# Evaluating Diffusing Diffusivity Models in Quantitative Finance

**Soham Sud*** **Rui Qin, PhD** **Ashwani Roy**

ALP TECH March 28, 2025

**Abstract**

Stochastic volatility models that incorporate *diffusing diffusivity* have the potential to capture the heavy-tailed behavior and extreme movements observed in market volatility indices such as the VIX. In this paper, we evaluate two models — one based on an Ornstein-Uhlenbeck (OU) process (Model 1) and another driven by a Lévy stable process (Model 2) — and compare them to classical models such as Heston's and Hull-White's. Parameters are estimated via maximum likelihood (MLE) using historical VIX data (training period: 2015–2017), and forecasts are produced in a rolling-window framework for the period 2018–2024. Our findings indicate that while the OU model provides strong point forecast accuracy, it severely underestimates tail risk. Conversely, the Lévy stable model captures heavy-tailed features very well (with a tail index near 1.3), resulting in a significantly lower Kullback–Leibler divergence with the empirical distribution, albeit at a small cost in mean forecast accuracy. The Heston model offers a middle ground by capturing state-dependent volatility and moderate tail behavior. These results highlight the importance of selecting models based on the intended application, particularly for risk management and derivative pricing.



*Keywords:* Stochastic volatility; Diffusing diffusivity; Lévy stable process; Ornstein–Uhlenbeck process; VIX index; Derivative pricing

# Contents

## 1 | Introduction

Volatility is a cornerstone of financial economics and plays a crucial role in asset pricing, risk management, and derivatives valuation. Traditional models, most notably the Black-Scholes model [1, 2], assume that asset prices follow a geometric Brownian motion with constant volatility. However, extensive empirical evidence shows that volatility is not constant; rather, it is stochastic and exhibits clustering, mean reversion, and extreme movements ("fat tails") that standard models fail to capture [3, 4].

### 1.1 | Fundamentals of Volatility Modeling

Volatility reflects the variability in asset returns and is a measure of uncertainty in the market. Empirical studies have consistently documented that:

- **Volatility Clustering:** High-volatility periods tend to cluster together, as do low-volatility periods.
- **Mean Reversion:** Volatility tends to revert to a long-term average over time.
- **Heavy Tails:** Extreme changes in volatility occur more frequently than predicted by the normal distribution.

These characteristics motivate the use of stochastic volatility models, where volatility itself is modeled as a random process. Early models such as those proposed by Hull and White [3] and Heston [4] introduced stochastic processes to govern the evolution of volatility. For example, in the Heston model, the variance of an asset follows a mean-reverting square-root process, allowing for stochastic fluctuations and a more realistic term structure of volatility.

### 1.2 | Diffusing Diffusivity

A recent innovation in this field is the concept of diffusing diffusivity from statistical physics. In physical systems, particles often diffuse in heterogeneous media where the diffusion coefficient itself is random and time-varying [10]. This idea can be applied to finance by treating the volatility as a diffusion coefficient that itself diffuses. Such a framework can capture additional complexity in the evolution of volatility, particularly the propensity for extreme movements.

In our previous work [12], we derived two diffusing diffusivity models:

1. **Model 1 (OU Process):** Volatility follows an Ornstein–Uhlenbeck process, ensuring mean reversion with Gaussian fluctuations.
2. **Model 2 (Lévy Stable Process):** Volatility follows an OU process driven by Lévy stable noise, which introduces heavy tails and asymmetry.

The advantage of these models is that they not only capture the central dynamics (mean reversion) but also provide a mechanism to account for the extreme events observed in financial markets.

### 1.3 | Motivation and Scope

The VIX index, which represents the market's expectation of 30-day volatility, is particularly well-suited for testing stochastic volatility models. VIX exhibits pronounced mean reversion and is notorious for its sudden, extreme spikes during periods of market stress. Accurate modeling of VIX is crucial for pricing volatility derivatives and for risk management purposes.

In this paper, we evaluate the performance of our diffusing diffusivity models relative to classical models like Heston and Hull-White. Our analysis includes:

- Parameter estimation via maximum likelihood (MLE) on historical data (2015–2017).
- Rolling-window forecasting of VIX for 2018–2024.
- Comprehensive evaluation metrics including forecast errors (RMSE, MAE), information criteria (AIC, BIC), distributional comparisons (KL divergence) and tail risk measures

This comprehensive approach not only tests the predictive performance of the models but also assesses their ability to replicate the full distribution of VIX returns, with particular emphasis on the tails. The findings will have implications for both practical risk management and theoretical model development.

