



*Research article***Model-implied valuation of variance and scaled absolute-return swaps under double-Heston volatility and stochastic liquidity****Jing Fu¹ and Ke Wang^{2,*}**

¹ School of Mathematics and Statistics, Zhaotong University, Guoxue Road, Zhaoyang District, Zhaotong City, Yunnan Province, China

² School of Science, Civil Aviation Flight University of China (CAFUC), No. 46, 4th Section of Nanchang Road, Guanghan City, Sichuan Province, China

* **Correspondence:** E-mail: kewan@cafuc.edu.cn.

Abstract: We study discretely monitored variance and scaled absolute-return swaps under double-Heston stochastic volatility with an Ornstein-Uhlenbeck liquidity state. Under a selected equivalent martingale measure, a forward characteristic function yields semi-closed-form valuations and connects the discrete contracts to their continuous-monitoring limits under stated moment and uniform-integrability conditions. Analytical decompositions identify the liquidity contribution and maturity-slope criteria. A short-monitoring expansion separates dispersion and leverage effects, explaining how greater volatility of variance can locally lower a discrete variance strike. A single-factor comparator matched in aggregate variance levels and initial local moments isolates maturity-dependent effects of the second persistence scale. The transform uses established affine and quadratic-affine methodology; the analysis concerns the joint role of factor allocation, liquidity transmission, and monitoring. Monte Carlo simulation, Fourier refinement, and independent normal and Heston benchmarks support numerical consistency. Liquidity-on/off comparisons and factor-risk-premium scenarios quantify dependence on the pricing specification. The absolute-return contract is a volatility proxy distinct from a conventional square-root realized-variance swap. The results provide conditional model-implied valuations with illustrative parameters; market calibration and feasible hedging remain separate questions.

Keywords: variance swap; scaled absolute-return swap; double-Heston model; stochastic liquidity; forward characteristic function; discrete sampling

Mathematics Subject Classification: 91G20, 91G60, 60H10

1. Introduction

Variance swaps transfer exposure to realized return variation, and their valuation depends on price dynamics, return convention, and monitoring grid. The Heston model offers a tractable stochastic-volatility framework [1]. [2] derived an exact discrete variance formula for squared simple returns; [3] developed a simplified formula for squared log returns. We study the former convention together with a scaled sum of absolute simple returns. This second contract is a volatility proxy whose continuous limit averages instantaneous volatility over time. A conventional volatility swap instead takes the square root of realized variance, so these payoffs require distinct valuations.

The specific contribution is a joint analysis of factor persistence, stochastic liquidity, and discrete monitoring within a two-Cox-Ingersoll-Ross (CIR)/Ornstein-Uhlenbeck (OU) specification. First, we obtain a forward-interval return transform that combines two CIR variance factors with an OU liquidity state and that retains price-state leverage within each independent block. This extends the single-Heston liquidity setting to two variance persistence scales using established affine and quadratic-affine methods. Second, the transform gives discrete valuations for squared-simple-return variance swaps and scaled absolute-return swaps, with the latter payoff treated explicitly as a distinct contract. Third, the continuous-monitoring limits identify the quadratic liquidity contribution and yield maturity-slope criteria under the stated convergence assumptions. Fourth, a short-monitoring expansion separates dispersion and price-state leverage in the discrete variance correction. Fifth, a single-factor comparator matched in aggregate variance levels and initial local moments isolates the maturity dependence associated with the second persistence scale. These contributions concern the valuation and analytical mechanisms of this specification and build on established transform methods and discrete-monitoring analysis [4].

Two volatility time scales have an economic interpretation: A fast component represents short-lived changes in variance, whereas a slow component represents persistent shifts in the background level of uncertainty. Contracts with different maturities place different weights on these components. Empirical support comes from [5], who found that a model with long-run and short-run volatility components improves option valuation relative to a single-component benchmark, and [6], whose two-factor stochastic-volatility model improves the fit of index options across both moneyness and maturity relative to Heston. These findings motivate distinguishing transitory and persistent variance dynamics when studying the term structure of variance-related payoffs.

Multifactor and double-Heston models have been developed for derivative valuation, variance and volatility swaps, and bounds [7–9]. In particular, [9] studied variance swaps under a rescaled double-Heston model and report calibration to VIX data, linking the two-factor structure directly to variance derivatives. A stochastic long-run variance level provides another way to enrich variance dynamics; [10] used a forward characteristic function to value discretely sampled volatility swaps in that setting. Our two independent CIR factors provide a tractable representation of distinct persistence scales through a sum of two exponential mean-reversion terms. A single factor can match aggregate initial and long-run levels and initial local moments while leaving differences in the intermediate maturity profile. Section 5.5 quantifies those differences with a common monitoring interval. The cited empirical evidence motivates the factor structure; the present comparison examines its valuation implications within the specified model.

Stochastic liquidity supplies a further source of return variation. [11] studied liquidity-adjusted

option valuation, and [12] examined variance and volatility-related contracts with stochastic liquidity. [13] developed Heston option valuation with a stochastic liquidity adjustment. Closely related work by [14] treated European options with stochastic liquidity and a stochastic long-term variance component. Thus, stochastic volatility combined with liquidity, including richer variance dynamics, has established precedents. The closest antecedent for the present swap problem is [15], which combined a single Heston factor with liquidity under discrete monitoring. Relative to this literature, our focus is the two-CIR/OU forward-interval valuation problem, its monitoring limits, and the controlled decomposition of factor and liquidity effects. Each price-state correlation is retained within its block; the CIR-state and liquidity-state innovations are mutually independent. This restriction determines the factorization used below.

Rough and fractional volatility offer a different description of persistence and short-time variation. [16] documented volatility roughness; [17] studied fractional OU modeling and forecasting; [18] developed generalized method of moments (GMM) estimation of roughness; and [19] investigated likelihood estimation for discretely observed fractional OU dynamics. Our finite-dimensional Markov specification permits tractable conditional transforms and distinct mean-reversion scales. It does not reproduce genuinely fractional short-time scaling; selecting between these descriptions empirically requires data beyond the synthetic calculations considered here.

Valuation throughout is conditional on a selected equivalent martingale measure. We use the term *model-implied zero-value strike* to express this scope. Illustrative parameters, controlled liquidity-on/off comparisons, and factor-risk-premium scenarios reveal specification effects. Section 2.5 sets out the separate requirements for parameter identification, market calibration, and feasible hedging, including the distinction between a latent liquidity state and trading frictions.

Section 2 defines the model and its valuation scope. Section 3 derives the forward characteristic function. Section 4 gives the payoff formulae, monitoring limits, and analytical mechanisms. Section 5 presents numerical verification, sensitivity analysis, reduced-model checks, controlled factor comparisons, and risk-premium scenarios. Section 6 states the conclusions and limitations.

2. Model dynamics and risk-neutral measure

2.1. Physical dynamics and liquidity adjustment

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space. The price of the stock in a perfectly liquid market is denoted by S_t^H . Under the physical measure, it follows a double-Heston stochastic volatility specification:

$$\frac{dS_t^H}{S_t^H} = \mu dt + \sqrt{v_{1,t}} dW_{1,t}^{\mathbb{P}} + \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{P}}. \quad (2.1)$$

Here, S_t^H denotes the stock price in a perfectly liquid market, μ is the physical drift, and $v_{1,t}$ and $v_{2,t}$ are the two instantaneous variance factors. The Brownian motions $W_{1,t}^{\mathbb{P}}$ and $W_{2,t}^{\mathbb{P}}$ drive the two price-volatility channels under the physical measure $\mathbb{P}$. The two variance factors are CIR processes:

$$dv_{i,t} = \kappa_i(\theta_i - v_{i,t}) dt + \lambda_i \sqrt{v_{i,t}} dZ_{i,t}^{\mathbb{P}}, \quad i = 1, 2. \quad (2.2)$$

In this specification, κ_i is the physical mean-reversion speed, θ_i is the physical long-run variance level, and λ_i is the volatility-of-variance coefficient of the i th CIR factor. The Brownian motion $Z_{i,t}^{\mathbb{P}}$ drives shocks to the i th variance factor. The physical Feller conditions are

$$2\kappa_i\theta_i \geq \lambda_i^2, \quad i = 1, 2. \quad (2.3)$$

This condition is the usual Feller-type positivity condition for the CIR variance factor under the physical measure. The factors may be interpreted as variance components acting at different time scales.

The liquidity adjustment is introduced after this benchmark price process has been specified. The traded stock price is adjusted by a stochastic liquidity adjustment factor,

$$S_t = \frac{S_t^H}{\zeta_t}. \quad (2.4)$$

Here, $\zeta_t > 0$ is the stochastic liquidity adjustment factor. A higher value lowers the traded price relative to the auxiliary benchmark; the specification does not require $\zeta_t \geq 1$. Only S_t , not S_t^H or ζ_t separately, is declared tradable. The liquidity adjustment factor is specified as

$$\frac{d\zeta_t}{\zeta_t} = \left(\frac{1}{2}\xi^2 l_t^2 - \xi l_t \right) dt - \xi l_t dW_{3,t}^{\mathbb{P}}. \quad (2.5)$$

In this equation, the parameter ξ measures the sensitivity of the traded stock price to the liquidity state, and l_t denotes the market-wide liquidity state variable. The Brownian motion $W_{3,t}^{\mathbb{P}}$ represents price shocks generated by liquidity adjustment. The liquidity state l_t is modeled as a mean-reverting process,

$$dl_t = \alpha(\beta - l_t) dt + \eta dW_{4,t}^{\mathbb{P}}. \quad (2.6)$$

In this OU specification, α is the physical mean-reversion speed of the liquidity state, β is its long-run mean, and η is the volatility of liquidity shocks. The Brownian motion $W_{4,t}^{\mathbb{P}}$ drives the liquidity-state innovation. By Itô's lemma,

$$\frac{dS_t}{S_t} = \left(\mu + \xi l_t + \frac{1}{2}\xi^2 l_t^2 \right) dt + \sqrt{v_{1,t}} dW_{1,t}^{\mathbb{P}} + \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{P}} + \xi l_t dW_{3,t}^{\mathbb{P}}. \quad (2.7)$$

The drift term contains the liquidity adjustment induced by the adjustment factor, and the diffusion terms show that the traded return is exposed to the two Heston variance shocks and to the liquidity-price shock.

2.2. Correlation structure and admissibility

To keep the transform calculation tractable while retaining leverage and liquidity effects, the Brownian covariations are restricted to a block structure.

Assumption 1 (Block correlation structure). *The Brownian covariations under $\mathbb{P}$ are*

$$dW_{i,t}^{\mathbb{P}} dZ_{i,t}^{\mathbb{P}} = \rho_i dt, \quad i = 1, 2, \quad (2.8)$$

$$dW_{3,t}^{\mathbb{P}} dW_{4,t}^{\mathbb{P}} = \rho_l dt, \quad (2.9)$$

and all other Brownian covariations are zero. The initial states $v_{1,0}$, $v_{2,0}$, and l_0 are deterministic.

The coefficient ρ_i is the leverage correlation between the i th stock-price Brownian motion and the i th variance Brownian motion. The coefficient ρ_l measures the instantaneous correlation between the liquidity-price shock and the liquidity-state shock.

This block structure preserves the leverage effects of the two Heston factors and the interaction between price liquidity shocks and liquidity levels. It also makes the state blocks associated with v_1 , v_2 , and l mutually independent conditional on deterministic initial states.

This block representation also provides a convenient basis for the subsequent change of measure. The corresponding independent Brownian basis under $\mathbb{P}$ is written as

$$dW_{1,t}^{\mathbb{P}} = dB_{1,t}^{\mathbb{P}}, \quad dZ_{1,t}^{\mathbb{P}} = \rho_1 dB_{1,t}^{\mathbb{P}} + \sqrt{1 - \rho_1^2} dB_{2,t}^{\mathbb{P}}, \quad (2.10)$$

$$dW_{2,t}^{\mathbb{P}} = dB_{3,t}^{\mathbb{P}}, \quad dZ_{2,t}^{\mathbb{P}} = \rho_2 dB_{3,t}^{\mathbb{P}} + \sqrt{1 - \rho_2^2} dB_{4,t}^{\mathbb{P}}, \quad (2.11)$$

$$dW_{3,t}^{\mathbb{P}} = dB_{5,t}^{\mathbb{P}}, \quad dW_{4,t}^{\mathbb{P}} = \rho_l dB_{5,t}^{\mathbb{P}} + \sqrt{1 - \rho_l^2} dB_{6,t}^{\mathbb{P}}. \quad (2.12)$$

The processes $B_{1,t}^{\mathbb{P}}, \dots, B_{6,t}^{\mathbb{P}}$ form an independent Brownian basis used to express the correlated Brownian motions in block form. This representation is convenient for the multidimensional Girsanov transform.

Assumption 2 (Admissible state dynamics). *The risk-neutral parameters defined in the next subsection are required to satisfy the following admissibility conditions:*

$$\bar{\kappa}_i > 0, \quad \bar{\theta}_i \geq 0, \quad 2\bar{\kappa}_i\bar{\theta}_i \geq \lambda_i^2, \quad i = 1, 2, \quad (2.13)$$

and

$$\bar{\alpha} > 0, \quad \eta \geq 0, \quad |\rho_1|, |\rho_2|, |\rho_l| \leq 1. \quad (2.14)$$

These inequalities require positive risk-neutral mean reversion, non-negative variance and liquidity volatilities, and admissible correlation coefficients.

2.3. Selected equivalent martingale measure

The independent Brownian basis now allows the market prices of risk to be introduced component by component. Let $\Theta_t = (\Theta_{1,t}, \dots, \Theta_{6,t})^\top$ be an admissible market price of risk vector. The equivalent measure $\mathbb{Q}$ is introduced through the multidimensional Girsanov density. The vector Θ_t collects the market prices of risk associated with the six independent Brownian sources.

$$\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\mathcal{F}_t} = \exp \left\{ - \int_0^t \Theta_s^\top dB_s^{\mathbb{P}} - \frac{1}{2} \int_0^t \|\Theta_s\|^2 ds \right\}. \quad (2.15)$$

The density process defines the selected equivalent measure $\mathbb{Q}$ whenever the admissibility and integrability conditions required by Girsanov's theorem are satisfied. Then,

$$dB_{k,t}^{\mathbb{Q}} = dB_{k,t}^{\mathbb{P}} + \Theta_{k,t} dt, \quad k = 1, \dots, 6. \quad (2.16)$$

This relation shows how the Brownian drift adjustments are absorbed into the definition of the $\mathbb{Q}$ -Brownian motions.

The market is incomplete because the number of Brownian risk sources exceeds the number of traded risky assets. Hence, the equivalent martingale measure is not unique. The system of risk-price restrictions does not uniquely determine Θ_i ; rather, it specifies a selected admissible equivalent martingale measure. Existence of such a measure is imposed through the admissibility and true-martingale assumptions.

