

OPTION PRICING UNDER THE LIFTED HESTON MODEL WITH JUMPS

Nithit BORVORNLUCK

Department of Banking and Finance, Faculty of Commerce and Accountancy,
Chulalongkorn University, Thailand; nitthit2000@gmail.com

ARTICLE HISTORY

Received: 3 July 2026

Revised: 7 July 2026

Published: 22 July 2026

ABSTRACT

This study develops and evaluates the Lifted Heston with Jumps (LHJ) model, a finite-dimensional Markovian option-pricing framework that augments the Lifted Heston approximation of rough volatility with asymmetric double-exponential jumps. The moment generating function of the log price is derived in closed transform form by combining the lifted Riccati system with the Kou jump transform, and option prices are recovered through Laplace inversion. Six nested specifications, namely Heston, Bates, rough Heston, rough Heston with jumps, Lifted Heston and Lifted Heston with jumps, are calibrated in implied-volatility space to S&P 500 index options spanning January 2017 to December 2023, with parameters estimated each Wednesday and evaluated out of sample on the following trading days. Errors are disaggregated jointly by moneyness and maturity and by bull, bear, calm and turbulent market regimes. The results show that the advanced models earn their accuracy on the put surface, where the heavy left tail of the index carries pronounced skew. The proposed model attains the lowest out-of-sample implied-volatility root mean squared error of the six models on both surfaces and holds the put lead across the bull, bear and turbulent regimes, because roughness corrects the at-the-money skew while the asymmetric jump supplies the short-maturity crash tail. It reaches this accuracy in roughly one quarter of the calibration time of the exact rough jump model, so the lifting delivers the accuracy of rough volatility with jumps at close to conventional cost.

Keywords: Rough Volatility, Lifted Heston Model, Jump Diffusion, Option Pricing, S&P 500 Index Options

CITATION INFORMATION: Borvornluck, N. (2026). Option Pricing Under the Lifted Heston Model with Jumps. *Procedia of Multidisciplinary Research*, 4(7),54.

INTRODUCTION

In derivatives markets, both retail and institutional investors rely on pricing models to determine the fair value of financial contracts. Accurate valuation is essential for ensuring market efficiency. To address the challenge of consistent and tractable option valuation, Black and Scholes (1973) proposed a groundbreaking model that laid the foundation for modern option pricing theory. The Black–Scholes model assumes that asset prices follow a geometric Brownian motion with constant volatility, allowing for a closed-form analytical solution for European option prices and establishing a cornerstone in quantitative finance.

Despite its analytical tractability, the Black–Scholes framework is misspecified. In reality, asset returns display non-Gaussian distributions with heavy tails (Mandelbrot, 1963; Fama, 1965), price discontinuities resulting from sudden jumps (Merton, 1976), and roughness in volatility, evidenced by slowly decaying autocorrelations (Ding et al., 1993; Cajueiro & Tabak, 2007). These stylized facts give rise to the well-known volatility smile, as well as persistent skewness and a complex term structure of implied volatility, all of which a constant-volatility, continuous-path model like Black–Scholes fails to capture.

To address the empirical evidence of non-constant and rough behavior in volatility, El Euch and Rosenbaum (2019) introduced the rough volatility framework, most notably the Rough Heston (RH) model. This model replaces the Markovian variance dynamics of the classical Heston model with a fractional kernel, allowing it to capture the short-maturity implied volatility smiles and the persistent roughness observed in realized volatility. At the same time, the Rough Heston model preserves transform-based analytical tractability for pricing vanilla options, making it both theoretically sound and practically applicable.

Despite its strong empirical performance, the Rough Heston (RH) model is computationally demanding due to its non-Markovian and non-local fractional kernel, which complicates calibration and simulation in practical applications. To overcome this limitation, the Lifted Heston (LH) model was proposed as a finite-dimensional Markovian approximation of the Rough Heston (RH) model by expressing the fractional kernel as a finite sum of exponentials. The Lifted Heston (LH) model preserves the essential roughness of volatility while restoring Markovian tractability (Abi Jaber, 2019). This approximation allows for efficient numerical implementation, faster calibration, and compatibility with conventional pricing frameworks.

While the Lifted Heston (LH) model provides a computationally efficient and realistic description of volatility dynamics, it may still struggle to accurately capture sudden market movements and extreme short-term fluctuations. To address these limitations, this study introduces a jump component from Kou (2002) into the Lifted Heston framework, leading to the development of the Lifted Heston with Jumps (LHJ) model. This extension captures discontinuities and heavy tails in returns, sharp short-maturity wings, and downside skew driven by crash risk. In doing so, Lifted Heston with Jumps (LHJ) model preserves Lifted Heston model's rough volatility structure while incorporating jump-driven discontinuities in returns, delivering a more comprehensive and practical framework for option pricing and risk management.

The objective of this study is to assess the performance of the proposed Lifted Heston with Jumps (LHJ) model relative to established benchmarks, namely Heston, Rough Heston (RH), Heston with Jump (Bates), Rough Heston with Jump (RHJ), and Lifted Heston (LH). Evaluation proceeds along two dimensions. First, accuracy measured by the deviation between implied volatility generated by model price and market-observed implied volatilities across strikes and maturities. Second, computational efficiency, measured by total calibration time required to fit each model. The empirical analysis uses S&P 500 (SPX) index options from 2017 to 2023, covering maturities from 7 days to 1 year, and spans the full moneyness range. Results are disaggregated jointly by moneyness, out of the money (OTM), at the money (ATM), in the money (ITM) and by maturity buckets, to reveal model performance across both dimensions.

LITERATURE REVIEWS

The Black-Scholes model (Black & Scholes, 1973) remains one of the most influential frameworks in the field of option pricing and hedging. However, its practical applicability has been widely questioned due to several unrealistic assumptions. For example, volatility clustering, roughness, and sharp short-maturity smiles.

To capture these features, Heston model was introduced, relaxing the constant-volatility assumption by allowing variance to follow a mean-reverting stochastic process (Heston, 1993). This framework can reproduce volatility clustering and implied-volatility skew/smile while remaining analytically tractable through its closed-form characteristic function.

Moreover, Gatheral et al. (2018) find that volatility is rough, with a low Hurst parameter, indicating short-range power-law dependence rather than smooth, memoryless dynamics. Rough-volatility models were introduced to capture this roughness, enabling them to reproduce key market features, particularly pronounced short-maturity implied-volatility smiles and a realistic term structure of skew.

The Rough Heston (RH) model bridges the gap between the classical Heston framework and Rough volatility model by replacing the Brownian-driven variance process with a rough volatility component (El Euch & Rosenbaum, 2019). Incorporating the empirically observed roughness of volatility while preserving semi-analytical tractability. The Rough Heston (RH) model successfully reproduces the entire implied volatility surface, the steep short-term skews, and captures the roughness of volatility in a parsimonious and theoretically well-founded manner. Despite, its strong empirical performance, the Rough Heston (RH) model is burdensome in practice. The non-Markovian fractional kernel turns pricing into a fractional Riccati–Volterra problem with history-dependent convolutions, leading to high computational cost.

To mitigate these computational burdens, the Lifted Heston (LH) model of Abi Jaber (2019) approximates the Rough Heston fractional kernel by a finite sum of exponentials, producing a K -dimensional Markovian variance state (a weighted sum of OU factors) that closely emulates rough behavior while restoring affine tractability. LH typically reproduces Rough Heston’s short-maturity skew and term-structure features at a fraction of the run-time.

