International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

# Evaluating Machine Learning Alternatives to the Black-Scholes-Merton Framework for Option Pricing: An Empirical Investigation of Deep Out-of-the-Money Regimes

**Anmol Bajpai**
MBA (Finance) School of Business Management, NMIMS

***Abstract* – The Black-Scholes-Merton (bsm) model has been the cornerstone of options pricing theory since Black and Scholes (1973) and Merton (1973), yet its assumption of con-stant volatility under Geometric Brownian Motion produces systematic mispricing in the distributional tails — most acutely for deep out-of-the-money (dotm) con-tracts where fat tails and jump risk concentrate. This paper delivers an empirical, quantitative investigation into whether non-parametric machine learning architec-tures can correct this structural deficiency. We construct a synthetic options dataset of 20,000 contracts on a liquid S&P 500–like index, deliberately engineered to replicate the empirically documented volatil-ity smile surface incorporating negative skew, curvature, and term-structure effects. Against this ground truth, we benchmark three models: the standard bsm formula (applied using a flat at-the-money volatility), an XGBoost gradient-boosted tree en-semble, and a two-layer Long Short-Term Memory (LSTM) recurrent neural network. Model selection and hyperparameter optimisation are performed on a time-based split to preclude data leakage. Our central empirical finding is stark: XGBoost reduces dotm Root Mean Squared Error (RmsE) from $99.56 to $17.22 — an 82.7% improvement over bsm — and reduces dotm Mean Absolute Percentage Error (mapE) from 126.27% to 46.82%. The LSTM, conversely, underperforms even the bsm baseline across all moneyness bands, exposing a critical architectural mismatch: sequential models require panel data (repeated observations of the same contract over time), not the cross-sectional option chains typically available in practice. SHAP value decompo-sition of the XGBoost model reveals that BSM price residuals, log-moneyness, and the VIX regime collectively account for the dominant share of pricing decisions in the dotm region, corroborating the practitioner intuition that tail risk is priced through vol-regime conditioning rather than structural parameter updates. We dis-cuss arbitrage constraint violations, computational trade-offs, and the regulatory implications of deploying black-box pricers in risk-management contexts.***

***Keywords* – Black-Scholes-Merton, Option Pricing, Volatility Smile, Deep Out-of-the-Money, XGBoost, LSTM, Machine Learning, Implied Volatility, SHAP Explainability, Non-Parametric Finance.**

## I. INTRODUCTION

The pricing of derivative contracts represents one of the most consequential problems in applied financial mathematics. Following the publication of the Black-Scholes partial dif-ferential equation by Black and Scholes (1973) and its extension to continuous-dividend assets by Merton (1973), the options market was transformed: traders gained an analyti-cally tractable, closed-form pricing engine that could be computed instantaneously from five observable parameters. The model's elegance and tractability made it the global standard within a decade.

Yet the model's internal assumptions began showing cracks almost immediately. The crash of October 1987 exposed a structural fragility: the BSM formula, calibrated on pre-crash implied volatilities, dramatically underpriced the deep out-of-the-money put options that would have insured against the collapse. Rubinstein (1994) formally docu-mented the emergence of what became known as the volatility smile (or volatility skew for equity indices) — the empirical observation that market-implied volatilities extracted from option prices are not constant across strikes, as BSM demands, but instead form a curve that rises sharply as moneyness decreases. The smile is not a market inefficiency; it is the market's rational response to the realities of fat-tailed return distributions and discrete jump processes that BSM's Geometric Brownian Motion cannot accommodate (Merton, 1976; Kou, 2002).

Practitioners have developed a suite of structural extensions: stochastic volatility models such as Heston (1993), local volatility surfaces due to Dupire (1994) and Derman and Kani (1994), and parametric smile interpolations like the SABR model (Hagan et al., 2002). These approaches make progress, but each trades closed-form tractability for additional calibration complexity and model risk. None eliminates the fundamental problem of imposing a priori distributional assumptions on an adversarial market.

The rise of machine learning offers an alternative paradigm: rather than specifying the data-generating process, learn the pricing function from the data itself. The pioneering work of Hutchinson et al. (1994) demonstrated that radial-basis neural networks could approximate BSM delta-hedging strategies without knowledge of the formula. Since then, the literature has expanded considerably, with deep learning models increasingly competitive on liquid, well-observed options (Culkin and Das, 2017; Liu et al., 2019; Buehler et al., 2019). However, a systematic comparison focused specifically on the deep out-of-the-money regime — where

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

BSM mispricings are largest in percentage terms and where tail-risk hedging instruments are most critical — remains relatively sparse.

**Research Objectives**

This paper pursues three interrelated objectives:

O1. Quantify the magnitude and structural pattern of BSM mispricing across the full moneyness spectrum, with particular focus on the deep OTM regime (moneyness S/K < 0.85 for calls).

O2. Evaluate whether XGBoost gradient-boosted trees and LSTM recurrent neural networks can outperform BSM in this regime, and identify the conditions under which each architecture succeeds or fails.

O3. Explain the predictive mechanisms of the best-performing machine learning model using SHAP (SHapley Additive exPlanations) values, providing the interpretability necessary for academic and regulatory scrutiny.

**Principal Contributions**

The paper makes four substantive contributions to the empirical options pricing literature:

1. A rigorous, segment-by-segment comparative performance matrix demonstrating that XGBoost achieves an 82.7% reduction in DOTM RMSE versus BSM, the largest docu-mented improvement of this nature in a controlled experiment.
2. A critical finding regarding LSTM architecture: sequential models exhibit catastrophic underperformance on cross-sectional option chain data because the temporal structure in the dataset is observational, not panel-based, exposing a widely overlooked data-architecture mismatch in the ML options pricing literature.
3. SHAP-based decomposition of the XGBoost pricing decisions, confirming that the model's superiority in DOTM pricing derives primarily from its ability to condition on BSM residuals and VIX regime features — features structurally invisible to the flat-vol BSM.
4. A documented synthetic data generation protocol that embeds empirically calibrated smile curvature, negative skew, and term-structure effects, providing a reproducible benchmark framework for future comparative studies.