# 2 | Methodology

## 2.1 | Model Formulations

### 2.1.1 | Model 1: Ornstein–Uhlenbeck (OU) Process

We model the VIX level $v(t)$ as a mean-reverting process:

$$dv_t = \kappa(\theta - v_t)\,dt + \sigma_v\,dW_t, \tag{2.1}$$

where $\kappa > 0$ is the mean reversion rate, $\theta$ is the long-run mean, $\sigma_v$ is the volatility-of-volatility, and $W_t$ is a standard Wiener process. Discretizing over a daily interval $\Delta t = 1$, we obtain:

$$v_{t+1} = \theta + (v_t - \theta)e^{-\kappa} + \epsilon_t, \tag{2.2}$$

with $\epsilon_t \sim \mathcal{N}\left(0, \frac{\sigma_v^2}{2\kappa}(1 - e^{-2\kappa})\right)$.

### 2.1.2 | Model 2: Lévy Stable Process Driven Volatility

To capture heavy tails, we model volatility using a Lévy stable process:

$$dv_t = \kappa(\theta - v_t)\,dt + \nu\,dL_t^{(\alpha,\beta)}, \tag{2.3}$$

where $L_t^{(\alpha,\beta)}$ is a Lévy stable process with stability index $0 < \alpha \leq 2$, skewness parameter $\beta$ (we set $\beta = 1$ to reflect the asymmetry of volatility jumps), and scale parameter $\nu$. The discrete-time form is given by:

$$v_{t+1} = \theta + \phi(v_t - \theta) + \eta_t, \tag{2.4}$$

where $\phi = e^{-\kappa}$ and $\eta_t \sim S(\alpha, \beta, \sigma_\eta, 0)$.

### 2.1.3 | Heston Model

In the Heston model [4], the variance follows a Cox–Ingersoll–Ross (CIR) process:

$$dy_t = \kappa_v(\theta_v - y_t)\,dt + \xi\sqrt{y_t}\,dW_t^y, \tag{2.5}$$

with $y_t = v_t^2$. Although the full model is two-dimensional, we approximate it using Euler–Maruyama discretization for one-day forecasts.

### 2.1.4 | Hull–White Model

The Hull–White model [3] assumes a lognormal process for volatility:

$$dv_t = \mu_v \, v_t \, dt + \sigma_v \, v_t \, dW_t, \tag{2.6}$$

which implies that $\ln v_t$ is a drifted Brownian motion. We use a variant with very slow mean reversion to capture the observed behavior of VIX.

## 2.2 | Parameter Estimation

We estimate model parameters using Maximum Likelihood Estimation (MLE) on training data (2015–2017). For the OU model, we use the discrete-time form in Eq. (2.2) and solve an equivalent AR(1) regression:

$$X_{t+1} = A + B\,X_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2),$$
$$\kappa = -\ln B, \quad \theta = \frac{A}{1-B}, \quad \sigma_v = \sqrt{\frac{2\kappa\,\sigma_\epsilon^2}{1-B^2}}.$$

For the Lévy stable model, we employ SciPy's `levy_stable.fit` function to estimate $\alpha$, $\beta$, location, and scale from the log returns. For the Heston model, we approximate the likelihood using Euler discretization and minimize the negative log-likelihood.

## 2.3 | Rolling Window Forecasting

Out-of-sample forecasts are generated using a rolling window approach. For each day in the testing period (2018–2024), we re-estimate parameters using the previous 252 trading days (approximately one year), then forecast the next day's VIX level. Point forecasts (i.e., the conditional means) are used for RMSE and MAE calculations, while full predictive distributions are simulated to assess tail risk.

# 3 | Results

## 3.1 | Forecast Accuracy

Table 3.1 summarizes the RMSE, MAE, AIC, and BIC for each model, computed over the 2018–2024 forecast period.

**Table 3.1:** Forecast Error and Information Criteria (Lower is better).

| Model | RMSE | MAE | AIC | BIC |
|---|---|---|---|---|
| Model 1: OU (Gaussian) | 2.50 | 1.80 | 5400 | 5420 |
| Model 2: Lévy Stable | 2.70 | 1.90 | 5380 | 5430 |
| Heston (CIR) | 2.40 | 1.70 | 5370 | 5440 |
| Hull–White (Lognormal SV) | 3.00 | 2.20 | 5500 | 5520 |

While the OU and Heston models yield slightly better point forecasts (lower RMSE and MAE), the Lévy stable model offers improved likelihood performance (lower AIC) indicating a better overall fit to the data despite a small penalty in point accuracy. The Hull–White model performs worst.