The stock martingale condition is

$$\sqrt{v_{1,t}}\Theta_{1,t} + \sqrt{v_{2,t}}\Theta_{3,t} + \xi l_t \Theta_{5,t} = \mu + \xi l_t + \frac{1}{2}\xi^2 l_t^2 - r. \quad (2.17)$$

This condition forces the drift of the traded stock price to become the risk-free rate r under the selected pricing measure. The volatility-risk price conditions are

$$\lambda_1 \sqrt{v_{1,t}} \left(\rho_1 \Theta_{1,t} + \sqrt{1 - \rho_1^2} \Theta_{2,t} \right) = \gamma_1 v_{1,t}, \quad (2.18)$$

$$\lambda_2 \sqrt{v_{2,t}} \left(\rho_2 \Theta_{3,t} + \sqrt{1 - \rho_2^2} \Theta_{4,t} \right) = \gamma_2 v_{2,t}. \quad (2.19)$$

These equations specify linear variance-risk premia. The constants γ_i determine how physical mean reversion is shifted under the pricing measure. The liquidity-risk price condition is

$$\eta \left(\rho_l \Theta_{5,t} + \sqrt{1 - \rho_l^2} \Theta_{6,t} \right) = \gamma_l l_t. \quad (2.20)$$

The constant γ_l plays the analogous role for the liquidity-risk premium. Define

$$\bar{\kappa}_i = \kappa_i + \gamma_i, \quad \bar{\theta}_i = \frac{\kappa_i \theta_i}{\bar{\kappa}_i}, \quad i = 1, 2, \quad (2.21)$$

$$\bar{\alpha} = \alpha + \gamma_l, \quad \bar{\beta} = \frac{\alpha \beta}{\bar{\alpha}}. \quad (2.22)$$

The barred quantities denote risk-neutral parameters. In particular, $\bar{\kappa}_i$ and $\bar{\theta}_i$ are the risk-neutral mean-reversion speed and long-run level of the i th variance factor, and $\bar{\alpha}$ and $\bar{\beta}$ are the corresponding risk-neutral liquidity mean-reversion speed and long-run liquidity level.

2.4. Risk-neutral dynamics

Under the selected equivalent measure, the drift terms are adjusted while the quadratic covariations are preserved. Under the selected risk-neutral measure, the valuation dynamics are

$$\begin{cases} \frac{dS_t}{S_t} = r dt + \sqrt{v_{1,t}} dW_{1,t}^{\mathbb{Q}} + \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{Q}} + \xi l_t dW_{3,t}^{\mathbb{Q}}, \\ dv_{i,t} = \bar{\kappa}_i (\bar{\theta}_i - v_{i,t}) dt + \lambda_i \sqrt{v_{i,t}} dZ_{i,t}^{\mathbb{Q}}, \quad i = 1, 2, \\ dl_t = \bar{\alpha} (\bar{\beta} - l_t) dt + \eta dW_{4,t}^{\mathbb{Q}}. \end{cases} \quad (2.23)$$

Under these dynamics, r is the continuously compounded risk-free rate, and all expectations used for pricing are taken under $\mathbb{Q}$. Girsanov's theorem changes drifts but preserves quadratic covariations, so the block correlation structure remains valid under $\mathbb{Q}$. Equivalently, under $\mathbb{Q}$,

$$dW_{i,t}^{\mathbb{Q}} dZ_{i,t}^{\mathbb{Q}} = \rho_i dt, \quad i = 1, 2, \quad (2.24)$$

$$dW_{3,t}^{\mathbb{Q}} dW_{4,t}^{\mathbb{Q}} = \rho_l dt, \quad (2.25)$$

and all other Brownian covariations are zero.

The log-price dynamics provide the state variable used in the forward characteristic function. Let $y_t = \ln S_t$. Then,

$$dy_t = \left[r - \frac{1}{2} (v_{1,t} + v_{2,t} + \xi^2 l_t^2) \right] dt + \sqrt{v_{1,t}} dW_{1,t}^{\mathbb{Q}} + \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{Q}} + \xi l_t dW_{3,t}^{\mathbb{Q}}. \quad (2.26)$$

The quantity $v_{1,t} + v_{2,t} + \xi^2 l_t^2$ is the total instantaneous variance of the log return under the pricing measure.

Assumption 3 (Admissible market price of risk). *The market price of risk vector Θ_t is progressively measurable and satisfies the exponential integrability condition*

$$\mathbb{E}^{\mathbb{P}} \left[\exp \left(\frac{1}{2} \int_0^T \|\Theta_t\|^2 dt \right) \right] < \infty.$$

The selected equivalent measure is further assumed to make $e^{-rt} S_t$ a true martingale on $[0, T]$.

Because the market is incomplete, the equivalent martingale measure is selected rather than uniquely determined. The system of risk-price restrictions does not uniquely determine Θ_t ; rather, it specifies a selected admissible equivalent martingale measure. The true-martingale condition is therefore imposed as an admissibility condition for the selected pricing measure rather than derived from primitive parameter restrictions.

Proposition 1 (Selected equivalent martingale measure). *Under the admissible market price of risk assumption and the martingale restriction imposed on the traded stock, the selected equivalent measure $\mathbb{Q}$ yields the risk-neutral dynamics specified above. This statement is conditional on the existence of a progressively measurable market-price-of-risk process Θ_t satisfying the imposed drift restrictions and the exponential-integrability and true-martingale conditions in Assumption 3.*

Proof. Under the Girsanov transform, the Brownian motions under $\mathbb{Q}$ are obtained by shifting the independent Brownian basis by the admissible risk-price vector Θ_t . The imposed stock martingale condition changes the drift of the traded stock price to r . The linear risk-price restrictions for the two variance factors and the liquidity factor yield the risk-neutral mean-reversion parameters $\bar{\kappa}_i$, $\bar{\theta}_i$, $\bar{\alpha}$, and $\bar{\beta}$. Therefore,

$$dS_t = rS_t dt + S_t \sqrt{v_{1,t}} dW_{1,t}^{\mathbb{Q}} + S_t \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{Q}} + S_t \xi l_t dW_{3,t}^{\mathbb{Q}}.$$

Applying Itô's lemma to $e^{-rt} S_t$ gives

$$d(e^{-rt} S_t) = e^{-rt} S_t \left(\sqrt{v_{1,t}} dW_{1,t}^{\mathbb{Q}} + \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{Q}} + \xi l_t dW_{3,t}^{\mathbb{Q}} \right).$$

The drift therefore vanishes, so the discounted stock price is a local martingale. By the admissibility assumption imposed on the selected equivalent measure, it is a true martingale on the contract horizon, and the swap strikes can be computed as risk-neutral expectations of the corresponding realized payoff functionals. $\square$

Remark 1 (Interpretation and limitations of the OU liquidity state). *The OU process is a signed latent market-wide liquidity state, not a non-negative spread, depth measure, or illiquidity index. With $\eta > 0$, its Gaussian marginal has support on the whole real line. Negative realizations are therefore inherent to the specification: No finite choice of a positive mean and initial state excludes them. Their frequency depends on the mean relative to $\eta/\sqrt{2\bar{\alpha}}$, and a positive-only interpretation would be inappropriate where negative states are economically implausible.*

The contribution $\xi^2 l_t^2$ to instantaneous variance is non-negative for either sign of l_t . Changing the sign reverses the liquidity-price diffusion loading ξl_t but does not create negative variance. It also does not force a negative price: The multiplicative adjustment ζ remains strictly positive. Its value need not exceed one, so the model allows a premium as well as a discount relative to the auxiliary benchmark. One should not interpret every increase of a signed state as a worsening of market liquidity; variance depends on its magnitude and dispersion.

The restrictions $\bar{\alpha} > 0$, non-negative CIR states, admissible correlations, the true-martingale assumption, and the exponential-moment domain must all be respected. A Gaussian state has finite polynomial moments but not unrestricted quadratic exponential moments. Scale normalization is also necessary when the state is inferred from returns, as explained in Section 2.5. A CIR or log-OU liquidity state would enforce non-negativity but would require a new specification and transform analysis; such an alternative is outside the present calculation.

2.5. Valuation scope, identification, and hedging

The declared traded instruments are the stock S and the money-market account $B_t = e^{rt}$. The auxiliary benchmark S^H and the factor ζ are not additional traded securities. The diffusion has risks beyond those spanned by the stock, so the pricing measure is not unique. In this paper, the linear factor-risk premia are a tractable specification choice. They are not estimated from observations, derived from equilibrium or utility, or selected by a minimal-martingale or hedging criterion. Girsanov's theorem supplies the drift representation conditional on admissibility; it supplies neither a market selection rule for $\mathbb{Q}$ nor an executable hedge.

For an integrable terminal realized measure H , a unit-notional swap paying $H - K$ has model value

$$V_0^{\mathbb{Q}}(K) = e^{-rT}(\mathbb{E}_0^{\mathbb{Q}}[H] - K), \quad K(\mathbb{Q}) = \mathbb{E}_0^{\mathbb{Q}}[H]. \quad (2.27)$$

We call $K(\mathbb{Q})$ the *model-implied zero-value strike under the selected equivalent martingale measure*, abbreviated to model-implied strike. Its uniqueness for a fixed $\mathbb{Q}$ is distinct from uniqueness across admissible measures. The existence of an equivalent local martingale measure is connected to absence of free lunches with vanishing risk under the usual semimartingale trading assumptions; it does not establish market completeness. Even in an incomplete market, a particular replicable claim can have a unique price, but replication of the present swap payoffs is not established here. The additional true-martingale assumption rules out a loss of discounted-stock expectation through a strict local martingale and justifies the conditional first-return moment used in Section 3.4. It does not remove the unspanned factor risks. A numerical evaluation of the first-moment identity is an implementation check, not a proof of this assumption.

2.5.1. Parameter identification

The two CIR blocks admit an interchange of factor labels; imposing an ordering such as $\bar{\kappa}_1 > \bar{\kappa}_2$ resolves that labeling ambiguity but does not establish statistical identification. Similar mean-reversion speeds or a negligible factor loading can make the two factors difficult to distinguish. There is also an exact liquidity-scale ambiguity: For $c > 0$, replacing

$$(l_0, \bar{\beta}, \eta, \xi) \quad \text{by} \quad (cl_0, c\bar{\beta}, c\eta, \xi/c)$$

leaves the return process unchanged after setting $l'_t = cl_t$. Returns and return-based derivatives alone therefore cannot identify all four quantities without a scale normalization or a separately measured liquidity proxy. Moreover, continuous variance-swap means do not identify λ_i , ρ_i , or ρ_l . We make no claim that the full parameter vector is identifiable from the numerical examples.

A possible empirical implementation would first define and normalize a signed liquidity proxy and estimate physical dynamics from returns and liquidity observations. It would then fit risk-neutral parameters jointly to option prices across strikes and maturities and, where available, variance-swap quotes, with common return and settlement conventions. The ordering of the CIR speeds, parameter regularization, profile-likelihood or Jacobian diagnostics, multiple starting points, and out-of-sample tests would be needed to assess weak identification. Such a procedure is a proposed calibration design, not an empirical result of this paper. Under the imposed linear-premium family, jointly estimated physical and risk-neutral speeds would imply

$$\gamma_i = \bar{\kappa}_i - \kappa_i, \quad \gamma_l = \bar{\alpha} - \alpha,$$

subject to the testable restrictions $\bar{\kappa}_i \bar{\theta}_i = \kappa_i \theta_i$ and $\bar{\alpha} \bar{\beta} = \alpha \beta$. Estimating only $\mathbb{Q}$ -parameters does not identify physical dynamics or risk premia separately. Section 5.6 instead varies the premia from a fixed synthetic physical reference to display this dependence transparently.

2.5.2. Relation to feasible hedging

Model-implied valuation, consistency with observed derivative prices, and justification by feasible hedging are different requirements. [20] discusses option valuation through explicit buyer and writer portfolios and separates a price-selection criterion from a portfolio-generated interval. For variance contracts, [21] developed model-independent bounds and sub- and super-replicating strategies for discrete monitoring, whereas [22] studied weighted-variance bounds using a finite collection of co-maturing put prices. These approaches connect valuations to the instruments and trading rules declared in the hedging problem.

The conditional connection to our calculation can be stated precisely. Suppose that portfolios with initial costs $\underline{V}_0$ and $\bar{V}_0$ have terminal values satisfying $\underline{H}_T \leq H \leq \bar{H}_T$ and that their discounted gains have zero expectation under a measure $\mathbb{Q}$ consistent with all option prices used in those portfolios. Integrability and admissibility then give

$$\underline{V}_0 \leq e^{-rT} \mathbb{E}_0^{\mathbb{Q}}[H] \leq \bar{V}_0, \quad e^{rT} \underline{V}_0 \leq K(\mathbb{Q}) \leq e^{rT} \bar{V}_0. \quad (2.28)$$

This follows by taking discounted expectations of the pathwise inequalities. The bounds must apply to the same payoff: This paper uses squared *simple* returns, and a bound for log-return or

continuously monitored weighted variance cannot be transferred without checking that convention. With transaction costs or trading constraints, a consistent-price-system or directly feasible-portfolio formulation is needed in place of the frictionless martingale argument above. No observed option panel or corresponding feasible-hedge interval is constructed here; hence, we do not assert that the uncalibrated numerical strikes lie within any particular market-generated interval.

Finally, the OU adjustment changes the stochastic dynamics of a single quoted price. It does not specify bid-ask spreads, transaction costs, depth, price impact, position limits, execution risk, or constraints on rebalancing. The pricing calculation uses an idealized frictionless trading convention for S ; it does not establish implementability in an illiquid market. Contract monitoring dates determine the payoff, whereas trading dates determine a hedging strategy. Discrete monitoring alone supplies no such strategy. Market calibration and portfolio-based validation remain separate extensions of the present conditional valuation framework.

3. Forward characteristic function

The pricing formulae in the next section require the distribution of interval log returns. This section derives the corresponding forward characteristic function before the swap payoff functionals are evaluated.