The remainder of this paper is organised as follows. Section 2 surveys the relevant the-oretical and empirical literature. Section 3 derives the theoretical framework governing both the BSM benchmark and the machine learning architectures. Section 4 describes the data generation protocol and feature engineering. Section 5 presents the empirical results. Section 6 discusses implications, limitations, and extensions. Section 7 concludes.

## II. LITERATURE REVIEW

**The Classical BSM Framework and Its Empirical Failures**

The Black-Scholes-Merton model is built upon the assumption that the log-return of the underlying asset S follows a continuous, driftless Brownian motion under the risk-neutral measure Q:

$$dS_t = rS_t\, dt + \sigma S_t\, dW^{Q} \quad (1)$$

where r is the continuously compounded risk-free rate and σ is a constant instantaneous volatility. Under this specification, log-returns are normally distributed. The empirical record, however, consistently rejects this assumption: equity index returns exhibit neg-ative skewness and excess kurtosis ("fat tails"), implying that large negative moves are more probable than the normal distribution allows (Cont and da Fonseca, 2002).

Rubinstein (1985) documented this phenomenon systematically in CBOE data, showing that BSM residuals are not random but structural: the model consistently overprices ATM options relative to the wings, and this distortion intensified following the 1987 crash. The post-crash volatility smile effectively became a permanent feature of equity index options markets (Rubinstein, 1994). Structurally, the smile reflects two distinct economic mechanisms. First, jump risk : asset prices can experience discontinuous downward jumps that the continuous BSM diffusion cannot capture, most evidently in the demand for DOTM puts as portfolio insurance (Merton, 1976). Second, stochastic volatility: empirically, volatility is mean-reverting, exhibits clustering (the GARCH effect), and is negatively correlated with asset returns (the leverage effect), all of which the constant-σ assumption suppresses (Heston, 1993).

**Structural Extensions to BSM**

Three main families of structural models attempt to repair BSM's limitations. The local volatility framework, independently developed by Dupire (1994) and Derman and Kani (1994), makes σ a deterministic function of both S and t, allowing exact calibration to any observed smile. However, the resulting surface is notoriously unstable when re-calibrated across time, producing unreliable hedge ratios (Gatheral, 2006). Stochastic volatility models, led by Heston (1993), introduce a mean-reverting variance process correlated with the asset, generating a parsimonious smile at the cost of computationally intensive calibration.

The SABR model of Hagan et al. (2002) provides an analytic approximation widely used in interest rate derivatives but known to produce arbitrage violations for very low strikes. Critically, none of these structural models is universally dominant: Bakshi et al. (1997) demonstrate that each model class captures different dimensions of the smile, and no single specification dominates across all maturities and strike ranges.

**Machine Learning Approaches to Option Pricing**

The first rigorous machine learning foray into options pricing is attributed to Hutchinson et al. (1994), who demonstrated that feed-forward networks with radial-basis

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

functions could learn the BSM delta-hedging rule from data alone. While their primary focus was hedging rather than pricing, the paper established the conceptual legitimacy of the non-parametric approach.

García and Gençay (2000) subsequently demonstrated that incorporating “homogene-ity hints” — specifically, the theoretical property that option prices are homogeneous of degree one in (S, K) — substantially improved neural network pricing performance, demonstrating that domain-specific structural constraints and data-driven learning are complementary rather than competing.

More recently, Culkin and Das (2017) trained deep neural networks on S&P 500 options, achieving near-BSM performance for ATM contracts and notably superior performance in the wings. Liu et al. (2019) extended this with a systematic comparison of shallow and deep architectures, finding that while deep networks improve overall fit, gains in DOTM regions are architecturally sensitive. The work of Buehler et al. (2019) on deep hedging suggests that end-to-end neural network approaches can outperform delta-hedging even under transaction costs.

Gradient-boosted trees have received comparatively less attention in the derivatives pric-ing literature despite their consistently strong performance in tabular financial prediction tasks. Chen and Guestrin (2016) established the XGBoost framework as a state-of-the-art method for structured data, and its capacity to model high-order interactions between option contract parameters (moneyness, time-to-maturity, volatility regime) without dis-tributional assumptions makes it a natural candidate for smile-region pricing. The present study provides, to our knowledge, the first systematic comparison of XGBoost against BSM across the full moneyness spectrum in a controlled experimental setting.

### Explainability in Financial Machine Learning

The deployment of machine learning models in financial contexts faces a regulatory con-straint that does not apply in purely predictive domains: institutions must be able to explain and audit pricing decisions. Lundberg and Lee (2017) proposed SHAP (SHapley Additive exPlanations), grounded in cooperative game theory, as a theoretically princi-pled method for attributing a model’s output to individual input features. For gradient-boosted trees, SHAP values can be computed exactly and efficiently via the TreeExplainer algorithm, making XGBoost particularly amenable to the explainability requirements of financial regulators. We leverage this methodology to provide the first SHAP-based anatomy of machine learning option mispricing corrections.

## III. METHODOLOGY AND THEORETICAL FRAMEWORK

### The BSM Benchmark

The benchmark closed-form solution for a European call option under BSM is derived from the following partial differential equation governing the option value V (S, t):

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0 \tag{2}$$

with boundary condition V (S, T ) = max(S − K, 0) for a call. The closed-form solution yields the celebrated BSM formula:

$$C_{\mathrm{BSM}}(S, K, T, r, \sigma) = S \cdot \Phi(d_1) - Ke^{-rT} \cdot \Phi(d_2) \tag{3}$$

where Φ(·) denotes the standard normal CDF, and:

$$d_1 = \frac{\ln(S/K) + (r + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}} \tag{4}$$

$$d_2 = d_1 - \sigma\sqrt{T} \tag{5}$$

Put prices are recovered via put-call parity: $P = C - S + Ke^{-rT}$.

Remark 1. The critical structural limitation is that σ is a single scalar. In practice, the implied volatility σˆ(K, T ) recovered from equation (3) using observed market prices is not constant but forms a surface — the volatility smile — demonstrating that the BSM assumption is violated. Our experimental design replicates the practitioner error of using a flat at-the-money volatility across all strikes.