While stochastic volatility explains gradual smile dynamics, equity returns also exhibit discontinuities that diffusions alone understate. For example, earnings surprises, macro shocks, and crash risk. Jump–diffusion models address this by adding Poisson jump component on the price (Merton, 1976), improving tail behavior and far-OTM pricing. In the Heston family, adding jumps (Bates, 1996) preserves characteristic-function tractability and materially enhances fit. Kou (2002) introduced double-exponential jumps which tighten errors in the tails of the smile and at short maturities.

Building on Kou’s framework, we extend the Lifted Heston model with jumps (LHJ) to combine dynamics that capture short-maturity behavior with improved tail fit, aiming for Rough Heston with Jump (RHJ)-level accuracy at near-Lifted Heston (LH) calibration cost.

In conclusion, in the existing literature the Rough Heston (RH) model and Rough Heston with Jump (RHJ) excel in fit but are costly, while Lifted Heston (LH) is fast but may miss jump-driven tails. This motivates our Lifted Heston with Jumps (LHJ), which aims to retain Lifted Heston (LH) model’s speed while capturing tail risk and pronounced short-maturity wings through a double-exponential jump component.

RESEARCH METHODOLOGY

This section outlines the modeling architecture used in our study. We begin with the Heston (1993) stochastic-volatility benchmark, add price jumps via the Bates (1996) extension to accommodate discontinuities, and then incorporate rough features through the Rough Heston model of El Euch and Rosenbaum (2019). To jointly capture abrupt price moves and rough volatility, we consider a Rough Heston with Jump (RHJ) specification. Because fractional, non-Markovian dynamics

complicate calibration and computation, we adopt the Lifted Heston (LH) framework of Abi Jaber (2019), which approximates the rough kernel with a finite sum of exponentials while preserving tractability. Finally, we extend LH with Kou's (2002) double-exponential jumps (LHJ) to obtain a flexible, computationally efficient platform for pricing and hedging under stochastic volatility with discontinuous shocks.

Heston model

The Heston (1993) model introduces stochastic volatility by allowing the instantaneous variance to follow a mean-reverting square-root process. Two correlated standard Brownian motions B_t and W_t with constant correlation ρ drive the asset price and the variance respectively. Under the risk-neutral measure $\mathbb{Q}$ the asset price S_t follows

$$dS_t = rS_t dt + S_t \sqrt{V_t} dB_t \quad (1)$$

where S_t is the asset price at time t and r is the risk-free interest rate, the drift of the asset under the risk-neutral measure. The variance process $\sqrt{V_t}$ follows

$$dV_t = \kappa(\theta - V_t)dt + \nu\sqrt{V_t}dW_t, \quad V_0 \geq 0 \quad (2)$$

where V_0 is the initial variance, κ is the speed of mean reversion, θ is the long-run variance level and ν is the volatility of variance. The moment generating function of the log price, used by the Laplace inversion method, is available in closed form and is given in Albrecher et al. (2007).

Bates model

The Bates (1996) model extends the Heston framework by adding jumps that arrive through an independent Poisson process N_t with intensity λ and have lognormal size G , whose logarithm $Y = \ln(G)$ is normally distributed with mean μ and variance δ^2 . The moment generating function of the log price retains the closed form of the Heston model with an additional jump term and its compensator, and is given in Bates (1996).

Rough Heston model

The Rough Heston model of El Euch and Rosenbaum (2019) replaces the Brownian-driven variance of the classical Heston model with a fractional kernel, allowing the volatility to exhibit the rough behavior observed in financial markets. The price dynamics keep the form of equation (1), and the variance process V_t follows

$$V_t = V_0 + \frac{1}{\Gamma(H + \frac{1}{2})} \left[\int_0^t (t-s)^{H-\frac{1}{2}} \kappa(\theta - V_s) ds + \int_0^t (t-s)^{H-\frac{1}{2}} \nu \sqrt{V_s} dW_s \right], \quad V_0 \geq 0 \quad (3)$$

where V_0 , κ , θ and ν keep the meanings they have in the Heston model. The fractional kernel $(t-s)^{H-\frac{1}{2}}$ makes the variance rough, with the Hurst exponent H controlling the degree of roughness. H lies between 0 and 0.5, smaller values produce rougher variance paths, and the Gamma function $\Gamma(H + \frac{1}{2})$ normalizes the kernel, which is otherwise not directly integrable.

Following El Euch and Rosenbaum (2019), the moment generating function of the log price, used by the Laplace inversion method, has no closed form but is exponentially affine and is characterized through a fractional Riccati Volterra equation.

Rough Heston with Jumps model

The Rough Heston with Jumps (RHJ) model adds price jumps to the rough dynamics in the same way that Bates extends Heston, with jumps that arrive through an independent Poisson process N_t with intensity λ and have size G , whose logarithm $Y = \ln(G)$ now follows the asymmetric double exponential distribution of Kou (2002) with density

$$f_Y(y) = p\eta_1 e^{-\eta_1 y} 1_{\{y \geq 0\}} + q\eta_2 e^{\eta_2 y} 1_{\{y < 0\}}, \quad \eta_1 > 1, \eta_2 > 0, \quad (4)$$

where $p, q \geq 0, p + q = 1$ are the probabilities of an upward and a downward jump, and the tail parameters η_1 and η_2 control the magnitudes of the two jump directions. The moment generating function of the log price, used by the Laplace inversion method, combines the El Euch and Rosenbaum (2019) fractional Riccati with the Kou (2002) jump transform and its compensator, in the same product form as the proposed Lifted Heston with Jumps model, so the jump enters in closed form while the diffusive part still requires the fractional Riccati solution.

Lifted Heston model

The Lifted Heston (LH) model of Abi Jaber (2019) approximates Rough Heston with a finite-dimensional Markovian structure, replacing the fractional kernel by a weighted sum of n exponentials. The factor count spans the whole family, since with $n = 1$ the model reduces to classical Heston and as $n \rightarrow \infty$ it converges to Rough Heston. This study sets the number of lifting factors to twenty, following Abi Jaber (2019), and carries the same setting into the jump extension. The price dynamics keep the form of equation (1) and the variance process V_t^n follows

$$V_t^n = g_0^n(t) + \sum_{i=1}^n C_i^n U_t^{n,i}, \quad V_0^n = g_0^n(0) \quad (5)$$

where $g_0^n(t)$ is the initial forward variance curve and the sum $\sum_{i=1}^n C_i^n U_t^{n,i}$ aggregates the n factors $U_t^{n,i}$, each entering with weight C_i^n , so that the kernel approximation improves as n grows. Each factor follows

$$dU_t^{n,i} = (-x_i^n U_t^{n,i} - \kappa V_t^n)dt + \nu \sqrt{V_t^n} dW_t, \quad U_0^{n,i} = 0, \quad i = 1, \dots, n \quad (6)$$

where x_i^n is the mean-reversion speed of factor $U_t^{n,i}$, and κ and ν keep their roles as the mean-reversion rate and the volatility of volatility. The moment generating function of the log price, used by the Laplace inversion method, has no closed form and is given in Abi Jaber (2019) through a finite system of ordinary Riccati differential equations.

Lifted Heston with Jumps model

This study adopts the Lifted Heston (LH) framework (Abi Jaber, 2019) incorporating Kou's (2002) double-exponential jumps. The asset price S_t^n is driven by two correlated standard brownian motions B_t and W_t with correlation constant ρ , and by the independent Poisson process N_t with intensity λ . The number of lifting factors n is used to approximate the fractional kernel of the Rough Heston model. Under risk-neutral measure $\mathbb{Q}$, the underlying process S_t^n follows:

$$dS_t^n = S_t^n(r - \lambda E[G - 1])dt + S_t^n \sqrt{V_t^n} dB_t + (G - 1)S_t^n dN_t, \quad S_0^n > 0, \quad (7)$$

where S_t^n is the price of assets at time t , r is the drift term of the asset price process, which in risk-neutral measure the asset grows at risk-free interest rate. The term G is a non-negative random variables, such that $Y = \ln(G)$ and Y has asymmetric double exponential distribution with the density given in equation (4), and the same upward and downward probabilities and tail parameters as in the Rough Heston with Jumps model.