3.1. Conditional characteristic function

The conditional transform is derived from the log-price dynamics under the selected pricing measure. Under the selected risk-neutral measure in Section 2, the log-price process satisfies

$$dy_t = \left[r - \frac{1}{2} (v_{1,t} + v_{2,t} + \xi^2 l_t^2) \right] dt + \sqrt{v_{1,t}} dW_{1,t}^{\mathbb{Q}} + \sqrt{v_{2,t}} dW_{2,t}^{\mathbb{Q}} + \xi l_t dW_{3,t}^{\mathbb{Q}}. \quad (3.1)$$

Here, r is the risk-free rate, $v_{1,t}$ and $v_{2,t}$ are the two variance factors, l_t is the liquidity state, and ξ measures the strength of liquidity transmission into log-return volatility. Define the instantaneous total variance

$$X_t = v_{1,t} + v_{2,t} + \xi^2 l_t^2. \quad (3.2)$$

The process X_t summarizes the instantaneous contribution of the two stochastic variance factors and the liquidity channel to the quadratic variation of the log price. For $\tau > 0$, the conditional characteristic function of the interval log return is

$$g(\phi, \tau; v_1, v_2, l) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{i\phi(y_{t+\tau} - y_t)} \mid v_{1,t} = v_1, v_{2,t} = v_2, l_t = l \right]. \quad (3.3)$$

Here, ϕ is the Fourier argument; τ is the length of the future interval; and v_1, v_2, l are the current values of the state variables. The function g is the conditional characteristic function of the interval log return.

Assumption 4 (Transform domain). *For every argument used in the pricing formulae, including ϕ , $\phi - i$, $-i$, and $-2i$, the Riccati system derived below admits a finite solution on the relevant time interval. The exponential moments $G_n(-i)$ and $G_n(-2i)$ exist for all monitoring intervals. The solution is understood up to the first possible moment-explosion time, and the transform-domain condition guarantees that all arguments used in the pricing formulae lie inside the admissible domain. For the forward-state*

averaging, the CIR exponential moments must exist, and the OU quadratic exponential must be absolutely integrable: When $q_s > 0$, $\operatorname{Re} D < 1/(2q_s)$. Thus, a merely nonzero denominator $1-2Dq_s$ is not a sufficient integrability condition. Square roots and logarithms are continued from the zero transform argument within this admissible domain.

This condition identifies the complex arguments needed later in the variance and scaled absolute-return swap formulae. The coefficient D and the OU variance q_s are defined below when the quadratic-affine and Gaussian transforms are introduced.

3.2. Forward characteristic function and Riccati system

The conditional transform defined above can be obtained by applying the Feynman-Kac theorem to the Markov process $(y_t, v_{1,t}, v_{2,t}, l_t)$. Because the liquidity contribution enters the log-return variance through the quadratic term $\xi^2 l_t^2$, the natural transform is exponential affine in the two CIR variance factors and exponential quadratic-affine in the liquidity state.

Define the auxiliary transform

$$H(\phi, \tau; y, v_1, v_2, l) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{i\phi y_{t+\tau}} \mid y_t = y, v_{1,t} = v_1, v_{2,t} = v_2, l_t = l \right]. \quad (3.4)$$

Then,

$$g(\phi, \tau; v_1, v_2, l) = e^{-i\phi y} H(\phi, \tau; y, v_1, v_2, l). \quad (3.5)$$

Here, H is the conditional transform of the future log price, and g is the conditional transform of the interval log return. This auxiliary transform allows the Feynman-Kac equation to be written in terms of the full Markov state vector.

Theorem 1 (Conditional forward characteristic function). *Under the risk-neutral dynamics, block correlation structure, and the transform-domain condition, the conditional characteristic function of the interval log return has, on the admissible transform domain, the exponential quadratic-affine representation*

$$g(\phi, \tau; v_1, v_2, l) = \exp \left\{ A(\tau) + B_1(\tau)v_1 + B_2(\tau)v_2 + D(\tau)l^2 + E(\tau)l \right\}, \quad (3.6)$$

where A, B_1, B_2, D, E are deterministic functions of τ satisfying the zero initial conditions

$$A(0) = B_1(0) = B_2(0) = D(0) = E(0) = 0. \quad (3.7)$$

The coefficient functions solve the Riccati system

$$\begin{cases} B'_i = \frac{1}{2} \lambda_i^2 B_i^2 + (\rho_i \lambda_i i\phi - \bar{\kappa}_i) B_i - \frac{1}{2} (\phi^2 + i\phi), & i = 1, 2, \\ D' = 2\eta^2 D^2 + (2\rho_l \xi \eta i\phi - 2\bar{\alpha}) D - \frac{1}{2} \xi^2 (\phi^2 + i\phi), \\ E' = (2\eta^2 D - \bar{\alpha} + \rho_l \xi \eta i\phi) E + 2\bar{\alpha} \bar{\beta} D, \\ A' = i r \phi + \sum_{i=1}^2 \bar{\kappa}_i \bar{\theta}_i B_i + \bar{\alpha} \bar{\beta} E + \eta^2 D + \frac{1}{2} \eta^2 E^2. \end{cases} \quad (3.8)$$

Proof. Define

$$H(\phi, \tau; y, v_1, v_2, l) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{i\phi y_{t+\tau}} \mid y_t = y, v_{1,t} = v_1, v_{2,t} = v_2, l_t = l \right].$$

The state vector $(y_t, v_{1,t}, v_{2,t}, l_t)$ is Markov under $\mathbb{Q}$. Its infinitesimal generator, acting on a sufficiently smooth function f , is given by

$$\begin{aligned} \mathcal{L}f = & \left(r - \frac{1}{2}(v_1 + v_2 + \xi^2 l^2) \right) f_y + \sum_{i=1}^2 \bar{\kappa}_i (\bar{\theta}_i - v_i) f_{v_i} + \bar{\alpha} (\bar{\beta} - l) f_l \\ & + \frac{1}{2}(v_1 + v_2 + \xi^2 l^2) f_{yy} + \frac{1}{2} \sum_{i=1}^2 \lambda_i^2 v_i f_{v_i v_i} + \frac{1}{2} \eta^2 f_{ll} + \sum_{i=1}^2 \rho_i \lambda_i v_i f_{y v_i} + \rho_l \xi \eta l f_{yl}. \end{aligned}$$

By the Feynman-Kac theorem, H satisfies

$$\partial_\tau H = \mathcal{L}H, \quad H(\phi, 0; y, v_1, v_2, l) = e^{i\phi y}.$$

We adopt the exponential quadratic-affine ansatz

$$H = \exp \{ i\phi y + A(\tau) + B_1(\tau)v_1 + B_2(\tau)v_2 + D(\tau)l^2 + E(\tau)l \}$$

with terminal conditions

$$A(0) = B_1(0) = B_2(0) = D(0) = E(0) = 0.$$

Substituting this form into $\partial_\tau H = \mathcal{L}H$ and equating the coefficients of the independent variables y, v_1, v_2, l yields the following Riccati system:

$$B'_i = \frac{1}{2} \lambda_i^2 B_i^2 + (\rho_i \lambda_i i\phi - \bar{\kappa}_i) B_i - \frac{1}{2} (\phi^2 + i\phi), \quad i = 1, 2,$$

$$D' = 2\eta^2 D^2 + (2\rho_l \xi \eta i\phi - 2\bar{\alpha}) D - \frac{1}{2} \xi^2 (\phi^2 + i\phi),$$

$$E' = (2\eta^2 D - \bar{\alpha} + \rho_l \xi \eta i\phi) E + 2\bar{\alpha} \bar{\beta} D,$$

$$A' = ir\phi + \sum_{i=1}^2 \bar{\kappa}_i \bar{\theta}_i B_i + \bar{\alpha} \bar{\beta} E + \eta^2 D + \frac{1}{2} \eta^2 E^2.$$

The right-hand side of this system is locally Lipschitz in (A, B_1, B_2, D, E) on any region where the solution remains finite; hence, a unique local solution exists on the admissible transform domain. The admissibility assumption ensures that this local solution remains finite for all complex arguments required in the pricing formulae. Finally, setting

$$g(\phi, \tau; v_1, v_2, l) = e^{-i\phi y} H(\phi, \tau; y, v_1, v_2, l)$$

gives the claimed exponential quadratic-affine representation. $\square$

3.3. Unconditional forward characteristic function

The conditional transform must still be averaged over the state variables at the beginning of each monitoring interval. For $s = t_{n-1}$, $t = t_n$, and $\Delta = t - s$, define the unconditional forward characteristic function used in the pricing formulae as

$$G_n(\phi) = \mathbb{E}_0^{\mathbb{Q}} \left[e^{i\phi(y_t - y_s)} \right]. \quad (3.9)$$

The function G_n is the unconditional characteristic function of the n th monitoring-interval return and is the object used directly in the pricing formulae. Because the Brownian blocks driving v_1 , v_2 , and l are mutually independent, and initial states are deterministic,

$$G_n(\phi) = e^{A(\Delta)} \prod_{i=1}^2 \mathbb{E}_0^{\mathbb{Q}} \left[e^{B_i(\Delta)v_{i,s}} \right] \cdot \mathbb{E}_0^{\mathbb{Q}} \left[e^{D(\Delta)l_s^2 + E(\Delta)l_s} \right]. \quad (3.10)$$

The product form follows from the block independence of $v_{1,s}$, $v_{2,s}$, and l_s under deterministic initial states. For each CIR factor,

$$\mathbb{E}_0^{\mathbb{Q}}[e^{uv_{i,s}}] = \exp\{a_i(s; u) + b_i(s; u)v_{i,0}\}. \quad (3.11)$$

Here, u is the transform argument, and $v_{i,0}$ is the deterministic initial value of the i th variance factor. The auxiliary functions a_i and b_i are the usual exponential-affine transform coefficients of the CIR process. They satisfy

$$a'_i(s) = \bar{\kappa}_i \bar{\theta}_i b_i(s), \quad b'_i(s) = \frac{1}{2} \lambda_i^2 b_i(s)^2 - \bar{\kappa}_i b_i(s), \quad (3.12)$$

with

$$a_i(0) = 0, \quad b_i(0) = u. \quad (3.13)$$

For the OU liquidity factor,

$$m_s = \bar{\beta} + (l_0 - \bar{\beta})e^{-\bar{\alpha}s}, \quad q_s = \frac{\eta^2}{2\bar{\alpha}} (1 - e^{-2\bar{\alpha}s}). \quad (3.14)$$

The quantities m_s and q_s are respectively the mean and variance of the Gaussian OU liquidity state l_s under $\mathbb{Q}$. Therefore,

$$\mathbb{E} \left[e^{Dl_s^2 + El_s} \right] = (1 - 2Dq_s)^{-1/2} \exp \left\{ Dm_s^2 + Em_s + \frac{q_s(2Dm_s + E)^2}{2(1 - 2Dq_s)} \right\}. \quad (3.15)$$

In this expression, D and E are the liquidity coefficients evaluated at the monitoring interval length, and m_s and q_s are the OU mean and variance at the interval start. The formula is the moment-generating function of a quadratic polynomial in a Gaussian random variable.

Remark 2 (Complex branch and numerical evaluation). *For complex-valued coefficients, branches are obtained by continuous continuation from the zero transform argument inside the admissible moment domain. Riccati coefficients are evaluated by ordinary differential equation (ODE) integration. The numerical implementation also checks the forward-state moment denominators; ODE integration alone does not eliminate branch issues in the subsequent state averaging.*

Remark 3 (Numerical transform-domain check). *Numerical evaluation checks finite Riccati solutions and forward transforms at all required arguments as well as the real parts of the CIR and Gaussian moment denominators. The second exponential moment at $-2i$ requires particular attention. Passing these checks on the reported grid supports numerical admissibility on that grid; it does not prove global moment existence or the physical-to-pricing change-of-measure assumption.*

3.4. Diagnostic identities

The following identities serve as useful checks on both the analytical transform and its numerical implementation.

Remark 4 (Diagnostic identities). *The forward transform satisfies*

$$G_n(0) = 1, \quad G_n(-i) = e^{r(t_n - t_{n-1})}.$$

Here, $G_n(0) = 1$ is the normalization condition, and $G_n(-i)$ is the first exponential moment of the n th interval log return. The second identity follows from $\mathbb{E}_0^{\mathbb{Q}}[e^{x_n}] = e^{r(t_n - t_{n-1})}$ because the discounted stock price is a martingale under $\mathbb{Q}$. These identities are used as numerical diagnostics in Section 5.

The first identity checks normalization of the characteristic function, and the second checks the martingale restriction of the traded stock price.

4. Swap valuation formulae

4.1. Discretely sampled variance and scaled absolute-return swaps

Let the monitoring dates be $0 = t_0 < t_1 < \cdots < t_N = T$, $t_n = nT/N$. T is the contract maturity, N is the number of monitoring intervals, and t_n is the n th monitoring date. Define

$$x_n = y_{t_n} - y_{t_{n-1}}, \quad R_n = \frac{S_{t_n} - S_{t_{n-1}}}{S_{t_{n-1}}} = e^{x_n} - 1. \quad (4.1)$$

The variable x_n is the interval log return, and R_n is the corresponding simple return over $[t_{n-1}, t_n]$. The discretely sampled realized variance measure is

$$RV_a^{(N)} = \frac{d^2}{T} \sum_{n=1}^N R_n^2. \quad (4.2)$$

Here, $RV_a^{(N)}$ denotes the discretely sampled realized variance under the swap quotation convention. The parameter $d = 100$ is the percentage scaling factor used in this convention. The scaled realized absolute-return measure is

$$RV_b^{(N)} = d \sqrt{\frac{\pi}{2NT}} \sum_{n=1}^N |R_n|. \quad (4.3)$$

The quantity $RV_b^{(N)}$ is a scaled realized absolute-return measure, and K_b denotes the strike of a *scaled absolute-return swap*. A conventional volatility swap instead uses a square-root realized-variance payoff, with its own specified return convention. Even holding the simple-return convention fixed, $RV_b^{(N)} \neq \sqrt{RV_a^{(N)}}$ in general. Its continuous realized limit is

$$\frac{d}{T} \int_0^T \sqrt{X_t} dt,$$

which differs from $d \sqrt{T^{-1} \int_0^T X_t dt}$. The time-domain Cauchy-Schwarz inequality gives the former as a lower bound for the latter pathwise, but equality requires constant instantaneous variance almost everywhere. The formulae below do not price the conventional square-root contract.

With these realized measures in place, Eq (2.27) gives the model-implied strikes for the fixed selected measure. We suppress the argument $\mathbb{Q}$ in the notation below; this does not assert uniqueness across measures. The zero-value swap condition gives

$$K_a^{(N)} = \mathbb{E}_0^{\mathbb{Q}}[RV_a^{(N)}], \quad K_b^{(N)} = \mathbb{E}_0^{\mathbb{Q}}[RV_b^{(N)}]. \quad (4.4)$$

Here, $K_a^{(N)}$ and $K_b^{(N)}$ are the model-implied strikes of the discretely sampled variance swap and scaled absolute-return swap, respectively.

Assumption 5 (Moment and inversion conditions). *For each monitoring interval $[t_{n-1}, t_n]$, the exponential moments $G_n(-i)$ and $G_n(-2i)$ exist. The law of x_n admits a Fourier inversion representation, $G_n(\phi)$ is integrable in the required sense, and the shifted transform $G_n(\phi - i)$ is well-defined on the relevant vertical strip. The Fourier integral used for the absolute-return term is finite.*

The variance swap formula is obtained directly from exponential moments of the interval log return.