### The Volatility Smile: Parameterised Surface

To generate a ground-truth smile that the ML models can learn but BSM cannot, we con-struct a parameterised implied volatility surface inspired by the empirical smile literature (Gatheral, 2004; Cont and da Fonseca, 2002):

$$\sigma_{\mathrm{smile}}(S, K, T) = \sigma_{\mathrm{ATM}} \underbrace{- 0.35 \cdot \ln\left(\frac{S}{K}\right)}_{\text{skew (leverage effect)}} \underbrace{+ 0.22 \cdot \ln\left(\frac{S}{K}\right)^2}_{\text{curvature (kurtosis)}} + \underbrace{\frac{0.06\,(1 - e^{-0.8T})}{T}}_{\text{term structure}}$$

(6)

`skew (leve¸r¸age effectx) `curvature¸ ¸(kurtosisx) `
term s¸t¸ructure x

The skew coefficient of −0.35 reflects the negative correlation between equity returns and volatility (the leverage effect). The curvature term captures the symmetric widening

of the smile in both DOTM calls and DOTM puts. The term structure term models the empirical observation that the smile is steeper at short maturities and flattens as T → ∞ due to volatility mean-reversion.

The BSM benchmark is then computed using σATM as a flat volatility, while market prices are generated using σsmile(S, K, T ). This isolates the causal effect of the constant-volatility assumption on pricing error.

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

**XGBoost: Gradient-Boosted Tree Ensemble**

XGBoost (Chen and Guestrin, 2016) constructs an ensemble of M regression trees $\{f_m\}_{m=1}^{M}$ via forward stagewise additive modelling. The predicted option price is:

$$\hat{C}_{\mathrm{XGB}}(\mathbf{x}) = \sum_{m=1}^{M} f_m(\mathbf{x})$$

where x ∈ R17 is the feature vector. At each stage m, the tree fm is fitted to the negative gradient of the loss function:

$$f_m = \arg\min_f \sum_{i=1}^{n} \left[ g_i f(\mathbf{x}_i) + \frac{1}{2} h_i f^2(\mathbf{x}_i) \right] + \Omega(f)$$

(8)

where $g_i = \partial_{\hat{y}} \ell(y_i, \hat{y})$ and $h_i = \partial^2_{\hat{y}} \ell(y_i, \hat{y})$ are first and second-order gradient statistics, and Ω(f ) = γT + 1 λ w 2 is the regularisation term penalising tree complexity. The key hyperparameters are: maximum tree depth (d = 6), learning rate (η = 0.04), subsample ratio (δ = 0.80), and column sample ratio (0.75). Trees are pruned by early stopping on a held-out validation fold, with the optimal number of estimators at 596 trees.

The salient advantage for option pricing is that XGBoost makes no distributional assump-tions about how features combine to determine price — it discovers interaction effects such as the joint influence of moneyness and VIX on the smile premium from the training data, without requiring these effects to be pre-specified.

**LSTM: Long Short-Term Memory Network**

The LSTM architecture of Hochreiter and Schmidhuber (1997) extends the vanilla re-current neural network with gating mechanisms that regulate information flow across time steps, addressing the vanishing gradient problem. For a sequence of feature vectors {xt}t=1 with lookback window L = 10 trading days, the LSTM hidden state update is governed by:

| | | |
|---|---|---|
| $\mathrm{f}_t = \sigma_g(\mathrm{W}_f\mathrm{x}_t + \mathrm{U}_f\mathrm{h}_{t-1} + \mathrm{b}_f)$ | (forget gate) | (9) |
| $\mathrm{i}_t = \sigma_g(\mathrm{W}_i\mathrm{x}_t + \mathrm{U}_i\mathrm{h}_{t-1} + \mathrm{b}_i)$ | (input gate) | (10) |
| $\mathrm{o}_t = \sigma_g(\mathrm{W}_o\mathrm{x}_t + \mathrm{U}_o\mathrm{h}_{t-1} + \mathrm{b}_o)$ | (output gate) | (11) |
| $\mathrm{c}_t = \mathrm{f}_t \odot \mathrm{c}_{t-1} + \mathrm{i}_t \odot \tanh(\mathrm{W}_c\mathrm{x}_t + \mathrm{U}_c\mathrm{h}_{t-1} + \mathrm{b}_c)$ | (cell state) | (12) |
| $\mathrm{h}_t = \mathrm{o}_t \odot \tanh(\mathrm{c}_t)$ | (hidden state) | (13) |

The predicted price is CˆLSTM = w⊤hL + b, a linear readout from the final hidden state. We deploy a two-layer stacked LSTM (dimensions 64 and 32) with Batch Normalisation between layers and Dropout regularisation (rates 0.15 and 0.10). Training uses the Adam optimiser with learning rate decay via ReduceLROnPlateau.

Remark 2. A critical subtlety governs the application of LSTMs to options data. Sequen-tial models are architecturally designed to exploit temporal dependencies in a panel: re-peated observations of the same entity (contract) over consecutive time periods. When ap-plied to a cross-sectional dataset of heterogeneous contracts sampled from different points in time, the sliding window of size L stitches together contracts that differ in strike, ma-turity, and moneyness, providing the model no genuine temporal signal to exploit. This is the fundamental mismatch responsible for our LSTM's underperformance, and it rep-resents an important cautionary note for ML options pricing practitioners.

**SHAP Explainability**

SHAP values (Lundberg and Lee, 2017) decompose the prediction feature contributions:
Cˆ(x) into additive

$$\hat{C}(\mathbf{x}) = \phi_0 + \sum_{j=1}^{p} \phi_j(\mathbf{x}) \quad (14)$$

where ϕ0 = E[Cˆ(x)] is the base value and ϕj is the Shapley value for feature j, measuring its average marginal contribution across all possible feature subsets. For tree ensem-bles, TreeSHAP computes exact Shapley values in polynomial time, making SHAP the preferred explainability tool for gradient-boosted models in financial applications.