Moreover, $E[G - 1]$ is the expected percentage change per jump, including $\lambda E[G - 1]$ into drift term ensures the asset price process martingale. The variance process V_t follows the Lifted Heston dynamics in equations (5) and (6), with the same lifting weights and mean-reversion speeds.

The moment generating function of the log price, used by the Laplace inversion method, combines the Abi Jaber (2019) lifted Riccati system with the Kou (2002) jump transform.

$$M_T(u) = \exp \left(u \ln S_0 + (r - \lambda(E^Q[e^Y] - 1))uT + \int_0^T g_0^n(T-s)F(u, y(s)) ds + \lambda T(\varphi_Y(u) - 1) \right) \quad (8)$$

where the aggregated factor process and the forward variance curve are

$$y(t) = \sum_{i=1}^n c_i^n \psi_i(t) \quad (9)$$

$$g_0^n(t) = V_0 + \kappa \theta \sum_{i=1}^n \frac{c_i^n}{x_i^n} (1 - e^{-x_i^n t}) \quad (10)$$

the Riccati nonlinearity is

$$F(u, y) = \frac{1}{2}(u^2 - u) + (\rho v u - \kappa)y + \frac{1}{2}v^2 y^2 \quad (11)$$

and each factor solves the Riccati ordinary differential equation

$$\frac{d\psi_i}{dt} = -x_i^n \psi_i(t) + F(u, y(t)), \quad \psi_i(0) = 0, \quad i = 1, \dots, n \quad (12)$$

where the jump transform is

$$\varphi_Y(u) = E^Q[e^{uY}] = p \frac{\eta_1}{\eta_1 - u} + q \frac{\eta_2}{\eta_2 + u} \quad (13)$$

with compensator

$$E^Q[e^Y] = p \frac{\eta_1}{\eta_1 - 1} + q \frac{\eta_2}{\eta_2 + 1} \quad (14)$$

where n is the number of lifting factors, c_i^n and x_i^n are the weights and mean-reversion speeds of the lifting scheme of Abi Jaber (2019), ψ_i is the solution of the i -th factor Riccati equation, y is the aggregated factor process, g_0^n is the forward variance curve, and F is the Riccati nonlinearity. Because the compensated lifted log price and the compound Poisson jump sum are driven by independent sources of randomness under the risk-neutral measure, the transform factorizes into an exponentially affine diffusive component, whose coefficients solve the Riccati system above, and a jump component that evaluates in closed form through the Kou jump transform. The full derivation is provided in Appendix A.

Data

This study utilizes historical data on S&P 500 Index (SPX) index options and S&P 500 Index (SPX) and U.S. Treasury bill yields, sourced via the Bloomberg Terminal. The sample period spans from January 2017 to December 2023.

The option data set includes daily option prices, strike prices, time to maturity, implied volatility and the underlying index closing prices.

To improve data quality, the dataset is filtered to include only liquid options. Specifically, for each maturity, the selection consists of the five in-the-money options, one at-the-money option, and the five out-of-the-money options.

After filtering, the final dataset contains 495 153 option quotes available for the analysis. Of these, 64 968 quotes form the in sample calibration set and 430 185 quotes form the out of sample test set. The call side provides 244 948 quotes, made up of 31 268 in sample and 213 680 out of sample. The put side provides 250 205 quotes, made up of 33 700 in sample and 216 505 out of sample.

Classification of options by moneyness and maturity

This study classifies options into nine buckets by crossing moneyness and date to maturity.

Moneyness of call options

At the money

$$0.95 \leq \frac{K}{S} \leq 1.05$$

In the money

$$\frac{K}{S} < 0.95$$

Out of the money

$$\frac{K}{S} > 1.05$$

Moneyness of put options

At the money

$$0.95 \leq \frac{K}{S} \leq 1.05$$

In the money

	$\frac{K}{S} > 1.05$
Out of the money	
	$\frac{K}{S} < 0.95$
Date to maturity	
Short	$DTM < 20$
Medium	$20 \leq DTM \leq 180$
Long	$180 < DTM \leq 252$

where K is strike price, S is the price of asset at time t and DTM is date to maturity
The number of option quotes falling into each of the nine buckets is summarized in Table 1, separated into the in-sample and out-of-sample sets for call and put options.

Table 1 Number of option quotes by bucket

Bucket	Call options			Put options		
	In sample	Out of sample	Total	In sample	Out of sample	Total
Short-ITM	436	6039	6475	305	3853	4158
Short-ATM	13 497	75 197	88 694	13 392	72 591	85 983
Short-OTM	1601	14 404	16 005	3238	21 653	24 891
Med-ITM	424	6442	6866	339	5565	5904
Med-ATM	10 975	68 135	79 110	11 135	68 397	79 532
Med-OTM	2560	21 722	24 282	3489	22 544	26 033
Long-ITM	57	1642	1699	62	2030	2092
Long-ATM	1325	14 567	15 892	1365	15 410	16 775
Long-OTM	393	5532	5925	375	4462	4837
Total	31 268	213 680	244 948	33 700	216 505	250 205

Table 1 reports the number of call and put option quotes in each moneyness and maturity bucket, split into the in-sample and out-of-sample sets.

Table 1 also shows that liquidity is concentrated at the money. The ATM buckets hold about three quarters of all quotes on both surfaces, the most liquid being the short-ATM group, while the ITM buckets are very thin, the least liquid being the long maturity in the money group at under one percent of the sample. The put sample carries somewhat more OTM volume than the call sample, consistent with demand for downside protection.

In-sample and out-of-sample design

In-sample parameters are calibrated each Wednesday, and out-of-sample results are evaluated on the subsequent trading days (Thursday to Tuesday). The out-of-sample pricing uses the parameters estimated in-sample. The reason is that the goal of accuracy testing is not only to fit the term structure of IV on the calibration date, but also to evaluate whether the model can reproduce how the term structure evolves over the next few days. A model that fits the surface perfectly today but fails immediately the following days is not truly describing the underlying volatility process.

Market regime classification

The out of sample period is split into bull and bear regimes. Each trading day is classified by the dating rule of Lunde and Timmermann (2004) applied to the weekly S&P 500 index, under which a fall of at least 10% from a peak marks a bear market and a rise of at least 15% from a trough marks a bull market. The resulting classification over the sample period is shown in Figure 1.



Figure 1 S&P 500 index over the sample period with bear market periods shaded, under the dating rule of Lunde and Timmermann (2004).

The out of sample period is also split by volatility level. Following Tampouris and Dritsaki (2026), each trading day is classified as turbulent when the VIX closes above its 75th sample percentile, which is around 22.8 index points, and as calm otherwise.

Pricing by Laplace inversion

We price all models using the Laplace inversion method to ensure fair accuracy and speed comparisons. Each specification provides its risk-neutral transform (moment generating function) required by the inversion.

The Laplace inversion method is used to recover option prices in the time domain from the model's moment generating function (MGF) or Laplace–Fourier transform under the risk-neutral measure. From Hubálek et al. (2006), any affine stochastic volatility model, including Heston, Heston with Jump (Bates), Rough Heston (RH), Rough Heston with Jump (RHJ), Lifted Heston (LH), or Lifted Heston with Jumps (LHJ), the discounted option price C can be written in transform space as

$$C = E^{\mathbb{Q}}[e^{-rT}(S_T - K)^+] = \frac{e^{-rT}}{2\pi i} \int_{\gamma-i\infty}^{\gamma+i\infty} \frac{K^{1-u}}{u(u-1)} M_T(u) du \quad (15)$$

where K is strike price, S_t is the price of asset at time t , r is risk-free interest rate, $M_T(u) = E^{\mathbb{Q}}[S_T^u]$ is the moment generating function (MGF) of the terminal log price $\ln S_T$ and $\gamma > 1$ is chosen within the strip of convergence.