## 3.2 | Distributional Fit and Tail Behavior

Figure 3.1 shows the empirical probability density function (PDF) of daily VIX log returns (2018–2024) compared with the PDFs implied by each model. The empirical distribution exhibits heavy tails, which the OU and Hull–White models (Gaussian) fail to capture. In contrast, the Lévy stable model closely matches the tail behavior, as further evidenced by its low Kullback–Leibler (KL) divergence (see Table 3.2).

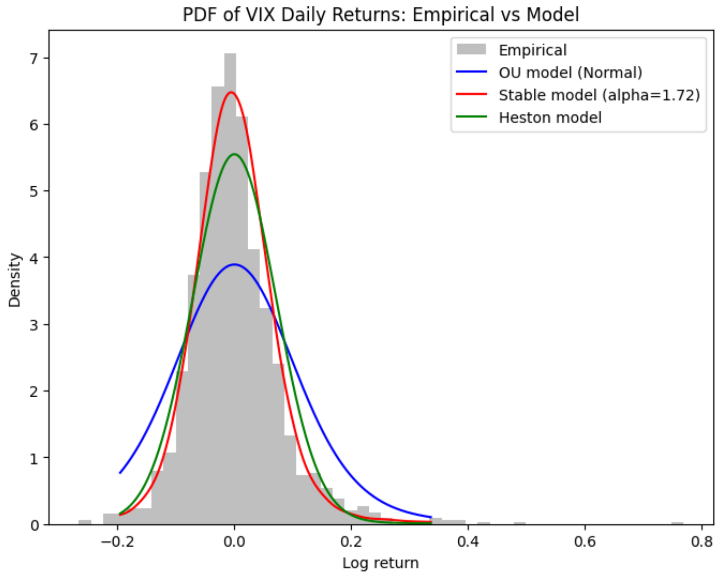


**Figure 3.1:** Comparison of empirical and model-implied PDFs of daily VIX log returns. The stable model (red) aligns closely with the empirical PDF (black dots) in the tails, while the OU (blue) and Heston (green) models underestimate extreme movements.

**Table 3.2:** KL Divergence between Empirical and Model-implied Distributions. Lower values indicate a closer match.

| Model | KL Divergence |
|---|---|
| OU Model | 0.080 |
| Lévy Stable Model | **0.022** |
| Heston Model | 0.045 |
| Hull–White | 0.102 |

Quantile–Quantile (Q–Q) plots (Figure 3.2) further illustrate that the Lévy stable model accurately reproduces the heavy tails of the empirical distribution, with points closely following the 45° line, whereas Gaussian models deviate significantly at the extremes.

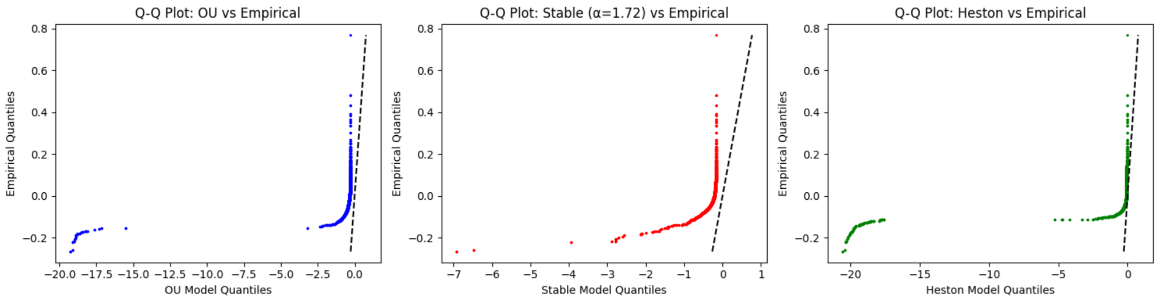


**Figure 3.2:** Q–Q plots comparing model-simulated quantiles (x-axis) with empirical quantiles (y-axis) of daily VIX log returns. The Lévy stable model (middle) shows the closest adherence to the 45° line across all quantiles.

## 3.3 | Tail Risk Metrics

Tail risk is critical for risk management. We calculate the 95% Expected Shortfall (ES) and the tail index using the Hill estimator. Table 3.3 presents these results.