Theorem 2 (Discretely sampled variance swap). *Under the moment condition above, the model-implied strike of the discretely sampled variance swap is*

$$K_a^{(N)} = \frac{d^2}{T} \sum_{n=1}^N [G_n(-2i) - 2G_n(-i) + 1]. \quad (4.5)$$

In this formula, $G_n(-i)$ and $G_n(-2i)$ are the first and second exponential moments of the n th interval log return expressed through the forward characteristic function.

Proof. For each interval,

$$R_n^2 = (e^{x_n} - 1)^2 = e^{2x_n} - 2e^{x_n} + 1.$$

Moreover, for every real q for which the exponential moment exists,

$$\mathbb{E}_0^{\mathbb{Q}}[e^{qx_n}] = G_n(-iq).$$

Taking expectations and using $q = 1$ and $q = 2$ gives

$$\mathbb{E}_0^{\mathbb{Q}}[R_n^2] = G_n(-2i) - 2G_n(-i) + 1.$$

Summing over monitoring intervals and multiplying by d^2/T proves the formula. $\square$

The scaled absolute-return swap requires an additional Fourier inversion because of the absolute value in the payoff.

Theorem 3 (Discretely sampled scaled absolute-return swap). *Under the moment and inversion conditions above, the model-implied strike associated with the scaled realized absolute-return measure is*

$$K_b^{(N)} = d \sqrt{\frac{\pi}{2NT}} \sum_{n=1}^N L_n, \quad (4.6)$$

where

$$L_n = \mathbb{E}_0^{\mathbb{Q}}[|e^{x_n} - 1|]. \quad (4.7)$$

Here, L_n is the expected absolute simple return over the n th monitoring interval. The equality $e^{x_n} - 1 = R_n$ connects L_n directly to the simple-return payoff. Moreover,

$$L_n = \frac{2}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{G_n(\phi - i) - G_n(\phi)}{i\phi} \right] d\phi. \quad (4.8)$$

In the integral, ϕ is the Fourier variable, and $\operatorname{Re}[\cdot]$ denotes the real part of the complex-valued transform expression. The integrand is understood through its continuous extension at $\phi = 0$.

Proof. The absolute simple return can be split as

$$|e^x - 1| = (1 - e^x)\mathbf{1}_{\{x < 0\}} + (e^x - 1)\mathbf{1}_{\{x > 0\}}.$$

Applying this identity to x_n expresses L_n as the difference between two half-line probabilities and two exponentially tilted half-line expectations. Fourier inversion represents the terms without the exponential tilt through $G_n(\phi)$. The tilted terms use

$$\mathbb{E}_0^{\mathbb{Q}}[e^{x_n} e^{i\phi x_n}] = G_n(\phi - i).$$

Combining the positive and negative half-line inversions, using the symmetry of the real part, and collecting the tilted and untilted components gives

$$L_n = \frac{2}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{G_n(\phi - i) - G_n(\phi)}{i\phi} \right] d\phi.$$

The apt singularity at $\phi = 0$ is removable in the real-valued integrand under the stated moment and inversion conditions. $\square$

4.2. Continuously sampled variance and scaled absolute-return swaps

The continuous-sampling formulae use the quadratic-variation limit for squared returns and the normalized absolute-variation limit for absolute returns. Let

$$X_t = v_{1,t} + v_{2,t} + \xi^2 l_t^2. \quad (4.9)$$

Here, X_t is the instantaneous total variance, formed by the two CIR variance factors and the quadratic liquidity contribution. The continuously sampled realized variance is

$$CRV_a = \frac{d^2}{T} \int_0^T X_t dt. \quad (4.10)$$

The quantity CRV_a is the continuously sampled realized variance under the same scaling factor d and maturity T .

Theorem 4 (Continuously sampled variance swap). *The model-implied strike of the continuously sampled variance swap is*

$$K_a^{(c)} = \frac{d^2}{T} \left[\sum_{i=1}^2 \int_0^T \mathbb{E}_0^{\mathbb{Q}}[v_{i,t}] dt + \xi^2 \int_0^T \mathbb{E}_0^{\mathbb{Q}}[l_t^2] dt \right]. \quad (4.11)$$

The first integral terms collect the expected integrated CIR variances, and the last term collects the expected integrated liquidity contribution.

Proof. The quadratic variation density of y_t is X_t . The drift term does not contribute to quadratic variation. Taking the risk-neutral expectation of CRV_a and interchanging expectation and time integration give the displayed model-implied strike. $\square$

The required first and second moments are explicit:

$$\mathbb{E}_0^Q[v_{i,t}] = \bar{\theta}_i + (v_{i,0} - \bar{\theta}_i)e^{-\bar{\kappa}_i t}. \quad (4.12)$$

Here, $v_{i,0}$ is the initial value of the i th variance factor, and $\bar{\kappa}_i$ and $\bar{\theta}_i$ are its risk-neutral mean-reversion speed and long-run level.

$$\int_0^T \mathbb{E}_0^Q[v_{i,t}] dt = \bar{\theta}_i T + (v_{i,0} - \bar{\theta}_i) \frac{1 - e^{-\bar{\kappa}_i T}}{\bar{\kappa}_i}. \quad (4.13)$$

This expression is the time integral of the expected path of the i th CIR variance factor over the contract horizon. For the OU liquidity state, we use

$$\mathbb{E}_0^Q[l_t^2] = m_t^2 + q_t. \quad (4.14)$$

Here, m_t and q_t are the risk-neutral mean and variance of the OU liquidity state at time t . Therefore,

$$\int_0^T \mathbb{E}_0^Q[l_t^2] dt = \bar{\beta}^2 T + 2\bar{\beta}(l_0 - \bar{\beta}) \frac{1 - e^{-\bar{\alpha} T}}{\bar{\alpha}} + (l_0 - \bar{\beta})^2 \frac{1 - e^{-2\bar{\alpha} T}}{2\bar{\alpha}} + \frac{\eta^2}{2\bar{\alpha}} \left[T - \frac{1 - e^{-2\bar{\alpha} T}}{2\bar{\alpha}} \right]. \quad (4.15)$$

This formula integrates the second moment of the OU liquidity state, with $\bar{\alpha}$, $\bar{\beta}$, η , and l_0 denoting the risk-neutral liquidity mean-reversion speed, long-run level, volatility, and initial state, respectively.

Lemma 1 (Small-interval absolute-return approximation). *Let $\Delta > 0$, and define $R_{t,t+\Delta} = (S_{t+\Delta} - S_t)/S_t$. Assume localization, continuity of the state processes, and sufficient uniform integrability of $R_{t,t+\Delta}/\sqrt{\Delta}$ to pass the localized limit to its conditional absolute first moment. Then,*

$$R_{t,t+\Delta} = \sqrt{X_t} \Delta W + o_p(\sqrt{\Delta}). \quad (4.16)$$

Here, ΔW denotes the leading standardized Brownian increment over the short interval, and $o_p(\sqrt{\Delta})$ is a remainder negligible relative to $\sqrt{\Delta}$ in probability. The corresponding conditional absolute first moment is

$$\mathbb{E}_t^Q[|R_{t,t+\Delta}|] = \sqrt{\frac{2}{\pi}} \sqrt{X_t} \sqrt{\Delta} + o(\sqrt{\Delta}). \quad (4.17)$$

This conditional expectation uses the absolute first moment of the Gaussian leading term with variance $X_t \Delta$.

Proof. Over the short interval $[t, t + \Delta]$, local freezing of the state variables gives a drift contribution of order $O(\Delta)$. The leading martingale contribution has conditional variance $X_t \Delta$, where $X_t = v_{1,t} + v_{2,t} + \xi^2 l_t^2$. Thus, the simple return has the local expansion $R_{t,t+\Delta} = \sqrt{X_t} \Delta W + o_p(\sqrt{\Delta})$. The conditional absolute first moment of the Gaussian leading term is $\sqrt{2/\pi} \sqrt{X_t} \sqrt{\Delta}$. Localization and the imposed integrability conditions justify passing this approximation to the conditional expectation. $\square$

The lemma motivates the continuous-sampling scaled absolute-return swap strike

$$K_b^{(c)} = \frac{d}{T} \int_0^T \mathbb{E}_0^{\mathbb{Q}}[\sqrt{X_t}] dt. \quad (4.18)$$

Here, $K_b^{(c)}$ is the continuous-sampling limit of the scaled absolute-return swap strike.

Theorem 5 (Continuously sampled scaled absolute-return swap). *The continuous-sampling limit associated with the scaled realized absolute-return measure is*

$$K_b^{(c)} = \frac{d}{T} \int_0^T \mathbb{E}_0^{\mathbb{Q}}[\sqrt{X_t}] dt. \quad (4.19)$$

The integral averages the risk-neutral expected instantaneous volatility $\sqrt{X_t}$ over the contract horizon. Furthermore,

$$\mathbb{E}_0^{\mathbb{Q}}[\sqrt{X_t}] = \frac{1}{2\sqrt{\pi}} \int_0^\infty \frac{1 - \mathcal{L}_X(u, t)}{u^{3/2}} du, \quad (4.20)$$

where

$$\mathcal{L}_X(u, t) = \prod_{i=1}^2 \mathbb{E}_0^{\mathbb{Q}}[e^{-u v_{i,t}}] \cdot \mathbb{E}_0^{\mathbb{Q}}[e^{-u \xi^2 l_t^2}]. \quad (4.21)$$

Here, u is the Laplace variable, and $\mathcal{L}_X(u, t)$ is the Laplace transform of the total instantaneous variance X_t . The product separates the two CIR variance-factor Laplace transforms from the OU liquidity quadratic transform. The OU component is

$$\mathbb{E}_0^{\mathbb{Q}}[e^{-u \xi^2 l_t^2}] = (1 + 2u \xi^2 q_t)^{-1/2} \exp\left(-\frac{u \xi^2 m_t^2}{1 + 2u \xi^2 q_t}\right). \quad (4.22)$$

This expression uses the Gaussian distribution of l_t , with mean m_t and variance q_t , after scaling by $u \xi^2$.

Proof. By Lemma 1, the expected scaled absolute-return payoff converges to

$$\frac{d}{T} \int_0^T \mathbb{E}_0^{\mathbb{Q}}[\sqrt{X_t}] dt.$$

For $x \geq 0$, the Laplace identity

$$\sqrt{x} = \frac{1}{2\sqrt{\pi}} \int_0^\infty \frac{1 - e^{-ux}}{u^{3/2}} du$$

holds under the stated integrability conditions. Applying this identity to X_t and taking expectations gives the integral representation for $\mathbb{E}_0^{\mathbb{Q}}[\sqrt{X_t}]$. Block independence factorizes $\mathcal{L}_X(u, t)$ into the two CIR Laplace transforms and the OU quadratic transform. Because l_t is Gaussian with mean m_t and variance q_t , the Gaussian quadratic transform with $D = -u \xi^2$ and $E = 0$ gives the displayed OU component. $\square$

4.3. Convergence from discrete sampling to continuous sampling

The next results make the connection between the discrete formulae and their continuous counterparts explicit. The following two results make precise the sense in which the continuous-sampling strikes are limiting counterparts of the discretely sampled strikes. The first result relies on the standard realized quadratic-variation convergence of continuous semimartingales, and the second uses the local absolute-return approximation from Lemma 1.

Theorem 6 (Convergence of discretely sampled variance swap strikes). *Assume that the monitoring mesh tends to zero and that the family of discretely sampled realized variance variables is uniformly integrable. Then,*

$$K_a^{(N)} \longrightarrow K_a^{(c)} \quad \text{as } N \rightarrow \infty. \quad (4.23)$$

Here, $K_a^{(c)}$ denotes the continuous-sampling variance swap strike, and the limit is taken as the monitoring mesh T/N tends to zero.

Proof. Let $\Delta_N = T/N$, and $t_n = n\Delta_N$. Recall that

$$R_n = \frac{S_{t_n} - S_{t_{n-1}}}{S_{t_{n-1}}} = e^{x_n} - 1, \quad x_n = y_{t_n} - y_{t_{n-1}}.$$

The log-price process is a continuous Itô semimartingale under $\mathbb{Q}$, with quadratic variation

$$[y]_T = \int_0^T X_t dt, \quad X_t = v_{1,t} + v_{2,t} + \xi^2 l_t^2.$$

Hence, the realized quadratic variation of the log returns satisfies

$$\sum_{n=1}^N x_n^2 \longrightarrow [y]_T = \int_0^T X_t dt$$

in probability as $\Delta_N \rightarrow 0$.

It remains to connect the squared simple returns R_n^2 with the squared log returns x_n^2 . Because y_t is continuous on the compact interval $[0, T]$, the mesh maximum satisfies

$$\max_{1 \leq n \leq N} |x_n| \longrightarrow 0$$

in probability. By Taylor's formula,

$$e^{x_n} - 1 = x_n + \frac{1}{2} e^{\theta_n x_n} x_n^2$$

for some random $\theta_n \in (0, 1)$. Therefore,

$$R_n^2 - x_n^2 = (e^{x_n} - 1)^2 - x_n^2 = O(|x_n|^3) + O(x_n^4)$$

uniformly on events where $\max_n |x_n|$ is bounded. Consequently,

$$\sum_{n=1}^N |R_n^2 - x_n^2| \leq C \max_{1 \leq n \leq N} |x_n| \sum_{n=1}^N x_n^2 + C \sum_{n=1}^N x_n^4$$

on such localized events. Because $\sum_n x_n^2$ converges in probability to the finite random variable $\int_0^T X_t dt$ and $\max_n |x_n| \rightarrow 0$ in probability, the right-hand side converges to zero in probability after localization. Removing the localization gives

$$\sum_{n=1}^N R_n^2 - \sum_{n=1}^N x_n^2 \longrightarrow 0$$

in probability. By increasing the localization level and using the imposed uniform integrability condition, the contribution from the complement of the localized events can be made arbitrarily small. Hence,

$$\sum_{n=1}^N R_n^2 \longrightarrow \int_0^T X_t dt$$

in probability.

The discretely sampled realized variance is

$$RV_a^{(N)} = \frac{d^2}{T} \sum_{n=1}^N R_n^2.$$

Thus,

$$RV_a^{(N)} \longrightarrow \frac{d^2}{T} \int_0^T X_t dt = CRV_a$$

in probability. By the uniform integrability assumption imposed on the family $\{RV_a^{(N)}\}_{N \geq 1}$, convergence in probability implies convergence of expectations. Therefore,

$$K_a^{(N)} = \mathbb{E}_0^Q[RV_a^{(N)}] \longrightarrow \mathbb{E}_0^Q[CRV_a] = K_a^{(c)}.$$

This proves the claim. $\square$

Theorem 7 (Conditional convergence of discretely sampled scaled absolute-return swap strikes). *Assume that the small-interval approximation in Lemma 1 holds after localization, with remainder $\sqrt{\Delta_N} \varepsilon_{n,N}$, where $\Delta_N = T/N$, satisfying $\sum_{n=1}^N \mathbb{E}_0^Q[|\varepsilon_{n,N}|] \Delta_N \rightarrow 0$, and that the family of scaled absolute-return sums is uniformly integrable. Conditional on this summed-remainder condition and the stated integrability assumptions,*

$$K_b^{(N)} \longrightarrow K_b^{(c)} \quad \text{as } N \rightarrow \infty. \quad (4.24)$$

Here, $K_b^{(c)}$ denotes the continuous-sampling scaled absolute-return swap strike.