Stability of the transform inversion for the lifted models

The inversion in equation (15) is exact, but its numerical evaluation is delicate for the two lifted specifications, whose transform is produced by a time march of the factor Riccati equations in equation (12). The march is stable only below a critical frequency inversely proportional to the step size, and once the maturity passes roughly 0.6 years that frequency falls inside the integration range, beyond which the computed transform stops decaying and erupts into sharp spikes. The corrupted values are finite and pass a finiteness guard untouched, and the damage concentrates in the long maturity deep in the money calls, where a modest price error inverts into a large implied volatility error. The jump transform multiplies the integrand by a bounded factor and leaves the threshold unchanged, so the instability originates in the lifted kernel rather than in the jump component. This study therefore truncates the sampled transform from the first frequency at which the required decay fails, and with this safeguard the computed price converges at every maturity.

Calibration objective

Model parameters Θ are calibrated in implied volatility space rather than directly to option prices. The Black–Scholes model is used as a mapping tool between prices and implied volatilities, so that all models are evaluated on a common, normalized scale across strikes and maturities. The objective is to minimize RMSE as

$$\min_{\Theta} RMSE(\Theta) = \sqrt{\frac{1}{N} \sum_{i=1}^N (IV_i^{mkt} - IV_i^{\Theta})^2} \quad (16)$$

where Θ denotes model parameters, N is the number of options, IV_i^{mkt} is the implied volatility of the asset as inverse solved from the Black-Scholes model using the price of option i on that day, and IV_i^{Θ} is the implied volatility of the asset as inverse solved from the Black-Scholes model by using the price of option i obtained from the pricing model.

RESEARCH RESULTS

This section reports the empirical performance of the six specifications, all calibrated to S&P 500 option-implied volatilities and evaluated out of sample. Puts and calls are assessed separately throughout, and particular attention is given to how the proposed LHJ model performs against the five benchmarks.

Calibrated parameters

Tables B1 and B2 in Appendix B report the calibrated parameters for the put and call surfaces in full, each entry the mean across the weekly calibrations with the standard deviation in parentheses, including the Bates mean log jump and the asymmetric upward probability and tail rates of the Kou jump for RHJ and LHJ.

Read across the weekly calibrations, the estimates within each model fall into three groups. The parameters that shape the smile, H , ρ and ν , hold narrow ranges around model specific levels, with standard deviations in Tables B1 and B2 well below their means, so the surface geometry each model builds is reproducible from week to week. The parameters that share the same work, λ , θ and κ , are identified only in combination, because several mixtures of them price the same surface. Their weekly estimates therefore move within wide ranges, with standard deviations comparable to or larger than their means, which is the signature of this trade off rather than of an unstable fit, since the smile parameters stay narrow in the same calibrations and each Wednesday's estimates continue to price the following days out of sample. The initial variance V_0 shows the largest relative variation in every model, but its path follows the spot variance that the short maturity at the money quotes pin down each week, so its movement is the market state changing rather than the fit drifting.

The volatility of volatility ν and the speed of mean reversion κ rise steadily as the models grow more flexible. For the pure rough and lifted specifications the roughness index H stays low, near 0.25 to 0.36 and well below 0.5. The value 0.5 is the smooth case that recovers ordinary Heston, so an H this low shows that the calibrated variance is genuinely rough. Rough variance moves most sharply over very short horizons, which is exactly where the observed smile is steepest, and because those sharp movements fade quickly the fit needs a large ν to reach that steepness. This is why ν rises from about 0.57 in Bates to 1.67 in LHJ on the put side in Table B1. The lifted models reproduce the rough kernel with a set of factors that revert at very different speeds, and the fast factors push the reported κ from around three in Heston to as much as thirteen in LHJ on the call side in Table B2. A high ν and a high κ are the signature of this rough texture in the parameters, the same roughness that a low H states directly. Once a jump is added the roughness index rises in the jump models, to about 0.40 for RHJ and 0.38 for LHJ on puts against 0.28 for RH and 0.25 for

LH, because the jump carries part of the skew that roughness would otherwise fit on its own, so the optimizer backs off the roughness.

The correlation ρ is negative throughout and traces how the skew is funded. In Tables B1 and B2 it sits near -0.63 in Heston and Bates, where leverage carries the whole skew. It eases to between -0.55 and -0.62 once roughness supplies part of the short end skew. The jump then splits the two sides. On the put surface ρ deepens furthest in RHJ and LHJ, reaching about -0.74 and -0.76, which runs against the easing seen when only one ingredient is present (Bates at -0.61, RH at -0.56). The jump substitutes for the roughness rather than for the correlation, pushing H up from 0.28 in RH to 0.40 in RHJ, and the smoother diffusion then needs a deeper ρ to carry the medium and long maturity skew the short-dated jump does not reach. On the call surface, where there is little downside skew for the jump to carry, ρ in the jump models stays near the rough level around -0.62 and does not deepen. The jump is itself asymmetric and tilted to the downside. The Bates mean log jump μ is about -0.10 on both surfaces, and for the Kou jump the downward tail parameter η_2 near 18 on puts sits well below the upward η_1 near 35, so the average downward jump is roughly twice the size of the upward one. The intensity λ remains loosely pinned, with a standard deviation as large as its mean, because roughness already produces heavy tails and competes with the jump for the same observations. The initial variance V_0 is by contrast the most consistent parameter across the six models, near 0.03 to 0.045 for all of them, because the short-ATM level pins it down in each week. Its standard deviation across the weekly calibrations is large only because it tracks the changing spot variance, not because the fit struggles to identify it. The long run variance θ is instead large and weakly identified, trading off against κ and the factor weights.

In-sample fit

Table 2 In-sample IVRMSE (%) and total calibration time (minutes)

Model	Put RMSE (%)	Put time (min)	Call RMSE (%)	Call time (min)
Heston	3.97	1.8	3.83	1.5
Bates	3.30	2.8	3.69	2.7
RH	2.31	350	3.11	660
RHJ	2.11	985	2.69	588
LH	2.52	142	3.14	203
LHJ	2.02	267	2.68	219

Table 2 reports the pooled in-sample IVRMSE together with the total calibration time for each model.

Table 3 In-sample IVRMSE (%) by bucket for put options

Model	short-ITM	short-ATM	short-OTM	med-ITM	med-ATM	med-OTM	long-ITM	long-ATM	long-OTM
Heston	15.30	4.15	6.28	6.68	2.58	1.79	2.73	2.11	1.55
Bates	15.11	3.83	2.65	6.60	2.38	1.44	3.66	1.85	1.42
RH	12.61	2.61	2.10	5.06	1.25	0.76	1.81	1.09	0.78
RHJ	11.87	2.48	1.31	4.40	1.11	0.57	1.35	1.00	0.58
LH	13.15	2.57	3.46	4.77	1.35	0.93	2.33	2.04	1.17
LHJ	11.19	2.45	1.03	4.50	1.09	0.58	0.93	0.78	0.60

Table 3 reports the in-sample IVRMSE breakdown by moneyness and maturity bucket for put options.