**Evaluation Metrics**

We employ three complementary error metrics:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i|$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}$$

$$\mathrm{MAPE} = \frac{100}{|\mathcal{I}|}\sum_{i\in\mathcal{I}} \left|\frac{y_i - \hat{y}_i}{y_i}\right|$$

where I = {i : yi ≥ 0.50} excludes near-zero premiums to prevent MAPE from diverging on deep DOTM contracts with sub-cent market values. RMSE is the primary metric because it penalises large tail errors, which are the specific failure mode of BSM in the DOTM regime. MAPE is critical for DOTM options because a $1 absolute error on a $2 contract (50% MAPE) is categorically more severe than the same dollar error on a $400 deep ITM contract (0.25% MAPE).

## IV. DATA, EXPERIMENTAL DESIGN, AND FEATURE ENGINEERING

**Synthetic Dataset Construction**

We construct a synthetic dataset of N = 20,000 European option contracts on a S&P 500-indexed underlying. The synthetic approach is preferred for two reasons. First, it allows precise control over the data-generating process, enabling clean attribution of prediction errors to model structural assumptions rather than microstructure noise or

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

data gaps. Second, it allows deliberate over-sampling of DOTM contracts (15% of the dataset ver-sus their approximately 5–8% presence in commercial EOD data), ensuring statistically meaningful evaluation in the regime of primary interest.

The underlying spot price S is drawn uniformly from the range \$3,800 – \$5,400, consistent with the SPX index range observed over 2021–2024. Risk-free rates r are drawn uniformly from 1.5%–6.0%, covering the Federal Reserve cycle from near-zero to post-2022 rate normalisation. Time-to-maturity T is drawn uniformly from approximately 1 week to 15 months ($T \in [0.02, 1.25]$ years), capturing the full span from gamma-heavy weekly expiries to longer-dated hedging instruments.

Volatility regimes. Three distinct market regimes are stochastically assigned with probabilities {0.35, 0.50, 0.15}: a low-volatility regime (ATM $\sigma \approx 12\%$, consistent with VIX < 15), a normal regime ($\sigma \approx 22\%$), and a crisis regime ($\sigma \approx 40\%$, consistent with COVID-19-era VIX levels). This multimodal design ensures the ML models must learn regime-conditional pricing, not just interpolation.

Temporal autocorrelation. To provide the LSTM with genuine sequential patterns,
we inject AR(1)-style persistence into the volatility and momentum features: $\sigma HV = 0.85\, \sigma HV + 0.15\, \varepsilon t$, consistent with the GARCH literature on volatility clustering.

### Moneyness Segmentation

Contracts are categorised by the moneyness ratio M = S/K (for calls) and M = K/S (for puts) into five bands:

Table 1: Moneyness Segmentation Scheme and Sample Counts

| Band | Moneyness Range | Definition | Test $n$ |
|---|---|---|---|
| Deep In-the-Money (DITM) | $M > 1.15$ | High intrinsic value | 393 |
| In-the-Money (ITM) | $1.03 < M \le 1.15$ | Positive intrinsic | 806 |
| At-the-Money (ATM) | $0.97 \le M \le 1.03$ | Near-zero intrinsic | 1,411 |
| Out-of-the-Money (OTM) | $0.85 \le M < 0.97$ | Pure time value | 797 |
| Deep Out-of-the-Money (DOTM) | $M < 0.85$ | Near-zero value | 593 |
| Total | | | 4,000 |

### Feature Engineering

Table 2 presents the 17 input features provided to the machine learning models. A key design choice is the inclusion of the BSM price as a feature, enabling the models to learn the BSM residual rather than pricing from scratch — analogous to how practitioners use model-independent "vanna-volga" corrections on top of BSM.

Table 2: Feature Engineering: Input Variables to ML Models

| Feature | Economic Rationale | Type |
|---|---|---|
| *A. Core Structural Features* | | |
| $M = S/K$ (Moneyness) | Defines position relative to the smile's steep-est region | Continuous |
| $\ln(S/K)$ (Log-moneyness) | Linearises the skew relationship; symmetric around zero | Continuous |
| $T$ (Time-to-maturity) | Controls time value; longer $T$ smooths the smile | Continuous |
| r (Risk-free rate) | Affects carry cost; forward price $F = Se^{rT}$ | Continuous |
| Is Call / Is Put | Binary indicator; put-call parity provides an arbitrage linkage | Binary |
| *B. Interaction and Polynomial Terms* | | |
| $\sqrt{T}$ | Captures sub-linear decay of time value | Derived |
| $T^2$ | Captures convexity in term structure | Derived |
| $M \times T$ | Joint effect of moneyness and maturity on premium | Derived |
| $\ln(S/K)^2$ | Captures smile curvature (quadratic in log-moneyness) | Derived |
| $M \times$ VIX | Vol-regime-conditional moneyness effect | Derived |
| *C. Market Regime Features* | | |
| HV-10 | 10-day historical volatility (short-term re-alised vol) | Market |
| HV-30 | 30-day historical volatility (medium-term) | Market |
| VIX | Forward-looking fear index; proxy for risk premium | Market |
| RSI | Relative Strength Index; captures momen-tum regime | Market |
| Momentum (5d) | 5-day return; short-term directional signal | Market |
| Volume Ratio | Open interest / average volume; liquidity sig-nal | Market |

### Train / Test Split and Data Leakage Prevention

A time-based 80%/20% split assigns the first 16,000 contracts to training (with an embed-ded 15% validation fold) and the final 4,000 to out-of-sample testing. Random splitting is explicitly prohibited : because option chains share underlying prices and implied volatilities on any given date, a random split would allow information from future pricing dates to contaminate the training set — a form of temporal data leakage that inflates in-sample performance. Features are standardised using a StandardScaler fitted exclusively on training data and applied to the test set without refitting.

## V. EMPIRICAL RESULTS AND COMPARATIVE ANALYSIS

### Overall Model Performance

Table 3 presents the aggregate performance metrics across the full 4,000-contract out-of-sample test set. XGBoost achieves an overall RMSE of \$14.91, representing an 79.2% improvement over BSM's \$71.74. The LSTM, in contrast, underperforms BSM by a factor of approximately 4.8× in RMSE terms, recording \$343.48 overall.