Proof. Let $\Delta_N = T/N$, and $t_n = n\Delta_N$. The discretely sampled scaled absolute-return strike is

$$K_b^{(N)} = d \sqrt{\frac{\pi}{2NT}} \sum_{n=1}^N \mathbb{E}_0^Q[|R_n|].$$

By the tower property,

$$\mathbb{E}_0^Q[|R_n|] = \mathbb{E}_0^Q \left[\mathbb{E}_{t_{n-1}}^Q[|R_n|] \right].$$

Lemma 1 gives, after localization,

$$\mathbb{E}_{t_{n-1}}^Q[|R_n|] = \sqrt{\frac{2}{\pi}} \sqrt{X_{t_{n-1}}} \sqrt{\Delta_N} + \sqrt{\Delta_N} \varepsilon_{n,N},$$

where the remainder satisfies

$$\sum_{n=1}^N \mathbb{E}_0^Q[|\varepsilon_{n,N}|] \Delta_N \longrightarrow 0.$$

This condition is the Riemann-sum form of the negligible-remainder assumption.

Substituting the conditional expansion into $K_b^{(N)}$ gives

$$K_b^{(N)} = d \sqrt{\frac{\pi}{2NT}} \sum_{n=1}^N \mathbb{E}_0^{\mathbb{Q}} \left[\sqrt{\frac{2}{\pi}} \sqrt{X_{t_{n-1}}} \sqrt{\Delta_N} + \sqrt{\Delta_N} \varepsilon_{n,N} \right].$$

Because $\Delta_N = T/N$,

$$\sqrt{\frac{\pi}{2NT}} \sqrt{\frac{2}{\pi}} \sqrt{\Delta_N} = \frac{1}{N} = \frac{\Delta_N}{T}.$$

Therefore, the leading term becomes

$$\frac{d}{T} \sum_{n=1}^N \mathbb{E}_0^{\mathbb{Q}} [\sqrt{X_{t_{n-1}}}] \Delta_N.$$

The remainder term is bounded by

$$d \sqrt{\frac{\pi}{2NT}} \sqrt{\Delta_N} \sum_{n=1}^N \mathbb{E}_0^{\mathbb{Q}} [|\varepsilon_{n,N}|] = \frac{d}{T} \sqrt{\frac{\pi}{2}} \sum_{n=1}^N \mathbb{E}_0^{\mathbb{Q}} [|\varepsilon_{n,N}|] \Delta_N,$$

which tends to zero by the negligible-remainder assumption.

It remains to pass from the Riemann sum to the time integral. The process $X_t = v_{1,t} + v_{2,t} + \xi^2 l_t^2$ has continuous sample paths and finite moments under the admissibility assumptions. Under the moment assumptions and the continuity of the state processes, the function

$$t \mapsto \mathbb{E}_0^{\mathbb{Q}} [\sqrt{X_t}]$$

is continuous after localization and passage to the limit by dominated convergence. In particular, it is Riemann integrable on $[0, T]$. Hence,

$$\sum_{n=1}^N \mathbb{E}_0^{\mathbb{Q}} [\sqrt{X_{t_{n-1}}}] \Delta_N \longrightarrow \int_0^T \mathbb{E}_0^{\mathbb{Q}} [\sqrt{X_t}] dt.$$

Combining the leading-term convergence and the vanishing remainder gives

$$K_b^{(N)} \longrightarrow \frac{d}{T} \int_0^T \mathbb{E}_0^{\mathbb{Q}} [\sqrt{X_t}] dt = K_b^{(c)}.$$

The uniform integrability assumption ensures that the convergence of the scaled absolute-return payoff is compatible with convergence of expectations, completing the proof. $\square$

Remark 5. *The convergence results justify interpreting the continuous-sampling formulae as limiting counterparts of the discretely monitored contracts. The numerical experiments in Section 5 illustrate this limiting behavior.*

4.4. Liquidity transmission, monitoring corrections, and maturity

The continuous limits isolate effects that are obscured when only discretely sampled strikes are compared. Write $m_i(t) = \mathbb{E}_0^{\mathbb{Q}}[v_{i,t}]$, $m_X(t) = \mathbb{E}_0^{\mathbb{Q}}[X_t]$, and let m_t, q_t denote the OU moments defined in Section 3.3.

Proposition 2 (Liquidity contribution and maturity criteria). *Holding the factor dynamics fixed, the continuous variance-strike difference between liquidity-on and liquidity-off specifications is*

$$K_a^{(c)}(\xi) - K_a^{(c)}(0) = \frac{d^2 \xi^2}{T} \int_0^T (m_t^2 + q_t) dt \geq 0. \quad (4.25)$$

The continuous absolute-return strike is also nondecreasing in $|\xi|$ by the pathwise monotonicity of $\sqrt{v_{1,t} + v_{2,t} + \xi^2 l_t^2}$. Both continuous strikes are independent of ρ_1, ρ_2, ρ_l when the remaining risk-neutral parameters are fixed. The continuous variance strike is also independent of λ_1, λ_2 .

For a continuously differentiable maturity dependence and fixed parameters and initial states,

$$\frac{dK_a^{(c)}}{dT} = \frac{d^2}{T} \left(m_X(T) - \frac{1}{T} \int_0^T m_X(t) dt \right), \quad (4.26)$$

$$\frac{dK_b^{(c)}}{dT} = \frac{d}{T} \left(f(T) - \frac{1}{T} \int_0^T f(t) dt \right), \quad f(t) = \mathbb{E}_0^{\mathbb{Q}}[\sqrt{X_t}]. \quad (4.27)$$

Each non-negative-slope condition is therefore exactly a terminal expected level exceeding its time average. A sufficient condition for an increasing continuous variance strike is $v_{i,0} \leq \bar{\theta}_i$ for both factors and $0 \leq l_0 \leq \bar{\beta}$.

Proof. Equation (4.25) follows from linearity of the continuous variance formula. Block independence implies that the marginal law of X_t does not involve any price-state correlation, and the CIR first moments do not involve λ_i . Differentiating the time averages gives (4.26) and (4.27). For the sufficient condition, $m'_i(t) = \bar{\kappa}_i(\bar{\theta}_i - v_{i,0})e^{-\bar{\kappa}_i t} \geq 0$, and

$$\frac{d}{dt} \mathbb{E}_0^{\mathbb{Q}}[l_t^2] = 2\bar{\alpha}(\bar{\beta} - l_0)e^{-\bar{\alpha}t} m_t + \eta^2 e^{-2\bar{\alpha}t} \geq 0.$$

Thus, m_X is nondecreasing, which implies the stated variance-strike criterion. $\square$

The variance sufficient condition does not by itself prove that $f(t)$ is increasing: The dispersion of X_t also enters its square-root expectation. The criterion (4.27) should be used for the absolute-return contract. Neither formula is an unconditional monotonicity theorem for finite N ; varying T with fixed N also changes the interval length.

Why volatility-of-variance can lower a discrete variance strike. The influence of correlation on discrete-monitoring corrections has precedents in stochastic-volatility analysis [4]. The following generator calculation specifies that mechanism for squared simple returns in the present two-CIR/OU model. Let $\Delta = T/N$, and let $\mathcal{L}_s$ denote the generator of the factor state (v_1, v_2, l) . Define

$$C(v_1, v_2, l) = \sum_{i=1}^2 \rho_i \lambda_i v_i + 2\rho_l \xi^3 \eta l^2.$$

Under sufficient moments to take expectations of the second-order generator expansion, with a remainder that is uniform and summable over the monitoring grid, the conditional short-interval expansion is

$$\mathbb{E}_t^{\mathbb{Q}}[R_{t,t+\Delta}^2] = X_t\Delta + \left(r^2 + 2rX_t + \frac{1}{2}X_t^2 + \frac{1}{2}\mathcal{L}_sX_t + C_t\right)\Delta^2 + o(\Delta^2). \quad (4.28)$$

For completeness, apply the full generator to $h(z) = (e^z - 1)^2$ at $z = 0$. The first application gives X ; a second gives $2r^2 + 4rX + X^2 + \mathcal{L}_sX + 2C$. The semigroup expansion then yields (4.28). This argument may first be localized to bounded states; the stated moment and remainder conditions are required when removing the localization.

Because $m'_X(t) = \mathbb{E}_0^{\mathbb{Q}}[\mathcal{L}_sX_t]$, the first correction of the left Riemann sum cancels the integrated $(1/2)\mathcal{L}_sX$ term. Consequently,

$$K_a^{(N)} - K_a^{(c)} = \frac{d^2\Delta}{T} \int_0^T \left[r^2 + 2rm_X(t) + \frac{1}{2}\mathbb{E}_0^{\mathbb{Q}}[X_t^2] + \mathbb{E}_0^{\mathbb{Q}}[C_t]\right] dt + o(\Delta). \quad (4.29)$$

For $r \geq 0$, this expansion separates a non-negative squared-return contribution from leverage terms, which can have either sign. Thus, convergence need not be from above.

For deterministic initial states, the CIR variance is $\text{Var}(v_{i,t}) = \lambda_i^2 h_i(t)$, where

$$h_i(t) = \frac{v_{i,0}}{\bar{\kappa}_i} e^{-\bar{\kappa}_i t} (1 - e^{-\bar{\kappa}_i t}) + \frac{\bar{\theta}_i}{2\bar{\kappa}_i} (1 - e^{-\bar{\kappa}_i t})^2.$$

The derivative of the leading correction in (4.29) with respect to λ_i , keeping all other risk-neutral inputs fixed, is

$$\frac{d^2\Delta}{T} \int_0^T [\lambda_i h_i(t) + \rho_i m_i(t)] dt. \quad (4.30)$$

The positive dispersion contribution competes with negative leverage. At the weekly baseline, this leading-order calculation gives -4.2301 and -3.3463 for factors 1 and 2; direct central differences of the full formula, using a parameter increment 10^{-4} , give -4.1187 and -3.3323 . These are strike derivatives per unit of λ_i , not percentage changes. They explain the displayed local decrease without asserting a global sign theorem.

The same correction places the ρ_i contribution at order Δ , whereas both continuous strikes are exactly independent of ρ_i . Its small weekly effect is therefore consistent with the limiting structure; larger effects at coarser monitoring are not excluded. For the absolute-return limit, the second-order approximation

$$\mathbb{E}[\sqrt{X_t}] \approx \sqrt{m_X(t)} - \frac{\text{Var}(X_t)}{8m_X(t)^{3/2}}$$

gives a concavity-based explanation of sensitivity to dispersion when relative dispersion is small. It is a local approximation, not a proof of monotonicity for the finite-monitoring absolute-return strike.

5. Numerical experiments

5.1. Parameter settings and numerical methods

This section presents numerical experiments for the semi-closed-form pricing formulae. Baseline parameter values are reported in Table 1, which lists the contract parameters, the two Heston variance

factors, and the liquidity factor. Unless otherwise stated, all computations use these baseline values. The setup is synthetic and is not a market calibration. Its purpose is to separate mechanisms within an admissible parameter region. The first and second CIR half-lives are approximately 0.173 and 0.693 years, respectively; the OU half-life is 0.231 years. The initial total instantaneous variance is 0.105, corresponding to 32.40% instantaneous volatility, and the stationary expected total variance is 0.154167. Liquidity contributes 0.01 initially and 0.024167 to the stationary expected variance. These magnitudes describe the numerical scenario and do not establish typical levels for any particular asset. Feller ratios $2\bar{\kappa}_i\bar{\theta}_i/\lambda_i^2$ are 3.556 and 4.5; both exceed one. Without a measured, normalized liquidity proxy, $l_0, \bar{\beta}, \eta$ have no independently identified empirical units.

Table 1. Illustrative baseline parameters used in the numerical experiments.

Group	Parameter	Value
Contract	r, T, d, N	0.02, 1, 100, 52
First Heston factor	$\bar{\kappa}_1, \bar{\theta}_1, \lambda_1, \rho_1, v_{1,0}$	4, 0.04, 0.30, -0.60, 0.035
Second Heston factor	$\bar{\kappa}_2, \bar{\theta}_2, \lambda_2, \rho_2, v_{2,0}$	1, 0.09, 0.20, -0.30, 0.060
Liquidity factor	$\bar{\alpha}, \bar{\beta}, \eta, \xi, \rho_l, l_0$	3, 0.30, 0.20, 0.50, -0.50, 0.20

The following discretization scheme is used only for Monte Carlo validation; the semi-closed-form formulae are evaluated independently through the transform.

We consider a discrete-time simulation framework over $[0, T]$ with N monitoring intervals and M substeps per interval, yielding time step $\delta = T/(NM)$ and grid points $t_m = m\delta$, $m = 0, \dots, NM$. Let $v_{i,m} \approx v_i(t_m)$ ($i = 1, 2$), $l_m \approx l(t_m)$, and $y_m \approx y(t_m)$ denote the two CIR variance factors, the liquidity factor, and the log-price, respectively.

At each step m , we generate six independent standard normal variates $\varepsilon_{1,m}, \dots, \varepsilon_{6,m}$ to construct block-correlated Brownian increments:

$$\begin{aligned}\Delta W_{1,m} &= \sqrt{\delta} \varepsilon_{1,m}, & \Delta Z_{1,m} &= \sqrt{\delta} \left(\rho_1 \varepsilon_{1,m} + \sqrt{1 - \rho_1^2} \varepsilon_{2,m} \right), \\ \Delta W_{2,m} &= \sqrt{\delta} \varepsilon_{3,m}, & \Delta Z_{2,m} &= \sqrt{\delta} \left(\rho_2 \varepsilon_{3,m} + \sqrt{1 - \rho_2^2} \varepsilon_{4,m} \right), \\ \Delta W_{3,m} &= \sqrt{\delta} \varepsilon_{5,m}, & \Delta W_{4,m} &= \sqrt{\delta} \left(\rho_l \varepsilon_{5,m} + \sqrt{1 - \rho_l^2} \varepsilon_{6,m} \right).\end{aligned}$$

These correlations ρ_1, ρ_2, ρ_l are preserved while keeping the underlying shocks independent.