Table 4 In-sample IVRMSE (%) by bucket for call options

Model	short-ITM	short-ATM	short-OTM	med-ITM	med-ATM	med-OTM	long-ITM	long-ATM	long-OTM
Heston	13.23	4.22	3.81	6.30	2.60	2.47	2.61	2.57	2.62
Bates	13.02	4.09	2.92	6.36	2.55	2.26	2.50	2.54	2.80
RH	12.45	3.54	3.28	5.39	1.55	1.48	1.08	1.39	1.31
RHJ	11.84	3.01	2.58	5.48	1.20	0.82	2.95	0.97	1.26
LH	12.29	3.55	3.43	5.30	1.58	1.62	1.66	1.26	1.21
LHJ	11.74	3.02	2.56	5.22	1.18	0.90	1.40	1.03	0.92

Table 4 reports the in-sample IVRMSE breakdown by moneyness and maturity bucket for call options.

Tables 2 to 4 report the in-sample fit, the pooled IVRMSE with the calibration time in Table 2 and the bucket breakdown for puts and calls in Tables 3 and 4. Two patterns stand out. First, every step up in model structure buys a tighter fit, on puts the pooled IVRMSE falls from 3.97 percent for Heston to 2.31 for RH and 2.02 for LHJ, and the same ordering holds on calls. Second, the two ingredients improve different parts of the surface. Roughness tightens the fit at the money, where the med-ATM put error drops from 2.58 percent for Heston to 1.25 for RH, while the jump acts in the short-OTM wing, where the put error falls to 1.03 percent for LHJ against 2.65 for Bates. These tighter fits come at a far higher calibration cost in Table 2.

Out-of-sample performance

Table 5 Overall out-of-sample IVRMSE (%)

Model	Put Option (%)	Call Option (%)
Heston	5.32	5.58
Bates	4.69	5.44
RH	4.69	5.31
RHJ	4.63	5.32
LH	4.95	5.36
LHJ	4.40	5.22

Table 5 reports the volume-weighted overall out-of-sample IVRMSE for each model.




Figure 2 Out-of-sample IVRMSE (%) by moneyness and maturity bucket for put options, one panel per model, darker cells marking smaller errors. Cell labels give the exact values, and the color scale is capped at 8 percent so that the noisy short-ITM cells do not compress the shading of the remaining buckets. Bold marks the most accurate model in each bucket.




Figure 3 Out-of-sample IVRMSE (%) by moneyness and maturity bucket for call options, one panel per model, darker cells marking smaller errors. Cell labels give the exact values, and the color scale is capped at 8 percent so that the noisy short-ITM cells do not compress the shading of the remaining buckets. Bold marks the most accurate model in each bucket.

Table 5 and Figures 2 and 3 present the out-of-sample IVRMSE, the overall error in Table 5 and the bucket breakdown in Figures 2 and 3. Every model loses accuracy once it prices out of sample, and the size of the loss depends on how many quotes a bucket holds. In the liquid med-ATM bucket the ranking survives, with the rough and lifted models still ahead at about 2.6 to 3.3 percent against 3.45 for Heston in Figure 2. In the thin med-ITM bucket the ranking inverts, the rough and lifted models deteriorate to 7.1 to 7.5 percent while Bates and Heston barely move, at 6.68 and 6.92. The difference is what the models were fitting. The thousands of at-the-money quotes pin down a persistent feature of the surface, so a tight fit there repeats on the following days, whereas the few hundred in-the-money quotes are noisy, and the flexible models spend part of their fit on noise that the following days do not reproduce. Their flexibility is therefore an asset in the data-rich center and a liability in the sparse wings.

Adding roughness improves accuracy because it corrects the part of the surface that the classical models fit worst. Implied volatility changes steeply across strikes near the money at short maturities, and a classical diffusion smooths that change while a rough one reproduces it. RH accordingly lowers the at-the-money error at every tenor, for example from 3.45 percent under Heston to 2.99 in the med-ATM put bucket of Figure 2, and RHJ improves on Bates in the same region, from 3.13 to 2.77, so the gain comes from roughness itself rather than from the jump. The overall gain in Table 5 is twice as large on puts, 0.63 points against 0.27 on calls, because roughness also repairs the steep short-OTM put wing, from 7.37 percent under Heston to 4.46 under RH, and the near-symmetric upside offers no comparable wing on calls.

Adding a jump component improves accuracy mainly in the short-OTM buckets, and it does so for every family. These options are priced almost entirely by the probability of a large move over a short horizon, which a continuous volatility path treats as nearly impossible, so a model without jumps underprices them while a jump supplies that tail directly. In Figure 2 the short-OTM put error falls from 7.37 percent for Heston to 4.62 for Bates, from 4.46 for RH to 3.86 for RHJ and from 4.83 for LH to 3.52 for LHJ. Because the improvement appears whether or not roughness is present, it can be credited to the jump itself. The benefit narrows as maturity lengthens, since over longer horizons a diffusion accumulates asymmetry of its own, and in the med-OTM put bucket the jump gain shrinks to 0.62 points for Bates over Heston and 0.09 for RHJ over RH.

The lifted models track the rough models they approximate. In Table 5 LH prices puts at 4.95 percent against 4.69 for RH and calls at 5.36 against 5.31, while in the jump pair LHJ reaches 4.40 percent against 4.63 for RHJ on puts and 5.22 against 5.32 on calls. The finite-factor approximation therefore recovers the behavior of the exact rough kernel closely enough that the choice between a rough model and its lifted form makes little difference to accuracy, and in the jump pair the lifted version is in fact the more accurate on both surfaces.

One column should be read with care. In the short-ITM bucket every model posts an error near 17 to 18 percent on both surfaces in Figures 2 and 3, because the bucket is thin and its quotes are noisy, so its pooled error is a measurement effect rather than a statement about model quality. That column therefore carries little ranking information, and the comparisons in this study rest on the remaining buckets.

These gains carry a heavy computational price, which Table 2 sets against the accuracy. Heston and Bates calibrate in under three minutes, the lifted models in roughly 140 to 270 minutes, and the rough models in several hundred, with RHJ on puts the most expensive at almost 1000 minutes. The lifted models therefore recover the rough accuracy at a fraction of the cost, which improves the trade off wherever rough dynamics are wanted.

Bull and bear market regimes

This section decomposes the out of sample performance into bull and bear market regimes, classified as described in the methodology. Every option quote inherits the regime of its trading day, and the IVRMSE is pooled within each regime.

Table 6 Overall out-of-sample IVRMSE (%) by market regime

Model	Put (Bull)	Put (Bear)	Call (Bull)	Call (Bear)
Heston	4.97	6.82	5.29	7.09
Bates	4.37	6.12	5.14	7.03
RH	4.76	6.56	5.69	7.02
RHJ	4.21	6.36	4.94	7.20
LH	4.54	6.70	5.07	6.89
LHJ	4.03	5.98	4.89	6.89

Table 6 reports the overall out-of-sample IVRMSE within the bull and bear market regimes.

Two patterns from the full sample reappear inside both regimes. First, every model is less accurate in the bear regime than in the bull regime, for example the best put model LHJ rises from 4.03 percent in the bull to 5.98 in the bear. A falling market steepens the volatility surface and makes it less stable from one day to the next, so a parameter set calibrated on a Wednesday describes the following days less well and the quotes themselves are noisier, which is consistent with the higher pricing errors reported for stressed markets by Augustyniak et al. (2023). Second, the jump mechanism identified above keeps working inside each regime, with the jump models repairing the short-maturity out-of-the-money buckets on both surfaces in the regime bucket breakdown reported in Borvornluck (2026).