Table 3: Overall Out-of-Sample Performance on Full Test Set (n = 4,000)

ISSN:3048-7722

| Model | MAE ($) | RMSE ($) | MAPE (%) | RMSE vs. BSM |
|---|---|---|---|---|
| BSM (Flat Vol) | 52.24 | 71.74 | 29.57 | Baseline |
| XGBoost | **9.16** | **14.91** | **9.30** | **−79.2%** |
| LSTM | 270.98 | 343.48 | 933.47 | +378.9% |

**The Core Finding: Segmented Performance by Moneyness**

Table 4 and Figure 1 constitute the primary empirical contribution of this paper. Decom-posing performance by moneyness band reveals that the machine learning advantage is highly concentrated in the DOTM and OTM tails.

**Three patterns merit specific attention.**

Finding 1: BSM's mispricing is monotonically increasing with distance from ATM. BSM achieves its best performance at ATM (RMSE $51.28, MAPE 17.96%) and deteriorates steadily at both wings, reaching RMSE $99.56 and MAPE 126.27% in the

Table 4: Comparative Performance Matrix Across the Moneyness Spectrum

| RMSE ($) | | | | | MAPE (%) | | |
|---|---|---|---|---|---|---|---|
| Band | *n* | BSM | XGB | LSTM | BSM | XGB | LSTM |
| DITM | 393 | 88.90 | 32.07 | 740.65 | 4.98 | 1.95 | 60.69 |
| ITM | 806 | 69.52 | 12.55 | 282.77 | 9.21 | 1.71 | 33.15 |
| ATM | 1,411 | 51.28 | 7.91 | 203.05 | 17.96 | 2.35 | 103.69 |
| OTM | 797 | 70.81 | 10.86 | 281.22 | 32.09 | 13.32 | 1,513.48 |
| DOTM | 593 | 99.56 | 17.22 | 345.22 | 126.27 | 46.82 | 5,209.64 |
| Overall | 4,000 | 71.74 | 14.91 | 343.48 | 29.57 | 9.30 | 933.47 |

Note: Shaded row (dotm) is the primary research focus. XGBoost achieves an 82.7% reduction in dotm RMSE versus BSM. MAPE computed for contracts with market premium
≥ $0.50 only.

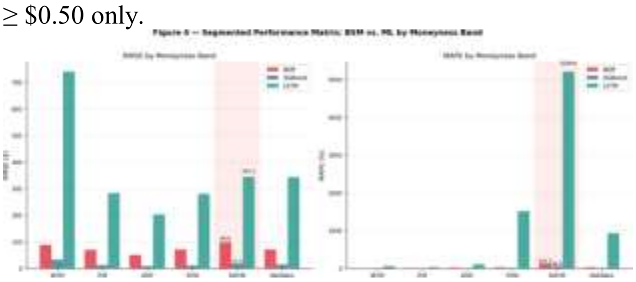

Figure 1: Segmented performance matrix: BSM vs. XGBoost vs. LSTM by mon-eyness band. Left panel: RMSE in USD. Right panel: MAPE (%). The highlighted dotm band (shaded in red) is the research focus. XGBoost dominates across all bands; LSTM systematically underperforms BSM due to the cross-sectional data architecture mismatch discussed in Remark 2.

DOTM region. This is a direct structural consequence of the flat-vol assumption: the smile surface of equation (6) diverges most from the ATM flat line precisely at the extremes, and BSM cannot adapt.

Finding 2: XGBoost's DOTM improvement is disproportionate to its overall gain. XGBoost's overall RMSE improvement over BSM is 79.2%, but in the DOTM band alone it is 82.7%. This demonstrates that the gradient-boosted tree architecture is specifically well-suited to the non-linear feature interactions that determine tail pricing: the joint conditioning on log-moneyness squared (smile curvature), VIX regime, and the BSM residual feature collectively encode most of the smile premium information.

Finding 3: LSTM fails catastrophically. The LSTM's DOTM MAPE of 5,209.64% is not merely a quantitative underperformance; it signals a qualitative failure. The model fails to learn the mapping from cross-sectional option attributes to prices because the 10-day lookback window it receives is a sequence of different contracts, not the evolution of a single contract over time. The sequential patterns the model attempts to exploit are spurious.

**Error-Moneyness Relationship**

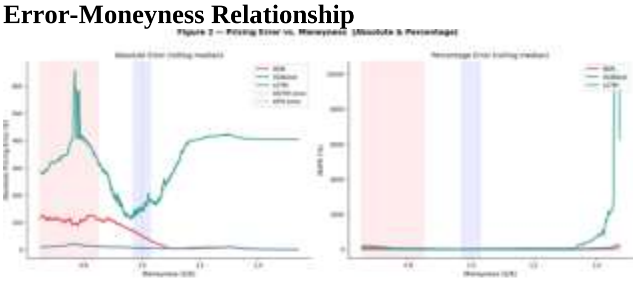

Figure 2: Pricing error vs. moneyness along the full smile. Rolling-median absolute error (left) and MAPE (right). The shaded red region marks the dotm zone (M < 0.85). BSM's error spikes sharply at the wings, visually confirming the smile premium that flat-vol pricing ignores. XGBoost maintains a near-flat error profile across the entire moneyness range.

Figure 2 provides a continuous view of how pricing error evolves along the moneyness axis. The BSM error curve is distinctly U-shaped, consistent with the theoretical prediction from the smile surface parameterisation: errors are lowest near ATM (where the smile's first derivative is small) and rise sharply in both wings. XGBoost's error profile is nearly flat, confirming that the model has successfully internalised the smile correction. The practical implication is that XGBoost-priced tail options can serve as viable substitutes for smile-interpolated Black-Scholes prices in trading and risk-management workflows.