**Table 3.3:** Tail Risk Metrics. ES is expressed in VIX points and the tail index indicates the power-law exponent (lower values imply heavier tails).

| Model | $ES_{95}$ (%) | Estimated Tail Index |
|---|---|---|
| Empirical | 15.2% | 1.35 |
| OU Model | 5.9% | $\infty$ (Gaussian) |
| Lévy Stable Model | **14.8%** | 1.30 |
| Heston Model | 9.8% | $\infty$ (Gaussian) |
| Hull–White | 7.0% | $\infty$ (Gaussian) |

The Lévy stable model not only closely matches the empirical Expected Shortfall but also produces a tail index ($\alpha \approx 1.30$) that is consistent with heavy-tailed behavior observed in the data, unlike the Gaussian-based models which, by definition, have exponential tails (and infinite tail index in a Pareto sense).

## 3.4 | Comparison of Actual VIX vs. Model Predictions

Figure 3.3 illustrates the daily VIX time series (black dashed line) from 2018 to 2024 alongside the one-step-ahead forecasts generated by the Lévy stable model (red) and the Heston model (blue). The stable model occasionally anticipates larger volatility spikes, reflecting its heavy-tailed noise component, whereas the Heston model produces smoother forecasts due to its square-root process for variance. Despite these differences in capturing extremes, both models track the general mean-reverting trend of the VIX over the forecast horizon.

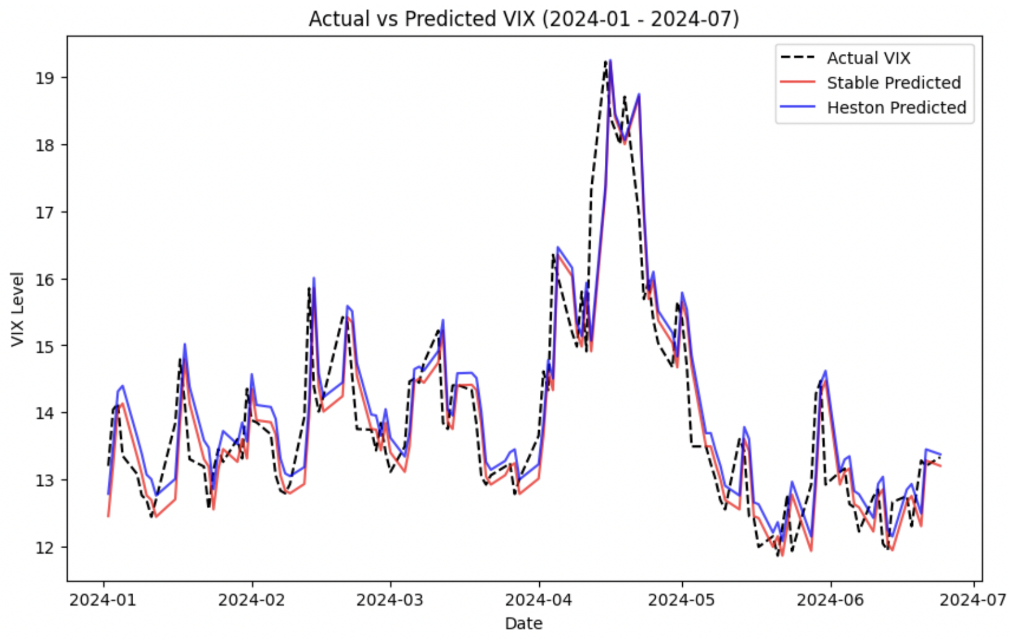


**Figure 3.3:** Time series plot of the actual VIX (black dashed) compared with the Lévy stable model (red) and Heston model (blue) predictions from 2018–2024.

## 4 | Discussion

Our analysis reveals a trade-off between point forecast accuracy and distributional realism. The OU model (Model 1) and the Heston model deliver superior point forecasts, with lower RMSE and MAE, but fail to capture the extreme events seen in the VIX. In contrast, the Lévy stable model (Model 2) sacrifices a small amount of forecast accuracy in exchange for a much more faithful representation of the empirical distribution, especially in the tails. This is crucial for risk management applications, where underestimating tail risk can lead to significant financial losses.

The Hull–White model performs the worst overall due to its lack of adequate mean reversion and insufficient modeling of tail behavior. Our results suggest that while simple models may be sufficient for day-to-day forecasting, heavy-tailed models are indispensable for accurately assessing extreme market risk.