The bucket breakdown within each market regime is reported in full in Borvornluck (2026). On the put surface the bull regime preserves the full-sample ordering, with the jump models in front, LHJ lowest at 4.03 percent ahead of RHJ at 4.21 and Bates at 4.37. The ordering largely survives the move into the bear regime, where LHJ remains the most accurate put model at 5.98 percent with Bates at 6.12 and RHJ at 6.36 close behind. The mechanism is the asymmetric jump. A bear market steepens the left wing of the surface, and the downward jump supplies exactly that steep tail, so the jump models fund the bear skew through the jump rather than straining the diffusion. RH and LH have no comparable protection and must carry the steeper bear skew through the diffusion alone, so they trail every model that carries a jump when the market falls, in Table 6 even the weakest jump model, RHJ at 6.36 percent, sits ahead of RH at 6.56 and LH at 6.70 on puts. The thin med-ITM bucket runs against this tide and stays with the classical models, with Bates the most accurate there in both regimes.

On the call surface the six models stay close in both regimes because the call wing holds almost no skew for roughness or jumps to capture. In the bull regime LHJ is lowest at 4.89 percent just ahead of RHJ at 4.94, and in the bear regime the six sit within about a third of a point of each other. Within that narrow field one movement stands out. RHJ slips from second in the bull regime to last in the bear, and the bucket breakdown traces the slip to the long maturities, where its long-

ATM error of 8.14 percent stands far above the 2.79 to 3.79 of the other five models. The exact rough jump model is therefore the least stable call pricer under stress, while its lifted counterpart LHJ stays at the front in both regimes.

In conclusion, the market regime scales the errors but does not overturn the put ranking. Every model loses accuracy in the bear regime because the surface it must track is steeper and less stable, yet the three jump models hold the first three put places in both regimes and LHJ leads both, at 4.03 percent in the bull and 5.98 in the bear, because only the models that carry a downward jump can supply the steeper bear-market left tail without straining their diffusion. The regime split therefore strengthens the case for LHJ, whose lead holds precisely when the market is hardest to price.

Calm and turbulent market regimes

This section repeats the decomposition under the volatility based classification described in the methodology, rather than a trend based one.

Table 7 Overall out-of-sample IVRMSE (%) by volatility regime

Model	Put (Calm)	Put (Turbulent)	Call (Calm)	Call (Turbulent)
Heston	4.19	7.62	4.49	7.87
Bates	3.50	7.03	4.44	7.59
RH	3.97	7.39	5.02	7.89
RHJ	3.05	7.40	4.25	7.54
LH	3.47	7.67	4.15	7.82
LHJ	3.08	6.88	4.22	7.33

Table 7 reports the overall out-of-sample IVRMSE within the calm and turbulent volatility regimes.

The separation between models lives in the calm regime, and there the ordering echoes the bull regime, with the advanced jump models in front. On puts RHJ and LHJ are the most accurate in calm conditions at 3.05 and 3.08 percent in Table 7. Turbulence rearranges this picture most for LH. Competitive at 3.47 percent in calm conditions, it posts the highest put error of the six at 7.67 on turbulent days, falling behind even its exact counterpart RH at 7.39. The finite-factor approximation is harmless on a calm surface but leaves the model short of the steep skew that extreme volatility demands, and LH has nothing else to supply it. LHJ shows the other half of the same argument, sharing the lifted kernel yet holding the lowest turbulent put error at 6.88 percent, because the asymmetric jump carries the stressed skew that the kernel alone cannot.

Taken together, the two lenses tell one story from two directions. They cut the sample differently, by trend and by volatility level, yet both place the advanced jump models ahead in the benign regime and both show LHJ keeping first place on puts under stress, while LH proves the most fragile and finishes last on turbulent days. The mechanism is the same asymmetric jump that carried the bear regime. The practical point is that when volatility is extreme the error is set by market conditions rather than by model structure, so the case for the advanced models rests mostly on the calm periods that make up the bulk of the sample, with the added assurance that the jump protects the downside fit when conditions deteriorate.

Performance of the Lifted Heston with Jumps model

The proposed lifted Heston with jumps (LHJ) model owes its put-surface accuracy to the match between its two ingredients and the two features that shape index puts. Roughness supplies the steep at-the-money skew that a classical diffusion smooths over, and the asymmetric jump supplies the short-maturity crash tail that no continuous path can reach. LHJ carries both, so it improves on Heston at the money at every maturity in Figure 2 and posts the lowest short-OTM put error of the six at 3.52 percent, against 4.46 for RH and 4.83 for LH, the same dynamics without a jump. Because the two ingredients repair different regions of the surface, their gains add rather than

overlap, and LHJ records the lowest put error in five of the nine buckets and the lowest overall error at 4.40 percent in Table 5. The lead survives every market condition, since stress concentrates in the left wing and demands exactly the tail the jump supplies, so LHJ stays the most accurate put model in the bull, the bear and the turbulent regimes and sits slightly behind RHJ in the calm one. The put lead is therefore not a fair-weather artifact but a structural property of the model.

On the call surface the six models converge, for the reason given above, that the index right tail is close to lognormal and offers little skew for roughness to trace and no crash for a jump to price. Within that narrow field LHJ still prices the call surface at 5.22 percent, narrowly the lowest of the six, and it prices the long-maturity deep-ITM bucket on par with the exact rough jump model. The model gains where the surface carries structure to capture and gives up nothing where it does not.

Set against these accuracy results is the calibration cost in Table 2, and here the advantage of LHJ is structural, the lifted kernel replaces the fractional Riccati equation of the exact rough models with a standard system of ordinary differential equations. LHJ therefore reaches its put accuracy in 267 minutes, roughly a quarter of the 985 the exact RHJ model requires, and gives up no accuracy in exchange.

The buckets in which the proposed model earns its lead are also the ones that carry the greatest economic weight. Short-OTM puts are the instruments through which the market insures against a crash, and their quotes are the inputs to hedging and to tail risk assessment. This study measures pricing accuracy only, so the lead of roughly a fifth relative to RH and a quarter relative to LH in Figure 2 shows that the proposed model prices these options more accurately, and nothing more. Whether good pricing performance leads to good hedging performance is left for future work.

DISCUSSION & CONCLUSION

Three findings survive every cut of the sample. First, the advanced models earn their accuracy on the put surface, where LHJ posts the lowest out-of-sample error at 4.40 percent with RHJ close behind, while on calls the six models finish within about a third of a point of one another because the near-lognormal right tail offers little for roughness or jumps to capture. Second, the two ingredients act in different places and independently of each other, roughness corrects the at-the-money skew that a classical diffusion smooths over, the jump supplies the short-maturity crash tail that no continuous path can reach, and each gain appears whether or not the other is present. Third, the advantage is anchored in the liquid part of the surface, the ranking survives where quotes are dense and inverts toward Bates and Heston in the thin in-the-money buckets. Across regimes LHJ leads the put surface in the bull, the bear and the turbulent markets and sits marginally behind RHJ in the calm one, whereas the plain lifted model posts the highest turbulent put error and the exact rough jump model slips to last on bear-market calls.

These results also locate the limits of the lifting itself. The approximation converges quickly in the number of factors, and this study prices with twenty factors, a setting at which the lifted models already track their exact rough counterparts to within a fraction of a percentage point in Table 5, so the residual kernel error is smaller than the differences between the models. Raising the factor count further buys little additional roughness while adding weights that trade off against the long run variance and the volatility of volatility, so the calibration grows slower and less well identified rather than more accurate, and the in-the-money buckets already show the flexible models fitting noise that the following days do not reproduce. The turbulent regime sharpens rather than weakens this reading. LH fails on turbulent days not for lack of factors but because no kernel supplies the stressed crash tail, so the remedy is the jump rather than a longer lifting. The factor count is therefore best set by the liquid surface with the tail left to the jump, which is the division of labor the proposed model implements.