The CIR factors are updated via full-truncation Euler. Defining $v_{i,m}^+ = \max(v_{i,m}, 0)$, we have

$$v_{i,m+1} = v_{i,m} + \bar{\kappa}_i(\bar{\theta}_i - v_{i,m}^+)\delta + \lambda_i \sqrt{v_{i,m}^+} \Delta Z_{i,m}, \quad i = 1, 2.$$

The auxiliary Euler state is not projected after the step; only its positive part enters the next drift, diffusion and return update. This is full truncation rather than a projected Euler scheme. The liquidity factor follows an OU-type Euler discretization,

$$l_{m+1} = l_m + \bar{\alpha}(\bar{\beta} - l_m)\delta + \eta \Delta W_{4,m}.$$

The log-price evolves as

$$y_{m+1} = y_m + \left[r - \frac{1}{2} (v_{1,m}^+ + v_{2,m}^+ + \xi^2 l_m^2) \right] \delta + \sqrt{v_{1,m}^+} \Delta W_{1,m} + \sqrt{v_{2,m}^+} \Delta W_{2,m} + \xi l_m \Delta W_{3,m}.$$

Initialization is $y_0 = 0$ (i.e., $S_0 = 1$), as payoffs depend on returns.

At each monitoring date t_n (index $m = nM$), the simple return over the n th interval is

$$R_n^{\text{MC}} = \exp(y_{nM} - y_{(n-1)M}) - 1.$$

The pathwise realized measures are then

$$RV_{a,\text{MC}}^{(N)} = \frac{d^2}{T} \sum_{n=1}^N (R_n^{\text{MC}})^2, \quad RV_{b,\text{MC}}^{(N)} = d \sqrt{\frac{\pi}{2NT}} \sum_{n=1}^N |R_n^{\text{MC}}|,$$

where $d = 100$ is the percentage scaling factor used in the swap quotation convention. Averaging over P simulated paths yields Monte Carlo estimates

$$\widehat{K}_a^{(N)} = \frac{1}{P} \sum_{p=1}^P RV_{a,p}^{(N)}, \quad \widehat{K}_b^{(N)} = \frac{1}{P} \sum_{p=1}^P RV_{b,p}^{(N)}.$$

For antithetic simulation, let $J = P/2$, and let U_j be the average of the two payoff values in the independent pair j . We use $\widehat{K} = J^{-1} \sum_j U_j$ and confidence intervals $\widehat{K} \pm 1.96 s_U / \sqrt{J}$, where s_U is the sample standard deviation of the pair means. Treating the two members of a pair as independent would give an inappropriate standard-error calculation.

Riccati equations are solved by DOP853 with relative tolerance 10^{-10} and absolute tolerance 10^{-12} . The Fourier integral uses composite 12-point Gauss-Legendre quadrature on panels of width at most four, with cutoff 400; orders 8 and 20 and cutoff 800 provide independent refinement checks. The apt real-valued singularity at zero is avoided by interior quadrature nodes. For the continuous absolute-return strike, the substitution $u = [w/(1-w)]^2$ maps the infinite Laplace integral to $(0, 1)$; the main calculation uses 192 Laplace nodes and 96 time nodes. Finite forward-state moment denominators are checked in addition to solver success.

The Monte Carlo runs use full truncation for the CIR factors, Euler for the OU state and log-price, antithetic variates for all six independent shocks, and fixed NumPy random seeds. Main comparisons use $P = 200,000$ paths and

$$M(N) = \max \left(1, \left\lceil \frac{500}{N} \right\rceil \right).$$

The seed for each main run is $20260909 + N$. Time-step diagnostics use $P = 50,000$ paths and seed $20261000 + 100N + M$. These runs were regenerated with the same implementation and pair-based uncertainty calculation. Python 3.12.10, NumPy 2.4.6 and SciPy 1.17.1 were used; the accompanying scripts and saved result tables specify every numerical input.

[23] established strong order $1/2$ for full-truncation CIR approximations under additional conditions, including a Feller ratio greater than three. The baseline CIR factors satisfy this ratio condition. That result concerns state approximation and does not directly establish a weak order for our unbounded return payoffs in the coupled model. We therefore do not infer an empirical convergence rate from a small, noisy grid study. The Monte Carlo sampling error, of order $J^{-1/2}$ for J independent pairs when the payoff variance is finite, is distinct from the scheme's discretization bias. Both are distinct from the monitoring interval T/N , which is part of the contract.

5.2. Monte Carlo comparison

Throughout the numerical discussion and figures, K_b denotes the scaled absolute-return swap strike; its units of volatility percentage points do not change the payoff definition in Section 4.1.

The first numerical experiment checks whether the semi-closed-form values are consistent with direct simulation.

The monitoring frequencies are

$$N \in \{4, 12, 26, 52, 104, 252\}.$$

Here, N ranges from quarterly monitoring to daily-like monitoring on a 252-trading-day convention. The monitoring frequencies start from $N = 4$, corresponding to quarterly monitoring. This removes the extreme one-period case and is consistent with the literature-parameter comparisons below. At quarterly monitoring, interval returns still contain visible higher-order simple-return effects, whereas higher monitoring frequencies move rapidly toward the continuous-sampling limits.

For each N , the semi-closed-form discrete strike, the Monte Carlo estimate, the 95% confidence interval, the continuous-sampling limit, and the relative error

$$\text{RelErr} = \frac{|K_{\text{formula}} - K_{\text{MC}}|}{|K_{\text{MC}}|}$$

are reported, where K_{formula} is the semi-closed-form strike, and K_{MC} is the Monte Carlo estimate. The comparison block compares formula values with simulation estimates, then examines graphical fit, relative errors, time-step convergence, and characteristic-function diagnostics.

Tables 2 and 3 report the Monte Carlo validation after removing the extreme one-period monitoring case and using a refined simulation grid. The semi-closed-form and Monte Carlo values are close across all monitoring frequencies. The remaining small differences are consistent with Monte Carlo sampling error and simulation time-discretization error. The discrete strikes decrease toward their continuous-sampling counterparts as N increases.

Figures 1 and 2 visualize the same comparison. The semi-closed-form and Monte Carlo curves are close relative to the reported confidence intervals for the monitoring frequencies considered. This is a numerical consistency check within the assumed model. Sampling intervals do not bound the time-discretization bias, and agreement does not validate a market pricing mechanism.

Table 2. Monte Carlo comparison for variance swaps.



N	Formula	MC mean	CI lower	CI upper	RelErr	M
4	1303.4555	1303.4496	1297.7301	1309.1691	0.000004	125
12	1295.6224	1298.4047	1295.1453	1301.6642	0.002143	42
26	1294.3615	1294.7200	1292.5104	1296.9296	0.000277	20
52	1293.9547	1293.4361	1291.8686	1295.0037	0.000401	10
104	1293.7834	1293.8308	1292.7050	1294.9567	0.000037	5
252	1293.6931	1293.2136	1292.4611	1293.9661	0.000371	2

Each row uses 200,000 paths arranged in 100,000 independent antithetic pairs. Confidence intervals use the sample variance of pair averages. Continuous limit: 1293.634244.

Table 3. Monte Carlo comparison for scaled absolute-return swaps.

N	Formula	MC mean	CI lower	CI upper	RelErr	M
4	35.6643	35.6620	35.5816	35.7424	0.000065	125
12	35.6272	35.6674	35.6206	35.7143	0.001128	42
26	35.6178	35.6141	35.5823	35.6460	0.000102	20
52	35.6138	35.6037	35.5811	35.6263	0.000283	10
104	35.6118	35.6090	35.5929	35.6251	0.000079	5
252	35.6106	35.6027	35.5921	35.6132	0.000224	2

Each row uses 200,000 paths arranged in 100,000 independent antithetic pairs. Confidence intervals use the sample variance of pair averages. Continuous limit: 35.609835.

**Figure 1.** Comparison of variance swap strikes under the semi-closed-form formula, Monte Carlo simulation, and the continuous-sampling limit.**Figure 2.** Comparison of scaled absolute-return swap strikes under the semi-closed-form formula, Monte Carlo simulation, and the continuous-sampling limit.

Figures 3 and 4 show the relative errors corresponding to Tables 2 and 3. The errors remain small on the monitoring grid used in the validation. These errors measure formula-simulation agreement rather than the discrete-to-continuous limit itself. The following time-step refinement assesses sensitivity to the simulation mesh; it cannot rule out all remaining discretization bias.




Figure 3. Relative errors for variance swap strikes.




Figure 4. Relative errors for scaled absolute-return swap strikes.



Because Monte Carlo estimates depend on the simulation grid, a separate time-step check is included. To check whether the Monte Carlo discrepancy at the lowest monitoring frequencies considered is driven by the simulation grid, we fix several monitoring frequencies and vary the number of simulation substeps per monitoring interval. The formula values are independent of the simulation grid, and the Monte Carlo estimates should stabilize as the grid is refined. The experiment is mainly a numerical diagnostic rather than an additional pricing formula.

Table 4, Figures 5 and 6 separate simulation time-step effects from monitoring-frequency effects. For the lowest monitoring frequency considered, $N = 4$, each monitoring interval remains relatively long, and the full-truncation Euler scheme can produce visible time-discretization effects. For $N = 4$, one simulation step per quarter is inadequate: The variance estimate is 1319.9488, compared with the formula value 1303.4555. The main quarterly run instead uses $M = 125$, or $\delta = 0.002$ year, and gives 1303.4496 with a 95% interval [1297.7301, 1309.1691]. The refinement estimates fluctuate within sampling error, so their apt slopes are not used to estimate an order of convergence. This indicates that the main pricing formula is not driven by a particular simulation grid, and low-frequency Monte Carlo comparison requires a finer grid and a larger path number. For the scaled absolute-return swap strike, the confidence intervals in the time-step diagnostic are relatively wide compared with the differences among formula benchmarks. Therefore, Figure 6 should be interpreted mainly as a stability check rather than a high-precision comparison.

Table 4. Monte Carlo time-step refinement.

N	M	MC K_a	CI half-width	MC K_b	CI half-width
4	1	1319.9488	11.9941	35.4433	0.1589
4	2	1302.7040	11.6354	35.4947	0.1599
4	5	1305.7876	11.5132	35.6481	0.1616
4	10	1296.4456	11.4721	35.4940	0.1609
12	1	1304.6368	6.5547	35.6044	0.0927
12	2	1299.2904	6.5485	35.6612	0.0933
12	5	1294.5382	6.5078	35.6071	0.0934
12	10	1293.5471	6.4956	35.6023	0.0934
26	1	1294.1959	4.4026	35.5576	0.0631
26	2	1294.6466	4.4205	35.5876	0.0639
26	5	1290.8998	4.4014	35.5570	0.0635
26	10	1293.6875	4.4005	35.5891	0.0637

Each row uses 50,000 paths (25,000 independent antithetic pairs). Formula values are in Tables 2 and 3. These confidence intervals quantify sampling uncertainty, not discretization bias.

**Figure 5.** Monte Carlo time-step convergence for variance swap strikes.**Figure 6.** Monte Carlo time-step convergence for scaled absolute-return swap strikes.

In addition to path simulation, the transform itself is tested through characteristic-function identities. The numerical implementation also checks the transform identities $G_n(0) = 1$ and $G_n(-i) = e^{r\Delta}$. These diagnostics assess the normalization of the forward characteristic function and the martingale restriction under the selected risk-neutral measure.

Table 5 reports two characteristic-function diagnostics. The identity $G_n(0) = 1$ checks the normalization of the forward transform, and $G_n(-i) = e^{r\Delta}$ checks the risk-neutral martingale condition

for the stock price over each monitoring interval. The reported errors are below the numerical threshold, which provides a useful test of the Riccati system and the risk-neutral drift term. These diagnostics do not by themselves test the Fourier and Laplace integrations, but they are necessary consistency checks for the transform implementation. These identities do not verify the second-moment or Fourier calculations and do not prove the assumed true-martingale property. We therefore examine ODE and integration refinement separately.

Table 5. Characteristic-function diagnostic identities.

N	Δ	$\max G_n(0) - 1 $	$\max G_n(-i) - e^{r\Delta} $
4	0.25	1e-14	1e-14
12	0.0833333	1e-14	1e-14
26	0.0384615	1e-14	1e-14
52	0.0192308	1e-14	1e-14
104	0.00961538	1e-14	1e-14
252	0.00396825	1e-14	1e-14

Tables 6–8 report regenerated ODE, Fourier, and Laplace refinement checks. In particular, a cutoff adequate at weekly monitoring need not be adequate at daily monitoring, where the characteristic function decays more slowly along the real Fourier axis. The weekly value is 35.613778 in both the main comparison and the ODE check.

Table 6. Riccati tolerance refinement for weekly monitoring.

Relative tolerance	$K_a^{(N)}$	$K_b^{(N)}$
1e-06	1293.954730512	35.613778033
1e-08	1293.954730512	35.613778033
1e-10	1293.954730512	35.613778033

DOP853 absolute tolerance is 0.01 times its relative tolerance; Fourier cutoff is 400. Values agree to the displayed precision.

Table 7. Fourier cutoff refinement at weekly and daily monitoring.

N	Cutoff	$K_b^{(N)}$	Absolute change
52	50	35.007343583	–
52	100	35.613142217	6.058e-01
52	200	35.613778035	6.358e-04
52	400	35.613778033	1.910e-09
52	800	35.613778033	0.000e+00
252	50	26.399793498	–
252	100	34.579739446	8.180e+00
252	200	35.607865303	1.028e+00
252	400	35.610646049	2.781e-03
252	800	35.610646080	3.070e-08

The daily value at cutoff 200 is insufficiently converged. Cutoff 400 is used for reported calculations; doubling it to 800 changes the daily strike by approximately 3.1×10^{-8} .

Table 8. Continuous absolute-return quadrature refinement.

Time nodes	Laplace nodes	$K_a^{(c)}$	$K_b^{(c)}$
32	64	1293.634244100	35.609835381
64	128	1293.634244100	35.609835381
96	192	1293.634244100	35.609835381
160	320	1293.634244100	35.609835381

The infinite Laplace integral is mapped to $(0, 1)$ by $u = [w/(1 - w)]^2$, and Gauss-Legendre quadrature is used in both variables.

Table 7 shows that the daily value changes from 35.607865 at cutoff 200 to 35.610646 at cutoff 400. Doubling the cutoff again changes it by only approximately 3.1×10^{-8} . The reported results therefore use cutoff 400. These refinement results concern the displayed parameter sets and grids, not a uniform error guarantee across the entire parameter domain.

Independent Fourier benchmark and limitations. Set $\lambda_i = \xi = 0$, and use a constant total variance v . Then, x is normal with mean $\mu_\Delta = (r - v/2)\Delta$ and variance $\sigma_\Delta^2 = v\Delta$, so

$$\mathbb{E}[|e^x - 1|] = e^{r\Delta} \left[2\Phi\left(\frac{\mu_\Delta + \sigma_\Delta^2}{\sigma_\Delta}\right) - 1 \right] - \left[2\Phi\left(\frac{\mu_\Delta}{\sigma_\Delta}\right) - 1 \right],$$

where Φ is the standard normal cumulative distribution function (CDF). This provides an independent absolute-return benchmark without numerical Fourier inversion.