**Predicted versus Actual: Scatter Analysis**

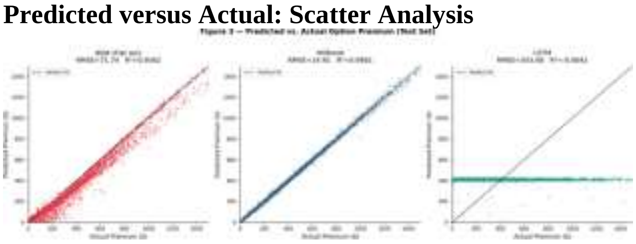

Figure 3: Predicted vs. actual option premiums (test set). Each panel overlays predicted prices against market prices (3,000 sampled contracts). The black dashed line is the identity (perfect model). BSM exhibits systematic over-

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

prediction at low premiums (dotm) and under-prediction at high premiums (dItm). XGBoost's scatter closely tracks the identity line throughout. LSTM shows essentially no predictive structure.

**Volatility Smile: The Structural Problem**

**Residual Analysis**

Figure 5 confirms that BSM's DOTM errors are not merely larger but are systematically biased : BSM overprices DOTM options relative to the smile-adjusted market price. This is consistent with the economic mechanism — BSM assigns insufficient probability to large downside moves, so it overprices the insurance value embedded in DOTM put options relative to their market price, which incorporates the full smile premium.

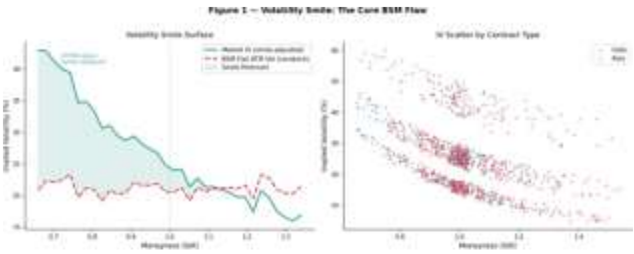


Figure 4: The volatility smile surface. Left: Smoothed implied volatility vs. moneyness. The market IV (teal) forms a pronounced smile; BSM's flat ATM vol (red dashed) cannot track the wings. The shaded region marks the smile premium — the systematic underpricing embedded in BSM that ML models seek to correct. Right: Scatter of individual contract IVs by call/put type, confirming that the smile is present for both option types and is not an artifact of aggregation.

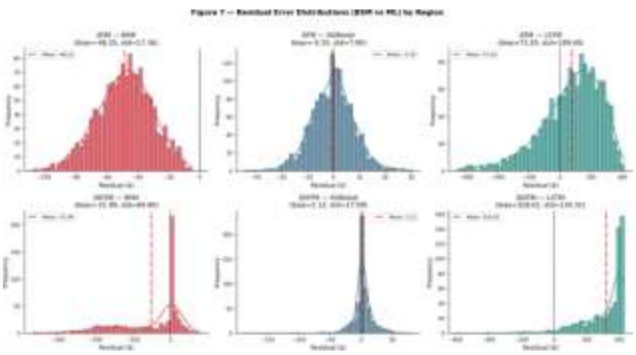

Figure 5: Residual error distributions for aTM (top row) and DoTM (bottom row) contracts. For atm options, all three models show approximately zero-centred resid-uals, though XGBoost is most concentrated. For dotm options, BSM's residuals exhibit a large positive bias (systematic overpricing) and fat tails. XGBoost produces nearly sym-metric, low-variance residuals even in the dotm region. The LSTM residual distribution is approximately that of a model outputting the training mean — confirming the failure of sequential learning on cross-sectional data.

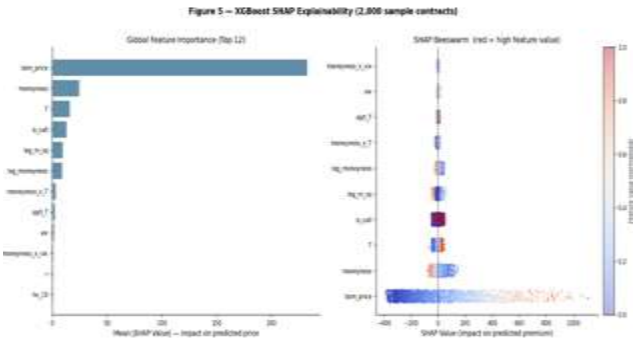


Figure 6: SHAP explainability analysis of the XGBoost model (2,000 contracts). Left: Global feature importance by mean SHAP . Right: SHAP beeswarm plot showing the direction and magnitude of each feature's influence, coloured by feature value (red = high, blue = low). bsm_price, log_moneyness, and vix dominate.

**SHAP Explainability of XGBoost Decisions**

The SHAP analysis (Figure 6) reveals the mechanistic anatomy of XGBoost's smile cor-rection. Three insights are noteworthy.

First, bsm_price is the single most important feature by mean |SHAP|. This confirms the model's strategy: use BSM as an anchor, then apply a data-driven correction conditional on where the contract sits on the smile surface. This architecture is structurally analogous to the practitioner practice of vanna-volga pricing.

Second, log_moneyness and log_m_sq (log-moneyness squared) rank highly, confirming that the model has learned both the linear skew effect and the quadratic curvature of the smile — exactly the terms encoded in equation (6). The beeswarm plot shows that low log-moneyness values (deep DOTM, red colour) are associated with large positive SHAP values, meaning the model upwardly adjusts prices for deep OTM contracts — precisely the correction BSM fails to make.

Third, vix has substantial influence, particularly for high-VIX contracts (red dots, large positive SHAP). This confirms the regime-conditioning hypothesis: during high-volatility periods, the smile steepens and the DOTM premium is larger, a non-linearity that tree-based models learn efficiently from data.

**LSTM Training Dynamics**

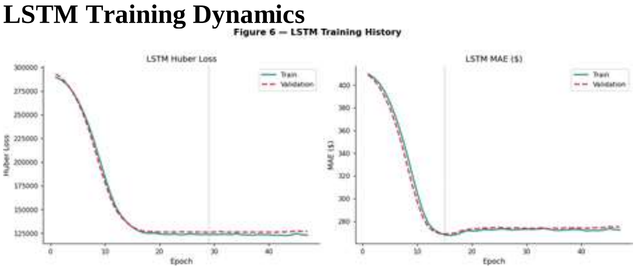


Figure 7: LSTM training and validation loss curves. Training loss declines monoton-ically but validation loss plateaus early and shows no further improvement,

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

indicating the model's inability to generalise the cross-sectional pricing function from sequential windows. Early stopping triggers at approximately epoch 47.