LHJ edges the exact rough jump model on both surfaces and reaches that accuracy in roughly a quarter of the calibration time, 267 minutes against 985 on puts. For practical use the results reduce

to one default and two exceptions, LHJ across the liquid surface in every market condition, with the far cheaper Bates model kept for the noisy in-the-money buckets and for long-dated puts in falling markets. Rough volatility with jumps is therefore no longer a trade between accuracy and tractability, the lifting delivers both, and that is the contribution of the proposed model. Moreover, the hedging comparison raised in the findings can be studied in future work.

ACKNOWLEDGEMENTS

This paper is based on the author's master's thesis at the Department of Banking and Finance, Faculty of Commerce and Accountancy, Chulalongkorn University. The author is grateful to the thesis advisor and the examination committee for their guidance and comments.

REFERENCES

- Abi Jaber, E. (2019). Lifting the Heston model. *Quantitative Finance*, 19(12), 1995–2013.
- Albrecher, H., Mayer, P., Schoutens, W., & Tistaert, J. (2007). The little Heston trap. *Wilmott Magazine*, 2007(January), 83–92.
- Augustyniak, M., Badescu, A., & Bégin, J.-F. (2023). A discrete-time hedging framework with multiple factors and fat tails: On what matters. *Journal of Econometrics*, 232(2), 416–444.
- Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *The Review of Financial Studies*, 9(1), 69–107.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- Borvornluck, N. (2026). Option pricing under the Lifted Heston model with jumps [Unpublished master's thesis]. Chulalongkorn University.
- Cajueiro, D. O., & Tabak, B. M. (2007). Long-range dependence and multifractality in the term structure of LIBOR interest rates. *Physica A: Statistical Mechanics and its Applications*, 373, 603–614.
- Ding, Z., Granger, C. W. J., & Engle, R. F. (1993). A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1), 83–106.
- El Euch, O., & Rosenbaum, M. (2019). The characteristic function of rough Heston models. *Mathematical Finance*, 29(1), 3–38.
- Fama, E. F. (1965). The behavior of stock-market prices. *The Journal of Business*, 38(1), 34–105.
- Gatheral, J., Jaisson, T., & Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, 18(6), 933–949.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2), 327–343.
- Hubálek, F., Kallsen, J., & Krawczyk, L. (2006). Variance-optimal hedging for processes with stationary independent increments. *The Annals of Applied Probability*, 16(2), 853–885.
- Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, 48(8), 1086–1101.
- Lunde, A., & Timmermann, A. (2004). Duration dependence in stock prices: An analysis of bull and bear markets. *Journal of Business & Economic Statistics*, 22(3), 253–273.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business*, 36(4), 394–419.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1–2), 125–144.
- Tampouris, A., & Dritsaki, C. (2026). Adaptive hierarchical hidden Markov models for structural market change. *Journal of Risk and Financial Management*, 19(1), 15.

APPENDIX A. DERIVATION OF THE LIFTED HESTON WITH JUMPS MOMENT GENERATING FUNCTION

This appendix derives the moment generating function of the log price under the Lifted Heston with Jumps model presented in the Research Methodology section. The argument has five steps. We first state the dynamics of the log price. We then split the log price into a continuous component and a jump component and exploit their independence to factor the moment generating function. We obtain the continuous factor from an exponential affine ansatz combined with Ito's formula, which yields a finite system of Riccati ordinary differential equations. We obtain the jump factor in closed form from the compound Poisson structure of the jumps. We finally combine the two factors.

A.1 Model and notation

Under the risk neutral measure Q the asset price follows the dynamics

$$dS_t = S_t \left[(r - \lambda m) dt + \sqrt{V_t^n} dB_t \right] + (G - 1) S_t dN_t$$

where the driving Brownian motion $B_t = \rho W_t + \sqrt{1 - \rho^2} W_t^\perp$ is correlated with W_t , the counting process N_t is a Poisson process with intensity λ , and the jump multipliers G are independent and identically distributed with $Y = \ln G$ following the Kou double exponential law f_Y . The mean jump compensator

$$m = E^Q[G - 1] = E^Q[e^Y - 1] = \varphi_Y(1) - 1$$

keeps the discounted asset price a martingale. The variance is a weighted sum of n mean reverting factors

$$V_t^n = g_0^n(t) + \sum_{i=1}^n C_i^n U_t^{n,i}$$

$$dU_t^{n,i} = (-x_i^n U_t^{n,i} - \kappa V_t^n) dt + v \sqrt{V_t^n} dW_t, U_0^{n,i} = 0, i = 1, \dots, n$$

with weights C_i^n and mean reversion speeds x_i^n , common mean reversion rate κ , volatility of variance v , and deterministic forward variance curve g_0^n . The Poisson process N and the jump marks $\{Y_k\}$ are independent of the Brownian pair $(W, W^\perp)$.

A.2 Dynamics of the log price

Applying Ito's formula for jump diffusions to $\ln S_t$ gives

$$\ln S_T = \ln S_0 + \int_0^T \left(r - \lambda m - \frac{1}{2} V_s^n \right) ds + \int_0^T \sqrt{V_s^n} dB_s + \sum_{k=1}^{N(T)} Y_k$$

The term $-\frac{1}{2} V_s^n$ comes from the quadratic variation of the continuous martingale part, and a jump of the asset price by the factor G contributes $\ln G = Y$ to the log price.

A.3 Separation of the continuous and jump parts

We define the jump stripped log price

$$\Lambda_T = \ln S_T - \sum_{k=1}^{N(T)} Y_k = \ln S_0 + \int_0^T \left(r - \lambda m - \frac{1}{2} V_s^n \right) ds + \int_0^T \sqrt{V_s^n} dB_s$$

The pair (Λ_t, V_t^n) is a lifted Heston log price that carries the compensated drift $r - \lambda m$. Because the Poisson process and its marks are independent of the Brownian motions that drive Λ and V^n , the variables Λ_T and $\sum_{k=1}^{N(T)} Y_k$ are independent. The moment generating function factors accordingly

$$M_T(u) = E^Q[e^{u \ln S_T}] = E^Q[e^{u \Lambda_T}] E^Q[e^{u \sum_{k=1}^{N(T)} Y_k}]$$

A.4 The continuous factor through the affine transform

We evaluate the continuous factor $E^Q[e^{u \Lambda_T}]$. Since $(\Lambda, U^{n,1}, \dots, U^{n,n})$ is an affine process, we posit that the conditional moment generating function has the exponential affine form

$$M(t) = E^Q[e^{u\Lambda_t} | \mathcal{F}_t] = \exp\left(\phi(t) + u\Lambda_t + \sum_{i=1}^n a_i(t)U_t^{n,i}\right)$$

where ϕ and a_i are deterministic with $a_i(t) = C_i^n \psi_i(T-t)$. The process $M(t)$ is a martingale, so the drift of its differential vanishes. We write the aggregated factor $y(\tau) = \sum_{i=1}^n C_i^n \psi_i(\tau)$ and recall $V_t^n = g_0^n(t) + \sum_i C_i^n U_t^{n,i}$. The quadratic covariations of the drivers are

$$d\langle \Lambda \rangle_t = V_t^n dt, d\langle U^{n,i}, U^{n,j} \rangle_t = v^2 V_t^n dt, d\langle \Lambda, U^{n,i} \rangle_t = \rho v V_t^n dt$$

