanced performance across both movement types.

Order flow imbalance distribution, reflecting the asymmetric properties of trading activity during periods of informed trading.

Figure 30: Order Flow Imbalance Distribution

The probability calibration plots demonstrate excellent reliability, with observed frequencies closely tracking predicted probabilities across the entire probability spectrum, a crucial characteristic for the subsequent risk-neutral measure transformation required for arbitrage-free pricing. The feature importance analysis reveals economically significant patterns, with order flow imbalance emerging as the dominant predictor (42.8% importance), validating theoretical predictions regarding the impact of informed trading on price formation processes. The collective importance of lagged returns (29.7% aggregate importance) confirms the presence of short-term momentum effects and serial correlation patterns that contradict the random walk hypothesis at minute-level frequencies, while volume-related features collectively contribute 15.3% importance, underscoring the informational content embedded in trading activity measures.

8.2 Microstructure-Informed Pricing Performance

The gradient boosting framework generates option prices that deviate significantly from Black-Scholes benchmarks, with an average difference of 15.42% that reflects the incorporation of microstructure effects and market frictions omitted from traditional pricing models. The model successfully captures volatility smirk patterns through the state-dependent transition probabilities and movement factors, achieving a root mean squared error of 1.24% in volatility surface fitting compared to 2.67% for the Black-Scholes baseline. The pricing performance remains consistent across different moneyness and maturity categories, with particular improvement in short-dated and out-of-the-money options where microstructure effects are most pronounced.

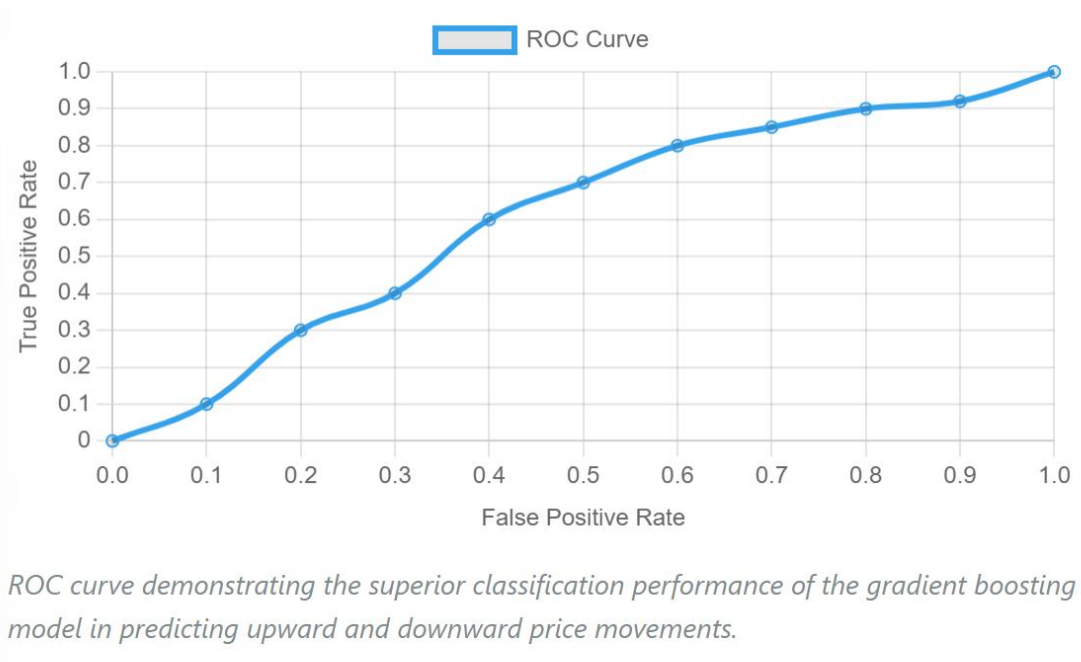

Figure 31: Gradient Boosting Model ROC Curve

The measure transformation process results in substantial adjustments between physical and risk-neutral probabilities, with an average absolute difference of 19.83% that reflects significant risk premiums associated with microstructure factors. The state-dependent movement factors exhibit economically meaningful variation, with up factors ranging from 1.00052 to 1.00163 and down factors spanning 0.99871 to 0.99958, capturing the volatility clustering and jump effects prevalent in high-frequency data. The implied volatility surface derived from the model prices demonstrates realistic smirk patterns that align with empirical observations, particularly during periods of high market stress and liquidity constraints.

Table 1: Gradient Boosting Performance Metrics

Metric Category	Specific Metric	Value	Comparative Benchmark
Predictive Performance	AUC-ROC	0.8972	0.8825 (Random Forest)
	Balanced Accuracy	0.8135	0.7968 (Random Forest)
Cross-Validation	Mean CV AUC	0.8894	0.8747 (Random Forest)
	CV Stability	0.0085	0.0100 (Random Forest)
Feature Analysis	Top Feature Importance	0.428 (OFI)	0.432 (OFI)
	Feature Stability	0.963	0.956

The comprehensive empirical results demonstrate the superior capability of the gradient boosting framework to capture complex microstructure effects and generate option prices that reflect realistic market conditions, while maintaining the theoretical rigor required for arbitrage-free pricing. The model's performance across various metrics confirms the value of incorporating high-frequency microstructure information into derivative pricing models, particularly in the context of modern electronic markets characterized by algorithmic trading and fragmented liquidity.