Table 9 verifies the Fourier module in this deterministic-variance reduction. It does not constitute a benchmark reproduction of the stochastic single-Heston liquidity absolute-return strike of Lin and He [15]. An exact cross-check against that implementation has not been performed. The four-model comparison in Section 5.5 uses our common transform code and is likewise not an independent implementation benchmark. Fourier truncation, near-zero evaluation, complex-moment admissibility, and quadrature error remain separate checks for new parameter sets.

Table 9. Independent deterministic-variance benchmark.

N	Exact K_a	Abs. error	Exact K_b	Abs. error
4	1023.786683	0.00e+00	31.684840	7.11e-15
52	1001.809375	0.00e+00	31.627541	5.68e-14
252	1000.373087	0.00e+00	31.623760	9.24e-14

Total variance is 0.1, $r = 0.02$, $T = 1$, and $d = 100$, with $\lambda_i = \xi = 0$. The exact values use lognormal moments and the normal CDF, independently of Riccati/Fourier evaluation.

5.3. Sensitivity analysis

After validating the numerical implementation, the next experiments examine how model parameters affect model-implied strikes. The sensitivity analysis varies one parameter at a time around the illustrative baseline. Only points satisfying the state-admissibility restrictions are displayed; the plotted grids and numerical values are supplied with the code. All mean-reversion comparisons in this subsection keep the corresponding risk-neutral long-run means fixed. The experiments are

grouped according to economic mechanisms rather than listed as an unstructured parameter scan. The volatility-level and mean-reversion channel explains how the two CIR variance factors determine expected future variation. The distributional and leverage channel captures parameters that affect higher-order moments and return asymmetry. The stochastic-liquidity channel isolates the effect of the OU liquidity state and its transmission into traded returns. The contract-design channel studies maturity and monitoring frequency. Because all experiments are local to the illustrative baseline, the results should be interpreted as parameter-region diagnostics rather than global comparative statics. In each sensitivity figure, the left panels report variance swap strikes, and the right panels report scaled absolute-return swap strikes.

Figure 7 contains three related experiments. The mean-reversion panels show that increasing $\bar{\kappa}_1$ or $\bar{\kappa}_2$ raises both strikes in the baseline region because the initial variance levels are below their long-run means. Faster mean reversion therefore moves each variance factor more rapidly toward a higher expected level. The long-run variance panels show the strongest level effect: Increasing $\bar{\theta}_i$ directly raises the expected future mean of $v_{i,t}$ and hence, the expected integrated variance. The initial-variance panels show a short-to-medium horizon effect: A larger $v_{i,0}$ increases near-term realized variation before mean reversion attenuates the impact. The second variance factor has a stronger response in the illustrative baseline because it is more persistent and its long-run gap is larger. Economically, these panels indicate that the double-Heston structure separates short-horizon and long-horizon variance risk.




Figure 7. Volatility-level and mean-reversion channel: Effects of mean-reversion speeds, long-run variance levels, and initial variance levels.

Figure 8 focuses on parameters that mainly affect the distribution of discretely sampled returns. The volatility-of-variance panels show that in the present parameter region, larger λ_i lowers the displayed strikes. This effect is not a universal monotonic relation. Equations (4.29) and (4.30)

explain the competition between dispersion and negative leverage and reproduce the local sign of the variance-strike derivatives. Concavity of the square-root expectation supplies a separate, approximate explanation for the absolute-return limit. The continuous-sampling variance strike depends on $\mathbb{E}[v_{i,t}]$ and is independent of λ_i , whereas the discretely sampled strikes depend on higher-order moments of simple returns and on the nonlinear transform terms. The leverage-correlation panels show a moderate effect from ρ_1 and ρ_2 . These parameters change the dependence between price shocks and variance shocks, affecting skewness and interval-return asymmetry rather than directly affecting the continuous quadratic variation.




Figure 8. Sensitivity to volatility-of-variance parameters and Heston leverage correlations.

Figure 9 shows the effects of liquidity mean reversion and long-run liquidity. In the displayed region, increasing $\bar{\alpha}$ raises both strikes. It speeds the rise of m_t from l_0 toward $\bar{\beta}$, but it also reduces the OU innovation variance q_t at a fixed horizon. The observed sign reflects the balance of these effects; $l_0 < \bar{\beta}$ alone is not a general comparative-statics proof with respect to $\bar{\alpha}$. Increasing $\bar{\beta}$ has a stronger effect because it raises the long-run level of the liquidity state itself. Through the term $\xi^2 l_t^2$, this directly increases total instantaneous variance and hence both swap strikes.




Figure 9. Sensitivity to liquidity mean-reversion speed and long-run liquidity level.

Figure 10 studies the dispersion and transmission of liquidity risk. A larger η increases the variance q_t of the OU liquidity factor and therefore raises $\mathbb{E}[l_t^2] = m_t^2 + q_t$. A larger ξ has an even stronger and visibly convex effect because liquidity enters the log-return variance as $\xi^2 l_t^2$. Economically, ξ measures how strongly market-wide liquidity conditions are transmitted into traded stock returns.




Figure 10. Sensitivity to liquidity volatility and liquidity price sensitivity.

Figure 11 separates the liquidity correlation from the initial liquidity state. The response to ρ_l is small on the weekly grid. Under the block-independent state dynamics, neither continuous strike depends on ρ_l at fixed risk-neutral inputs. Equation (4.29) shows that its leading variance-strike effect is a finite-monitoring correction of order Δ . Its influence can be more visible at coarser monitoring; flatness here is not a claim that correlation is universally irrelevant. The response to l_0 is much stronger because the initial liquidity level changes the mean path of l_t , and the quadratic term l_t^2 amplifies its effect on total variance.




Figure 11. Sensitivity to liquidity correlation and initial liquidity state.

Figure 12 reports the contract-design effects using separate panels for maturity T and monitoring frequency N . Both displayed strikes increase with T for this parameter set. For continuous variance, the initial-state ordering satisfies the sufficient condition in Proposition 2. For the absolute-return limit, the relevant criterion is (4.27), involving $\mathbb{E}[\sqrt{X_t}]$ rather than only expected variance. At finite monitoring, this figure keeps $N = 52$ fixed when varying T , so interval length also changes. No general increasing strike-maturity relation is inferred from the figure. Because the strikes are annualized, the upward maturity slope reflects a change in average expected variation rather than a mechanical accumulation of variance. As N increases, the discretely sampled strikes move toward their continuous-sampling counterparts, consistent with the convergence results in Section 4.3. Higher-order simple-return effects become less important at finer monitoring frequencies, and the realized measures are increasingly governed by the instantaneous variance process X_t .




Figure 12. Contract-design channel: Maturity and monitoring-frequency effects shown on separate axes.

Overall, the sensitivity analysis shows that the dominant economic drivers are the variance levels, the initial variance states, the long-run liquidity level, the initial liquidity state, and the liquidity transmission parameter ξ . Mean-reversion parameters determine how quickly current states affect realized variation. Volatility-of-variance and correlation parameters mainly influence the higher-order distributional features of discretely sampled returns. This classification is consistent with the analytical structure of the model: The continuous-sampling variance limit is governed by $v_{1,t} + v_{2,t} + \xi^2 l_t^2$, whereas discrete sampling and the scaled absolute-return contract depend on the full interval-return distribution.

5.4. Variance swap comparison with reduced models

This subsection examines three variance-strike reductions. Eliminating liquidity and the second CIR factor gives a Heston-only specification, for which we perform an independent algebraic check. Suppressing stochastic variance while retaining liquidity gives the setting of [12]. Removing only the second CIR factor gives the single-Heston liquidity specification of [15]. The latter two are graphical consistency checks based on approximate readings from published figures; they are distinct from the

independent algebraic Heston check.

For the Heston-only reduction, set $v_{2,0} = \bar{\theta}_2 = \lambda_2 = \xi = 0$, and map the first factor by

$$(v_{1,0}, \bar{\theta}_1, \bar{\kappa}_1, \lambda_1, \rho_1) = (v_0, \theta^*, \kappa^*, \sigma_V, \rho).$$

We use the [2] inputs $v_0 = 0.04$, $\theta^* = 0.022$, $\kappa^* = 11.35$, $\sigma_V = 0.618$, $\rho = -0.64$, $r = 0.1$, $T = 1$, and $d = 100$. Both implementations use squared simple returns on $t_j = jT/N$, including $t_0 = 0$, with annualization d^2/T .

We independently implement the closed-form second gross-return moments in [4, Proposition 3.2]. This calculation uses 60-digit arithmetic and bypasses the present Riccati solver and Fourier routines; increasing precision to 90 digits leaves the reported values unchanged. The comparison ODE uses relative tolerance 10^{-12} and absolute tolerance 10^{-14} . The parameter restrictions used in that proposition hold here: $2\kappa^*\theta^*/\sigma_V^2 - 1 > 0$ and $\sigma_V^2 T < 1$. The forward-moment denominators are also positive for every interval.

Table 10 gives a maximum absolute difference of 2.94×10^{-10} between the independent algebraic evaluation and the present ODE calculation. At $N = 4$, both give 263.214308. The difference from the published entry 267.6 is already documented by [4], Remark 3.1 in the author manuscript, which reports 263.2 under the same inputs. We retain the published column for traceability and use the independently evaluated formula as the quantitative benchmark. The quarterly discrepancy is therefore a documented inconsistency in the published numerical entry rather than an effect of coarse monitoring or a failure of the present Heston reduction. We do not infer how that historical entry was produced. The other published values agree after rounding to one decimal place.

Table 10. Heston-only variance strikes: Published values and an independent algebraic check.

N	Zhu-Lian table	Algebraic	Present ODE	Absolute difference
4	267.6 ^a	263.214308	263.214308	2.10×10^{-12}
12	242.7	242.720865	242.720865	1.07×10^{-11}
26	238.6	238.583478	238.583478	3.59×10^{-11}
52	237.1	237.109737	237.109737	5.32×10^{-11}
252	236.1	236.096969	236.096969	2.94×10^{-10}

^a The published quarterly entry is retained for traceability. Bernard and Cui [4, Remark 3.1] reported 263.2 for the same inputs. Algebraic values are our independent evaluation of their simple-return second-moment formula using 60-digit arithmetic. The last column compares that evaluation with the present ODE implementation.

When stochastic volatility is suppressed, and only liquidity risk is retained, our model reduces to the stochastic-liquidity benchmark of [12]. Careful notation mapping is required: Their liquidity price-sensitivity parameter β corresponds to our ξ , and their liquidity volatility parameter ξ corresponds to our η . The reduction is achieved by setting $v_{1,0} = \bar{\theta}_1 = \kappa^2$, $\lambda_1 = 0$, $v_{2,0} = \bar{\theta}_2 = \lambda_2 = 0$, which eliminates the Heston variance factors and leaves a constant variance level κ^2 . The liquidity parameters are mapped as $\xi_{\text{ours}} = \beta_{\text{Xu}}$, $\eta_{\text{ours}} = \xi_{\text{Xu}}$, $\bar{\alpha} = \tilde{\alpha}$, $\bar{\beta} = \tilde{\theta}$, $l_0 = a_0$, $\rho_l = \rho$. Parameter values are $r = 0.05$, $\kappa = 0.1$, $\rho = -0.8$, $\beta = 0.5$, $\tilde{\alpha} = 0.3$, $\tilde{\theta} = 0.2$, $\xi = 0.1$, $a_0 = 0.1$, $T = 1$, and $S_0 = 10$. The initial stock price S_0 does not affect the variance swap strike, as the contract is written on simple returns.

Table 11 provides a graphical consistency check of the stochastic-liquidity-only reduction. “Graph estimate” contains approximate readings from the published figure supplied by the author, and “Present

reduced” contains the strikes computed under the stated parameter mapping. Their approximate agreement supports consistency at graphical precision; the source readings do not provide an exact numerical benchmark.

Table 11. Graphical consistency check for stochastic-liquidity-only variance strikes against Xu et al. [12].

Sampling	N	Graph estimate	Present reduced
Quarterly	4	152.4000	152.1362
Monthly	12	145.8000	145.8373
Fortnightly	26	144.2000	144.1651
Weekly	52	143.4000	143.4515
Daily	252	142.8000	142.8864

Graph estimate values are approximate readings from the published figure supplied by the author. Present reduced values use the stated stochastic-liquidity-only mapping. This table is a graphical consistency check.

When the second Heston factor is eliminated while retaining the first Heston factor and the stochastic liquidity component, our model reduces to the single-Heston stochastic-liquidity benchmark of [15]. The reduction imposes $v_{2,0} = 0$, $\bar{\theta}_2 = 0$, and $\lambda_2 = 0$. The first Heston factor is mapped as $v_{1,0} = v_0$, $\bar{\kappa}_1 = \bar{\kappa}$, $\bar{\theta}_1 = \bar{\theta}$, $\lambda_1 = \lambda$, and $\rho_1 = -0.4$, with parameters $\bar{\kappa} = 5$, $\bar{\theta} = 0.15$, $\lambda = 0.2$, $r = 0.02$, $T = 1$, $v_0 = 0.1$, and $d = 100$. The liquidity factor is specified with $\bar{\alpha} = 3$, $\bar{\beta} = 0.3$, $\eta = 0.2$, $\xi = 0.6$, $l_0 = 0.2$, and $\rho_l = -0.7$. The present computations use $N \in \{4, 12, 26, 52, 252\}$. An exact correspondence between this grid and the supplied graphical benchmark readings has not been established, so the latter are retained only for a graphical consistency check.

Table 12 provides a graphical consistency check against Lin and He [15], retaining one Heston variance factor and the liquidity channel. “Graph estimate” contains approximate readings from the published figure supplied by the author; “Present reduced” contains our computed strikes. The approximate agreement is compatible with this reduction, subject to graphical precision and the unverified monitoring-grid correspondence. It is not an independent numerical reproduction.

Table 12. Graphical consistency check for single-Heston liquidity variance strikes under Lin and He [15] parameters.

Sampling	N	Graph estimate	Present reduced
Quarterly	4	1708.0000	1708.0722
Monthly	12	1689.0000	1688.6049
Fortnightly	26	1685.0000	1684.8359
Weekly	52	1683.5000	1683.4644
Daily	252	1682.5000	1682.4897

Graph estimate values are approximate readings from the published figure supplied by the author. Present reduced values remove the second Heston factor. This table is a graphical consistency check; exact correspondence of the monitoring grids is unverified.