Applying Ito's formula to $M(t) = e^{\Phi(t)}$ with $\Phi(t) = \phi(t) + u\Lambda_t + \sum_i a_i(t)U_t^{n,i}$ and setting the drift to zero gives

$$\phi'(t) + u(r - \lambda m) + \sum_{i=1}^n (a_i'(t) - x_i^n a_i(t)) U_t^{n,i} + V_t^n F(u, y(T-t)) = 0$$

where the quadratic nonlinearity collected from the first and second order terms is

$$F(u, y) = \frac{1}{2}(u^2 - u) + (\rho v u - \kappa)y + \frac{1}{2}v^2 y^2$$

We substitute $V_t^n = g_0^n(t) + \sum_i C_i^n U_t^{n,i}$ and match the coefficient of each factor $U_t^{n,i}$, which is a free state variable. This produces the n dimensional Riccati system

$$a_i'(t) - x_i^n a_i(t) + C_i^n F(u, y(T-t)) = 0$$

With $a_i(t) = C_i^n \psi_i(T-t)$ and the change of variable $\tau = T-t$ this becomes

$$\psi_i'(\tau) = -x_i^n \psi_i(\tau) + F(u, y(\tau)), \psi_i(0) = 0, i = 1, \dots, n$$

The initial condition $\psi_i(0) = 0$ follows from the terminal requirement $M(T) = e^{u\Lambda_T}$, which forces the factor coefficients to vanish at $\tau = 0$. Matching the remaining state independent terms gives

$$\phi'(t) = -u(r - \lambda m) - g_0^n(t)F(u, y(T-t))$$

We integrate from $t = 0$ to $t = T$ with the terminal value $\phi(T) = 0$ and change the variable of integration to $s = T-t$, which yields

$$\phi(0) = u(r - \lambda m)T + \int_0^T g_0^n(T-s)F(u, y_s)ds$$

where $y_s = \sum_{i=1}^n C_i^n \psi_i(s)$. Evaluating the ansatz at $t = 0$ with $U_0^{n,i} = 0$ and $\Lambda_0 = \ln S_0$ gives the continuous factor

$$E^Q[e^{u\Lambda_T}] = \exp\left\{u \ln S_0 + (r - \lambda m)uT + \int_0^T g_0^n(T-s)F(u, y_s)ds\right\}$$

The forward variance curve follows from the factor dynamics. Taking expectations in the factor equation and solving the resulting linear ordinary differential equation gives

$$g_0^n(t) = V_0 + \kappa\theta \sum_{i=1}^n \frac{C_i^n}{x_i^n} (1 - e^{-x_i^n t})$$

which is the input curve used throughout the study, with long run variance θ and initial variance V_0 .

A.5 The jump factor through the compound Poisson structure

We evaluate the jump factor. Conditioning on the number of jumps $N(T)$, which is Poisson with mean λT , and using that the marks Y_k are independent and identically distributed gives

$$\begin{aligned} E^Q[e^{u \sum_{k=1}^{N(T)} Y_k}] &= \sum_{n \geq 0} \frac{e^{-\lambda T} (\lambda T)^n}{n!} (E^Q[e^{uY}])^n = e^{-\lambda T} \exp(\lambda T E^Q[e^{uY}]) \\ &= \exp\{\lambda T (\varphi_Y(u) - 1)\} \end{aligned}$$

The jump transform follows by integrating the Kou density

$$\begin{aligned}\varphi_Y(u) &= E^Q[e^{uY}] = \int_{-\infty}^{\infty} e^{uy} f_Y(y) dy = \int_0^{\infty} p \eta_1 e^{-(\eta_1-u)y} dy + \int_{-\infty}^0 q \eta_2 e^{(\eta_2+u)y} dy \\ &= \frac{p\eta_1}{\eta_1 - u} + \frac{q\eta_2}{\eta_2 + u}\end{aligned}$$

which is finite for $\text{Re}(u) \in (-\eta_2, \eta_1)$. The compensator constant of Section A.1 is the value at $u = 1$, namely

$$m = \varphi_Y(1) - 1 = \frac{p\eta_1}{\eta_1 - 1} + \frac{q\eta_2}{\eta_2 + 1} - 1$$

A.6 The moment generating function

Combining the continuous factor of Section A.4 and the jump factor of Section A.5 through the factorization of Section A.3 gives the moment generating function

$$M_T(u) = \exp \left\{ u \ln S_0 + (r - \lambda E^Q[e^Y - 1])uT + \int_0^T g_0^n(T-s)F(u, y_s)ds + \lambda T(\varphi_Y(u) - 1) \right\}$$

This is the moment generating function stated in equation (8), which completes the derivation.

APPENDIX B. CALIBRATED PARAMETERS

Table B1 Calibrated parameters for put options

Parameter	Heston	Bates	RH	RHJ	LH	LHJ
V_0	0.039 (0.039)	0.034 (0.037)	0.036 (0.061)	0.033 (0.058)	0.035 (0.059)	0.031 (0.056)
θ	0.062 (0.022)	0.055 (0.018)	0.148 (0.147)	0.148 (0.165)	0.143 (0.151)	0.190 (0.193)
κ	3.20 (2.01)	3.31 (3.00)	5.71 (5.99)	5.48 (6.15)	6.59 (6.33)	6.46 (6.84)
ν	0.611 (0.204)	0.565 (0.198)	1.059 (0.580)	0.929 (0.514)	1.353 (0.644)	1.665 (0.554)
ρ	-0.631 (0.139)	-0.605 (0.205)	-0.560 (0.339)	-0.739 (0.321)	-0.546 (0.323)	-0.757 (0.301)
λ	—	0.194 (0.303)	—	1.046 (1.612)	—	1.039 (1.620)
μ	—	-0.101 (0.046)	—	—	—	—
δ	—	0.126 (0.062)	—	—	—	—
p	—	—	—	0.421 (0.374)	—	0.445 (0.387)
η_1	—	—	—	35.33 (18.56)	—	36.81 (17.54)
η_2	—	—	—	17.59 (19.10)	—	18.12 (19.23)
H	—	—	0.284 (0.167)	0.403 (0.133)	0.252 (0.130)	0.382 (0.100)

Table B1 reports the calibrated parameters for the put surface, each entry the mean across the weekly calibrations with the standard deviation in parentheses.

Table B2 Calibrated parameters for call options

Parameter	Heston	Bates	RH	RHJ	LH	LHJ
V_0	0.034 (0.045)	0.030 (0.037)	0.039 (0.050)	0.041 (0.058)	0.036 (0.052)	0.039 (0.058)
θ	0.050 (0.018)	0.046 (0.017)	0.063 (0.077)	0.039 (0.071)	0.069 (0.089)	0.063 (0.101)
κ	2.91 (2.72)	2.86 (2.57)	6.57 (5.88)	10.40 (7.65)	7.67 (6.45)	13.16 (7.65)
ν	0.511 (0.225)	0.467 (0.161)	0.719 (0.399)	0.849 (0.554)	1.200 (0.646)	1.301 (0.632)
ρ	-0.643 (0.136)	-0.644 (0.152)	-0.615 (0.271)	-0.619 (0.409)	-0.600 (0.254)	-0.625 (0.381)
λ	—	0.229 (0.528)	—	1.394 (1.541)	—	1.313 (1.483)
μ	—	-0.096 (0.055)	—	—	—	—
δ	—	0.110 (0.042)	—	—	—	—
p	—	—	—	0.520 (0.375)	—	0.533 (0.372)
η_1	—	—	—	39.86 (14.65)	—	40.60 (14.11)
η_2	—	—	—	28.18 (19.05)	—	28.69 (19.77)
H	—	—	0.355 (0.139)	0.308 (0.177)	0.359 (0.112)	0.290 (0.152)

Table B2 reports the calibrated parameters for the call surface, each entry the mean across the weekly calibrations with the standard deviation in parentheses.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Conflicts of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.



Copyright: © 2026 by the authors. This is a fully open-access article distributed under the terms of the Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0).