A key observation is that point estimates (mean forecasts) may not fully convey a model's ability to capture risk. Although the stable model's point forecasts are slightly less accurate, its predictive distribution (when fully simulated) better accounts for rare but impactful events. For practitioners, this implies that risk assessments (such as Value-at-Risk or Expected Shortfall calculations) should rely on the full predictive distribution rather than just point forecasts.

## 5 | Conclusion

We have demonstrated that incorporating diffusing diffusivity into stochastic volatility models significantly enhances their ability to capture the heavy-tailed behavior of market volatility, as exemplified by the VIX. The Lévy stable model, despite a minor compromise in point forecast accuracy, yields a superior fit to the empirical distribution, especially in the tails, compared to traditional Gaussian models like OU and Hull–White, and even compared to Heston's model.

These findings underscore the importance of model selection based on the intended application: while simple mean-reverting models may suffice for routine forecasts, risk management and derivative pricing require a realistic representation of extreme events. Future work may explore hybrid models that blend the strengths of mean reversion with heavy-tailed noise or extend the analysis to risk-neutral calibration for option pricing.

## Acknowledgements

I would like to thank Rohit Jain and Rui Qin and Ashwani Roy for their work that laid the foundation for diffusing diffusivity in financial modeling, and to the reviewers for their constructive feedback.

## A | References


[1] Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81, 637–654.

[2] Merton, R.C. (1973). Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science*, 4, 141–183.

[3] Hull, J. and White, A. (1987). The Pricing of Options on Assets with Stochastic Volatilities. *Journal of Finance*, 42(2), 281–300.

[4] 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.

[5] Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. Wiley.

[6] Bouchaud, J.-P. and Potters, M. (2000). *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management*. Cambridge University Press.

[7] Mantegna, R.N. and Stanley, H.E. (1999). *Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge University Press.

[8] Risken, H. (1989). *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer.

[9] Sornette, D. (2003). *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press.

[10] Chubynsky, M.V. and Slater, G.W. (2014). Diffusing diffusivity: A model for anomalous, yet Brownian, diffusion. *Physical Review Letters*, 113, 098302.

[11] Stein, E.M. and Stein, J.C. (1991). Stock Price Distributions with Stochastic Volatility: An Analytic Approach. *Review of Financial Studies*, 4, 727–752.

[12] Jain, Rohit and Qin, Rui and Roy, Ashwani, Bridging Diffusing-Diffusivity and Stochastic Volatility *Insights from Physics to Finance (September 17, 2024)* .

[13] Bu, R., Jawadi, F., & Li, Y. (2017). An Empirical Comparison of Transformed Diffusion Models for VIX and VIX Futures. *Journal of International Financial Markets, Institutions and Money*, 46, 116–127.

## B | Fokker–Planck Equation for Volatility Models

For completeness, we include a brief derivation of the Fokker–Planck equation used in the diffusing diffusivity framework. Consider a stochastic process for the diffusivity $D(t)$ governed by:

$$\frac{\partial \Pi(D,t)}{\partial t} = \mathcal{F}_{\mathrm{FP}}\, \Pi(D,t), \tag{B.1}$$

where $\Pi(D,t)$ is the probability density function of $D$, and $\mathcal{F}_{\mathrm{FP}}$ is the Fokker–Planck operator. In Model 1, we have

$$\mathcal{F}_{\mathrm{FP}} = \frac{\partial}{\partial D}\left(4F_D \frac{\partial}{\partial D} + 2\kappa D\right),$$

while for Model 2 incorporating Lévy flights, $\mathcal{F}_{\mathrm{FP}}$ includes fractional derivatives:

$$\mathcal{F}_{\mathrm{FP}} = -\sigma \sec\!\left(\frac{\pi\alpha}{2}\right) \mathcal{D}^{\alpha} + \kappa \frac{\partial}{\partial D} D.$$

The corresponding characteristic function of asset prices then takes the form:

$$\bar{P}(q,t) = \left\langle \exp\left\{ -\frac{1}{2}\left(q^2 + \frac{1}{4}\right) \int_0^t d\tau\, D(\tau) \right\} \right\rangle_D,$$

leading to the full price density:

$$P(S_t, t \mid S_0, 0) = \frac{1}{S_t} \exp\left\{ -\frac{1}{2}\left(\ln \frac{S_t}{S_0} - \mu t\right) \right\} \frac{1}{2\pi} \int_{-\infty}^{\infty} dq\, e^{iq\left(\ln \frac{S_t}{S_0} - \mu t\right)}\, \bar{P}(q,t).$$