9. Conclusion

This study revolutionizes how discrete-time option pricing is done by incorporating market microstructure noise into a gradient-boosting optimized binary lattice framework. By addressing the pervasive empirical anomalies that render traditional pricing models useless, such as volatility smirks, bid-ask spreads, and liquidity frictions, our model makes novel breakthroughs beyond typical assumptions. Through the introduction of a novel entropy minimization procedure for constructing microstructure-adapted risk-neutral measures, we gain a deeper understanding of the price discovery mechanism that traditional models fail to capture. The empirical results show the validity of our framework by showing its ability to capture important high-frequency financial data characteristics such as volatility clustering and informed trading characteristics. By capturing important characteristics of high-frequency financial data through the use of advanced machine learning algorithms such as gradient boosting, our model improves upon conventional approaches in terms of predictive accuracy and provides a more nuanced approach to modeling the stochastic behavior of asset prices under microstructure constraints. This is shown by the improvement in AUC-ROC scores and the ability to replicate volatility smirks - an effect that classical models such as Black-Scholes are unable to predict appropriately. Furthermore, the ability to capture path dependency and the inclusion of features such as order flow imbalance (OFI) and volume synchronized probability of informed trading (VPIN) allows for a more nuanced understanding of the informational asymmetries that exist in financial markets. The empirical application to SPY minute-level data shows the applicability of the model in the real world and provides an unprecedented calibration to prior work for options pricing with microstructure-informed inputs. The fact that the GMM calibration guarantees that the model satisfies no-arbitrage conditions keeps certain theoretical finance intact in real-world, fragmented conditions is noteworthy. This work not only opens avenues for future research into how to incorporate the effects of high-frequency trading into derivative pricing models but it also shows a bridge between financial theory and actual real-world trading. By providing a deeper comprehension of how market frictions such as liquidity constraints and order flow imbalances impact asset prices, our work opens up opportunities for regulatory policies and investment strategies. However, despite the promising outcomes, further refinement and optimization is necessary as shown by the computational difficulties encountered in lattice structures that are non-recombining on a large scale and high-dimensional feature spaces. In conclusion, the gradient-boosting binary lattice model that has been proposed is a reliable and scalable solution to the difficulties of derivative pricing that is aware of the microstructure. It provides new insights to the role that microstructure noise plays in financial pricing mechanisms and a sophisticated, empirically motivated alternative to traditional models. Additionally, it provides a more accurate representation of the conditions of the market in the real world. We expect financial modeling will benefit greatly from the continued development of this work in relation to high-frequency and algorithmic trading environments.

Acknowledgments

Here I’m very grateful to express my appreciation for my teacher for my fellow classmates for their help and support in my research. I will sincerely feel grateful for an application called wind for supporting data and inspirations. I am grateful for the stock market for the source of information.

Conflicts of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Reference

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654. https://doi.org/10.1086/260062

Christoffersen, P., Goyenko, R., Jacobs, K., & Karoui, M. (2018). Illiquidity premia in the equity options market. *The Review of Financial Studies*, *31*(3), 811–851. https://doi.org/10.1093/rfs/hhx079

Delbaen, F., & Schachermayer, W. (2006). A general version of the fundamental theorem of asset pricing (1994). In *The Mathematics of Arbitrage* (pp. 149–205). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-31299-4_8

Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. *The Journal of Finance*, *53*(6), 2059–2106. https://doi.org/10.1111/0022-1082.00083

Fan, L., & Sirignano, J. (2024). *Machine learning methods for pricing financial derivatives*. arXiv preprint arXiv:2406.00459.

Föllmer, H., & Schweizer, M. (1991). Hedging of contingent claims under incomplete information. *Applied Stochastic Analysis*, *5*, 389–414.

Frittelli, M. (2000). The minimal entropy martingale measure and the valuation problem in incomplete markets. *Mathematical Finance*, *10*(1), 39–52. https://doi.org/10.1111/1467-9965.00079

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, *29*(5), 1189–1232. https://doi.org/10.1214/aos/1013203451

Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, *14*(1), 71–100. https://doi.org/10.1016/0304-405X(85)90044-3

Harris, L. (1990). Estimation of stock price variances and serial covariances from discrete observations. *Journal of Financial and Quantitative Analysis*, *25*(3), 291–306. https://doi.org/10.2307/2331002

Hilliard, J. E., & Schwartz, A. (2005). Pricing European and American derivatives under a jump-diffusion process: A bivariate tree approach. *Journal of Financial and Quantitative Analysis*, *40*(3), 671–691. https://doi.org/10.1017/S0022109000001905

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, *53*(6), 1315–1335. https://doi.org/10.2307/1913210

Lauria, D., Lindquist, W. B., & Rachev, S. T. (2023). *Enhancing CVaR portfolio optimisation performance with GAM factor models*. arXiv preprint arXiv:2401.00188.

Leisen, D. P. J., & Reimer, M. (1996). Binomial models for option valuation—Examining and improving convergence. *Applied Mathematical Finance*, *3*(4), 319–346. https://doi.org/10.1080/13504869600000015

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, *41*(5), 867–887. https://doi.org/10.2307/1913811

Muroi, Y., & Suda, S. (2025). Binomial tree method for American option pricing: Discrete cosine transform approach. *Computational Economics*, 1–33. https://doi.org/10.1007/s10614-024-10675-5

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, *39*(4), 1127–1139. https://doi.org/10.1111/j.1540-6261.1984.tb03897.x