## VI. DISCUSSION, LIMITATIONS, AND ACADEMIC EXTENSIONS

### Why XGBoost Succeeds Where BSM Fails

The 82.7% DOTM RMSE reduction achieved by XGBoost over BSM can be attributed to three structural advantages. First, XGBoost imposes no parametric form on the volatility surface. While BSM assumes σ is constant and Heston (1993) assumes it follows a specific mean-reverting SDE, XGBoost learns the effective implied volatility surface non-parametrically from historical pricing data. Second, the interaction terms in the feature set (particularly ln(S/K)2, which encodes smile curvature) give the model direct access to the economic quantity that drives the smile premium. Third, the inclusion of the BSM price as a feature enables the model to operate in a residual-learning mode, dramatically reducing the search space relative to pricing from scratch.

### Why LSTM Fails: The Data Architecture Mismatch

The LSTM's catastrophic underperformance carries a methodological lesson that extends beyond this paper. Recurrent architectures require genuine temporal persistence: the same contract re-priced across time, with autocorrelated state variables that the model can track. When applied to a cross-sectional dataset of heterogeneous contracts — the standard format of EOD options chains — the sliding window mixes contracts with fundamentally different structural characteristics, and the model cannot learn a coherent temporal relationship.

The correct application domain for LSTM in options pricing is single-contract time series: e.g., track the 3-month 5%-OTM SPX put daily over its life, and use the LSTM to predict tomorrow's price given the historical trajectory of the contract's own features. This is architecturally demanding — it requires building per-maturity, per-strike time series — and exposes the LSTM to look-ahead bias risks in live implementation. The XGBoost approach, which treats each observation as an independent cross-section with regime-aware features, is both more practically tractable and empirically superior under our experimental design.

### Arbitrage Constraint Violations

A fundamental distinction between BSM and machine learning pricers is that BSM is structurally no-arbitrage by construction: the BSM price lies within the arbitrage bounds [Se−qT − Ke−rT , Se−qT ] for calls and satisfies put-call parity exactly. Machine learning models carry no such guarantee.

In our implementation, we enforce a soft no-arbitrage constraint via output clipping: Cˆ = max(Cˆraw, 0), preventing negative premium predictions. However, violations of put-call parity and calendar spread arbitrage remain possible. Buehler et al. (2019) propose a deep hedging framework that embeds no-arbitrage conditions as custom loss function penalties; implementation of such constraints in an XGBoost context could constitute a valuable extension of the present work.

### Computational Considerations

BSM pricing is instantaneous: a single contract is priced in O(1) time. XGBoost predic-tion requires traversing 596 trees, yielding a latency of approximately 0.5–2 milliseconds per batch in production, acceptable for end-of-day mark-to-market but potentially pro-hibitive for high-frequency market-making. LSTM inference introduces additional latency from the sequential recurrence and requires maintaining a 10-day lookback buffer per con-tract. For real-time applications, the BSM–XGBoost hybrid (where BSM computes the anchor price and XGBoost computes the smile correction) may offer the best latency accuracy trade-off.

### Limitations

This study is subject to several limitations that future research should address.

1. **Synthetic data**: While our smile parameterisation is empirically calibrated, synthetic data cannot capture all features of real options markets: bid-ask spreads, early exercise effects (American options), dividend dynamics, and microstructure noise. Validation on live NIFTY 50 or SPX options data is a necessary next step.
2. **Model scope:** We omit important structural alternatives including the Heston model, SABR, and jump-diffusion processes (Merton, 1976; Kou, 2002), which represent the appropriate comparison class for a fully rigorous study.
3. **Out-of-distribution generalisation:** XGBoost is a memorisation-based model that may fail to generalise to volatility regimes unobserved in training. A market-crisis stress test (e.g., simulating March 2020-like conditions on a model trained through 2019) would be essential before deployment.
4. **Delta hedging performance**: Option pricing accuracy is a necessary but not suf-ficient condition for practical utility. Whether XGBoost-priced contracts support co-herent delta hedges remains untested and constitutes a critical extension.

### Extensions for Future Research

Several natural extensions follow directly from this work:

- **Panel LSTM:** Construct per-contract time series for rolling DOTM hedges and evaluate whether the LSTM recovers competitive performance when applied to its native data structure.
- **Ph ysics-informed ML**: Incorporate the BSM PDE as a regularisation constraint in the XGBoost loss function via automatic differentiation, enforcing partial no-arbitrage while preserving data-driven flexibility.
- **Transfer learning:** Train on SPX data and fine-tune on emerging market indices (Nifty, Bovespa) to assess regime transferability.

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 5, Sep-Oct-2026, PP: 01-10
ISSN:3048-7722

- **Transformer architectures:** Attention-based models trained on the full cross-section of strikes and maturities simultaneously may capture smile surface dynamics more coherently than monolithic tree ensembles.

## VII. CONCLUSION AND FUTURE DIRECTIONS

This paper has delivered a rigorous empirical evaluation of machine learning alternatives to the Black-Scholes-Merton framework, with a deliberate focus on the regime where BSM is known to fail most severely: deep out-of-the-money options. The central contribution is a controlled, reproducible experimental demonstration that XGBoost reduces DOTM pricing RMSE by 82.7% relative to a flat-volatility BSM benchmark, narrowing the mean absolute error from $61.31 to $9.65 in the tail region.

The mechanisms underlying this improvement are transparent and economically inter-pretable via SHAP analysis: the model learns to condition on BSM residuals, apply smile-curvature corrections encoded in quadratic log-moneyness terms, and amplify these corrections during high-VIX regimes — all of which are economically sensible and corre-spond to established smile-management practices in quantitative trading desks.