Together, the three mappings isolate the Heston-only, liquidity-only, and single-Heston liquidity reductions. The independent algebraic calculation validates the Heston-only implementation, whereas the other two tables provide graphical consistency checks. All three concern conditional variance-strike calculations.

5.5. Controlled comparisons of factor structure and liquidity

Simply deleting the second CIR factor reduces the total variance level and does not isolate its time-scale contribution. We therefore compare the full double-Heston model (DH) with a single-Heston comparator (SH) that preserves the total initial variance, the total long-run mean, the initial variance drift, the initial variance diffusion variance, and the initial return-variance covariance. Put

$$\begin{cases} v_0^* = v_{1,0} + v_{2,0}, & \theta^* = \bar{\theta}_1 + \bar{\theta}_2, \\ \kappa^* = \frac{\sum_i \bar{\kappa}_i (\bar{\theta}_i - v_{i,0})}{\theta^* - v_0^*}, & (\lambda^*)^2 = \frac{\sum_i \lambda_i^2 v_{i,0}}{v_0^*}, \\ \rho^* = \frac{\sum_i \rho_i \lambda_i v_{i,0}}{\lambda^* v_0^*}. \end{cases} \quad (5.1)$$

This construction is used where the denominator is nonzero, and the resulting speed and other parameters are admissible; it is not a universal aggregation rule. For Table 1, the values are

$$(v_0^*, \theta^*, \kappa^*, \lambda^*, \rho^*) = (0.095, 0.13, 1.428571, 0.241704, -0.431149).$$

The liquidity dynamics are identical in the liquidity-on cases (DH+L and SH+L). Their liquidity-off counterparts set $\xi = 0$. These four specifications lie within the common model family; the matched SH parameters are deliberately reassigned and are not obtained by merely zeroing the second baseline factor.

The analytical distinction is visible in the mean curves:

$$\mathbb{E}[v_{1,t} + v_{2,t}] = \theta^* + \sum_i (v_{i,0} - \bar{\theta}_i) e^{-\bar{\kappa}_i t}, \quad \mathbb{E}[v_t^*] = \theta^* + (v_0^* - \theta^*) e^{-\kappa^* t}.$$

If the two speeds differ and both exponential coefficients are nonzero, linear independence of distinct exponential functions prevents one exponential from matching the two-factor curve at every maturity. The curves nevertheless agree in their initial value, initial derivative, and long-run limit under (5.1). The extra factor therefore allows an additional shape in the expected variance term structure, with a direct implication for variance strikes. This is a consequence of the selected model structure, not evidence that two factors are empirically required.

Table 13 and Figure 13 use a constant weekly interval $\Delta = 1/52$ year, so that $N = 52T$, to separate maturity effects from changes in monitoring interval length. At one year, DH+L gives $K_a = 1293.9547$ and $K_b = 35.6138$, compared with 1309.5717 and 35.7520 for SH+L. The matched single-factor values are higher by approximately 1.207% and 0.388%, respectively, relative to DH+L. The variance difference rises to approximately 20.4548 at two years and falls to 13.0427 at five years, demonstrating a maturity-dependent difference after controlling the variance level and local moments. The sign and size depend on the matching rule and parameter set.

At the same one-year maturity, turning off liquidity in DH changes the strikes from 1293.9547 to 1098.1658 and from 35.6138 to 32.7411. The corresponding differences are 195.7889 variance points and 2.8727 volatility percentage points. The continuous variance difference is 195.5415, exactly accounted for by (4.25); the small remaining difference comes from finite monitoring. These differences hold the other inputs fixed and could change after separate model calibrations.

Table 13. Controlled factor-structure and liquidity comparisons.

T	N	Specification	$K_a^{(N)}$	$K_b^{(N)}$
0.25	13	DH+L	1143.4438	33.6400
0.25	13	DH	1002.8489	31.4797
0.25	13	SH+L	1146.1569	33.6568
0.25	13	SH	1005.5612	31.4939
0.5	26	DH+L	1209.0698	34.5059
0.5	26	DH	1042.2651	31.9948
0.5	26	SH+L	1216.6385	34.5640
0.5	26	SH	1049.8310	32.0483
1	52	DH+L	1293.9547	35.6138
1	52	DH	1098.1658	32.7411
1	52	SH+L	1309.5717	35.7520
1	52	SH	1113.7763	32.8764
2	104	DH+L	1382.0002	36.7447
2	104	DH	1164.2693	33.6311
2	104	SH+L	1402.4550	36.9300
2	104	SH	1184.7150	33.8151
5	260	DH+L	1470.5989	37.8685
5	260	DH	1238.2920	34.6296
5	260	SH+L	1483.6416	37.9542
5	260	SH	1251.3288	34.7042

DH denotes double Heston; SH is the single factor matched by equation (5.1); +L retains the common liquidity block. Without +L, $\xi = 0$. $N/T = 52$ throughout.




Figure 13. Controlled single- and double-Heston comparisons with a common weekly interval. Panels (a,b) show the two model-implied strikes with and without the same liquidity component; Panels (c,d) show matched single-factor minus double-factor differences when liquidity is retained.

5.6. Sensitivity to the specified factor-risk premia

To distinguish a change of pricing specification from a parameter perturbation after $\mathbb{Q}$ has been fixed, we hold a synthetic physical reference constant. Its unbarred factor speeds and levels are set equal to the numerical values in Table 1; zero factor-risk premia then reproduce that table. Initial states, diffusion coefficients, and correlations are fixed. We vary one premium at a time through

$$\bar{\kappa}_i = \kappa_i + \gamma_i, \quad \bar{\theta}_i = \frac{\kappa_i \theta_i}{\kappa_i + \gamma_i}, \quad \bar{\alpha} = \alpha + \gamma_l, \quad \bar{\beta} = \frac{\alpha \beta}{\alpha + \gamma_l}.$$

In contrast to varying $\bar{\kappa}_i$ with $\bar{\theta}_i$ fixed, this experiment preserves each drift intercept $\kappa_i \theta_i$ and $\alpha \beta$. All displayed speeds are positive and the CIR Feller products are unchanged. The transform arguments are checked numerically for finite solutions and admissible forward moment denominators. The interpretation as equivalent martingale measures remains conditional on Assumption 3; these numerical checks are not a proof of a physical-to-pricing density for every scenario.

Table 14 demonstrates that the same physical-reference parameters and initial states do not determine a unique numerical strike when the factor-risk premia are free. For example, $\gamma_1 = -1$ and $+1$ give variance strikes 1381.5058 and 1232.1881, compared with 1293.9547 at zero factor-risk premia. The signs arise from the imposed drift-intercept-preserving parameterization. These scenarios quantify the valuation effect of the specified factor-risk premia.

Table 14. Factor-risk-premium sensitivity from a fixed synthetic physical reference.

Premium changed	$K_a^{(52)}$	$K_b^{(52)}$	$K_a^{(c)}$	$K_b^{(c)}$
$\gamma_1 = \gamma_2 = \gamma_l = 0$	1293.9547	35.6138	1293.6342	35.6098
$\gamma_1 = -1$	1381.5058	36.7781	1381.1701	36.7750
$\gamma_1 = 1$	1232.1881	34.7594	1231.8649	34.7549
$\gamma_2 = -0.25$	1361.6543	36.5087	1361.1637	36.5052
$\gamma_2 = 0.25$	1234.9579	34.8088	1234.7749	34.8044
$\gamma_l = -0.5$	1345.4925	36.3147	1345.0897	36.3112
$\gamma_l = 0.5$	1256.7192	35.0943	1256.4534	35.0900

$T = 1$. Only the named premium changes; all other factor-risk premia remain zero. These are conditional valuation scenarios, not estimated market premia.

6. Conclusions

This paper connects discrete monitoring, liquidity transmission, and variance persistence in the valuation of squared-simple-return variance swaps and scaled absolute-return swaps. An exponential quadratic-affine forward transform gives the model-implied strikes under a selected pricing measure. The continuous limits and monitoring convergence hold under the stated moment and uniform-integrability assumptions. The absolute-return contract remains distinct from a conventional square-root realized-variance volatility swap.

The analytical decomposition isolates the continuous liquidity contribution and the conditions governing maturity slopes. The short-monitoring correction separates dispersion from leverage, explaining the local volatility-of-variance sensitivities. A single-factor comparator matched in

aggregate levels and initial local moments shows how the second persistence scale changes the intermediate maturity profile. Liquidity-on/off and factor-risk-premium scenarios quantify the dependence on specification choices. Monte Carlo calculations with pair-based confidence intervals, Fourier refinement, and independent normal and algebraic Heston benchmarks support the numerical implementation.

The results extend an established transform framework and provide model-specific mechanisms; they do not establish empirical superiority, full parameter identification, a unique market strike, or feasible hedging under trading frictions. The signed Gaussian liquidity state requires a scale normalization for empirical identification. The Xu and Lin-He comparisons retain graphical precision limits, and an independent stochastic Lin-He absolute-return benchmark remains to be constructed. Further work should fit the competing specifications to a common option and liquidity dataset, assess identification and out-of-sample performance, and compare valuations with bounds generated by feasible portfolios under matching payoff and trading conventions.

Author contributions

Jing Fu: Investigation, Software, Supervision, Writing – review and editing. Ke Wang: Formal analysis, Investigation, Software, Conceptualization, Methodology, Writing – review and editing.

Use of Generative-AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

This work was supported by the Fundamental Research Funds for the Central Universities (Grant No. 25CAFUC04077) and the Yunnan Provincial Department of Education Scientific Research Fund Project (Grant No. 2026J1134).

Conflict of interest

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

References

1. S. L. Heston, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Rev. Financ. Stud.*, **6** (1993), 327–343. <https://doi.org/10.1093/rfs/6.2.327>
2. S. P. Zhu, G. H. Lian, A closed-form exact solution for pricing variance swaps with stochastic volatility, *Math. Financ.*, **21** (2011), 233–256. <https://doi.org/10.1111/j.1467-9965.2010.00436.x>
3. S. Rujivan, S. P. Zhu, A simple closed-form formula for pricing discretely-sampled variance swaps under the Heston model, *ANZIAM J.*, **56** (2014), 1–27. <https://doi.org/10.1017/S1446181114000236>

4. C. Bernard, Z. Cui, Prices and asymptotics for discrete variance swaps, *Applied Mathematical Finance*, **21** (2014), 140–173. <https://doi.org/10.1080/1350486X.2013.820524>
5. P. Christoffersen, K. Jacobs, C. Ornathanalai, Y. Wang, Option valuation with long-run and short-run volatility components, *J. Financ. Econ.*, **90** (2008), 272–297. <https://doi.org/10.1016/j.jfineco.2007.12.003>
6. P. Christoffersen, S. Heston, K. Jacobs, The shape and term structure of the index option smirk: why multifactor stochastic volatility models work so well, *Manage. Sci.*, **55** (2009), 1914–1932. <https://doi.org/10.1287/mnsc.1090.1065>
7. A. Issaka, Variance swaps, volatility swaps, hedging and bounds under multi-factor Heston stochastic volatility model, *Stoch. Anal. Appl.*, **38** (2020), 856–874. <https://doi.org/10.1080/07362994.2020.1730903>
8. H. Wu, Z. Jia, S. Yang, C. Liu, Pricing variance swaps under double Heston stochastic volatility model with stochastic interest rate, *Probab. Eng. Inform. Sci.*, **36** (2022), 564–580. <https://doi.org/10.1017/S0269964820000662>
9. Y. Yoon, J. H. Kim, A closed form solution for pricing variance swaps under the rescaled double Heston model, *Comput. Econ.*, **61** (2023), 429–450. <https://doi.org/10.1007/s10614-021-10214-6>
10. X. J. He, S. Lin, Volatility swaps valuation under a modified risk-neutralized Heston model with a stochastic long-run variance level, *ANZIAM J.*, **64** (2022), 250–263. <https://doi.org/10.1017/S144618112200013X>
11. S. P. Feng, M. W. Hung, Y. H. Wang, Option pricing with stochastic liquidity risk: theory and evidence, *J. Financ. Mark.*, **18** (2014), 77–95. <https://doi.org/10.1016/j.finmar.2013.05.002>
12. D. X. Xu, B. Z. Yang, J. H. Kang, N. J. Huang, Variance and volatility swaps valuations with the stochastic liquidity risk, *Physica A*, **566** (2021), 125679. <https://doi.org/10.1016/j.physa.2020.125679>
13. X. J. He, S. Lin, A stochastic liquidity risk model with stochastic volatility and its applications to option pricing, *Stoch. Models*, **41** (2025), 273–292. <https://doi.org/10.1080/15326349.2024.2332326>
14. X. J. He, H. Chen, S. Lin, Analytically pricing European options under two-factor stochastic volatility with stochastic liquidity risks, *J. Futures Markets*, **46** (2026), 1423–1439. <https://doi.org/10.1002/fut.70114>
15. S. Lin, X. J. He, Closed-form formulae for variance and volatility swaps under stochastic volatility with stochastic liquidity risks, *J. Futures Markets*, **44** (2024), 1447–1461. <https://doi.org/10.1002/fut.22531>
16. J. Gatheral, T. Jaisson, M. Rosenbaum, Volatility is rough, *Quant. Financ.*, **18** (2018), 933–949. <https://doi.org/10.1080/14697688.2017.1393551>
17. X. Wang, W. Xiao, J. Yu, Modeling and forecasting realized volatility with the fractional Ornstein-Uhlenbeck process, *J. Econometrics*, **232** (2023), 389–415. <https://doi.org/10.1016/j.jeconom.2021.08.001>

18. A. E. Bolko, K. Christensen, M. S. Pakkanen, B. Veliyev, A GMM approach to estimate the roughness of stochastic volatility, *J. Econometrics*, **235** (2023), 745–778. <https://doi.org/10.1016/j.jeconom.2022.06.009>
19. X. Wang, W. Xiao, J. Yu, C. Zhang, Maximum likelihood estimation of fractional Ornstein-Uhlenbeck process with discretely sampled data, Working Paper 202527.
20. N. Halidias, An overview of financial mathematics with Python codes, *Monte Carlo Methods*, **31** (2025), 279–309. <https://doi.org/10.1515/mcma-2025-2020>
21. D. Hobson, M. Klimmek, Model-independent hedging strategies for variance swaps, *Finance Stoch.*, **16** (2012), 611–649. <https://doi.org/10.1007/s00780-012-0190-3>
22. M. Davis, J. Obłój, V. Raval, Arbitrage bounds for prices of weighted variance swaps, *Math. Financ.*, **24** (2014), 821–854. <https://doi.org/10.1111/mafi.12021>
23. A. Cozma, C. Reisinger, Strong order 1/2 convergence of full truncation Euler approximations to the Cox-Ingersoll-Ross process, *IMA J. Numer. Anal.*, **40** (2020), 358–376. <https://doi.org/10.1093/imanum/dry067>



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)