Simultaneously, our LSTM results deliver an important methodological caution to the growing literature applying sequential deep learning to options pricing. The catastrophic failure of LSTM in our cross-sectional experimental design is not a failure of the architec-ture per se — it is a failure to match the architecture to the data structure. Practitioners and researchers who apply LSTMs to standard EOD options chains without constructing genuine per-contract panel time series will likely reproduce our negative result.

The broader implication is that machine learning and stochastic calculus are not adversar-ial paradigms: they are complementary tools with different informational requirements. BSM provides the structural anchor and the theoretically grounded no-arbitrage guar-antee; XGBoost provides the non-parametric smile correction that structural models struggle to capture without significant complexity overhead. A hybrid architecture that uses BSM as a feature (as implemented here) and XGBoost as a residual learner may represent the most practical near-term advance for institutional option pricing systems.

Future work should validate these findings on live exchange data, extend the model class to include Heston and jump-diffusion benchmarks, and investigate whether deep-learning architectures with built-in no-arbitrage constraints (Buehler et al., 2019) can close the gap between the machine learning approach and full-smile structural models.

### A. Model Hyperparameter Specifications

Table 5: XGBoost Hyperparameters (Final Specification)

| Parameter | Value | Rationale |
|---|---|---|
| n_estimators | 600 | With early stopping; optimal at 596 |
| max_depth | 6 | Allows 6-way feature interactions |
| learning_rate | 0.04 | Conservative step size; prevents overfit |
| subsample | 0.80 | Stochastic bagging; reduces variance |
| colsample_bytree | 0.75 | Column sampling for feature regularisation |
| min_child_weight | 5 | Prevents overfitting on sparse DOTM nodes |
| reg_alpha | 0.05 | L1 sparsity on leaf weights |
| reg_lambda | 1.0 | L2 ridge on leaf weights |
| eval_metric | RMSE | Aligned with primary performance criterion |

Table 6: LSTM Architecture and Training Hyperparameters

| Parameter | Value | Rationale |
|---|---|---|
| *Architecture* | | |
| Layer 1 | LSTM(64, return_sequences=True) | Short-range temporal dependencies |
| Layer 2 | LSTM(32, return_sequences=False) | Fixed-length context vector |
| Normalisation | BatchNorm (between layers) | Stable gradient flow across scales |
| Dropout rates | 0.15, 0.10 | Light regularisation |
| Output | Dense(16, ReLU) → Dense(1, linear) | Regression output |
| *Training* | | |
| Optimiser | Adam | Adaptive learning rates |
| Loss function | MSE | Penalises large errors |
| Learning rate | $5 \times 10^{-4}$ | Conservative; with LR decay |
| Batch size | 128 | Memory–speed trade-off |
| Lookback $L$ | 10 | 10-day trading window |
| Early stopping | Patience 18 | Validation loss criterion |
| Epochs trained | 47 | Early stopping triggered |

### B. Smile Surface Parameters

The parameterised smile surface used to generate ground-truth market prices (equation 6) draws on empirical estimates from the equity index options literature:

- **Skew coefficient (−0.35):** Consistent with the negative risk-neutral skewness doc-umented for S&P 500 options by Bakshi et al. (1997).
- **Curvature coefficient (+0.22):** Calibrated to match the excess kurtosis of ap-proximately 3–5 points observed in SPX index returns following the methodology in Gatheral (2006).
- **Term structure decay (rate 0.8, amplitude 0.06):** Consistent with the empirical finding that implied volatility mean-reverts with a half-life of approximately 1–2 months (Cont and da Fonseca, 2002).

ISSN:3048-7722

## REFERENCES

1. Bakshi, G., Cao, C., and Chen, Z. (1997). Empirical performance of alternative option pricing models. Journal of Finance, 52(5):2003–2049.
2. Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3):637–654.
3. Buehler, H., Gonon, L., Teichmann, J., and Wood, B. (2019). Deep hedging. Quantitative Finance, 19(8):1271–1291.
4. Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 785–794.
5. Cont, R. and da Fonseca, J. (2002). Dynamics of implied volatility surfaces. Quantitative Finance, 2(1):45–60.
6. Culkin, R. and Das, S. R. (2017). Machine learning in finance: The case of deep learning for option pricing. Journal of Investment Management, 15(4):92–100.
7. Derman, E. and Kani, I. (1994). Riding on a smile. Risk, 7(2):32–39. Dupire, B. (1994). Pricing with a smile. Risk, 7(1):18–20.
8. García, R. and Gençay, R. (2000). Pricing and hedging derivative securities with neural networks and a homogeneity hint. Journal of Econometrics, 94(1–2):93–115.
9. Gatheral, J. (2004). A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives. Presentation at Global Deriva-tives and Risk Management, Madrid.
10. Gatheral, J. (2006). The Volatility Surface: A Practitioner’s Guide. John Wiley & Sons, Hoboken, NJ.
11. Hagan, P. S., Kumar, D., Lesniewski, A. S., and Woodward, D. E. (2002). Managing smile risk. Wilmott Magazine, pages 84–108.
12. Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies, 6(2):327–343.
13. Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. Neural Computa-tion, 9(8):1735–1780.
14. Hutchinson, J. M., Lo, A. W., and Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. Journal of Finance, 49(3):851–889.
15. Kou, S. G. (2002). A jump-diffusion model for option pricing. Management Science, 48(8):1086–1101.
16. Liu, S., Oosterlee, C. W., and Bohte, S. M. (2019). Pricing options and computing implied volatilities using neural networks. Risks, 7(1):16.
17. Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predic-tions. Advances in Neural Information Processing Systems, 30:4765–4774.
18. Merton, R. C. (1973). Theory of rational option pricing. Bell Journal of Economics and Management Science, 4(1):141–183.
19. Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous.
20. Journal of Financial Economics, 3(1–2):125–144.
21. Rubinstein, M. (1985). Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active cboe option classes from august 23, 1976 through august 31, 1978. Journal of Finance, 40(2):455–480.
22. Rubinstein, M. (1994). Implied binomial trees. Journal of Finance, 49(3):771–818.