

Monte Carlo Simulation of European Option Prices under the Heston Stochastic-Volatility Model

Yichen Shao

Maxwell, Syracuse University, Syracuse, USA
yshao28@syr.edu

Abstract. When volatility is not assumed to be constant, there are no closed-form solutions for the prices of European options; hence, numerical estimation is needed. In this paper, the accuracy and convergence of Monte Carlo simulation for pricing European call options under the Heston model is investigated. The empirical volatility process was characterised on the basis of the daily closing prices of the SPDR S&P 500 ETF Trust for the period January 2024 to December 2025 (502 trading days), and a standard literature parameter setting is used for the full market calibration. The paths of stock prices and variances were simulated using the Full Truncation Euler scheme with 252 time steps, while the terminal payoffs were discounted and averaged under the risk-neutral measure. These simulations were performed 20 times for 1,000, 10,000 and 100,000 paths and compared to the semi-analytical Heston price of 9.2425. As the number of paths increased, the mean standard error and the root mean square error decreased in the manner of $1/\sqrt{N}$ rate of contraction with the relative bias being $< 0.25\%$ throughout. The Black–Scholes special-case test was used to correctly replicate the closed-form price to within 0.073%. The estimator is thus correct at this scale but an additional discretisation error remains which cannot be eliminated by any number of additional paths.

Keywords: Option pricing, Heston model, stochastic volatility, Monte Carlo simulation

1. Introduction

Derivatives are extensively employed in financial economics and investment banking for investment, hedging and risk management purposes, with their value directly tied to the value of the underlying asset, making option pricing one of the most relevant topics in financial economics. The price of a European option can be obtained also from the Black–Scholes model [1] as a closed-form solution, under the assumption of constant volatility. However, markets aren't the same every day. Financial time series are in general characterized by stochastic volatility, meaning that the conditional variance changes over time, and high volatility periods are followed by more high volatility periods [2]. The Heston model overcomes this restriction by enabling the variance to be stochastic and mean-reverting and permitting correlation between asset-price and variance shocks [3]. They are better able to reflect the volatility clustering and the leverage effect than a constant-volatility specification, but at the price of analytical and numerical convenience.

Monte Carlo simulation is a general method to estimate expectation if it is not possible to calculate it directly. The way to difficult problems in mathematical physics was first formulated by Metropolis and Ulam [4] and later extended to molecular and chemical systems by systematic random sampling [5]. Boyle applied the technique to finance where he uses the method of averaging the discounted terminal payoffs over a large number of asset-price paths to price options [6]. More recent research has made the computation of path simulation more efficient, e.g., by multilevel approaches which decrease the computational effort for the same accuracy [7], and the estimator theory that standard errors and variance reduction rely on is now well developed [8]. However, the accuracy of a simulated option price is dependent upon the number of paths, the discretisation scheme and the treatment of the numerical boundary problems. Stochastic variance: A naive Euler update may result in a negative variance and any finite number of paths has statistical error. Convergence analysis and uncertainty quantification, and verification against a sound reference, are thus not optional.

In this study, Monte Carlo simulation method is used for calculating the European call option prices under the Heston model. The data of option chains are available in public markets but are limited and have large bid–ask spreads, so a typical set of the parameters from the academic literature is adopted and the exercise is one of assessing the numerical accuracy rather than of empirical market fit. The paths of the stock prices and the paths of the variances are simulated using the Full Truncation Euler scheme [9] that does not allow the simulated variance to be negative and the terminal payoffs are averaged and discounted under the risk-neutral measure. The simulated prices are compared against two independent benchmarks: the semi-analytical Heston price using the same parameters [10] and the Black–Scholes closed-form price in the special case where the variance is fixed. Using convergence analysis, uncertainty quantification, and the aforementioned two benchmark tests, this study quantifies the rate of convergence of the simulation-based price and the percentage of the error that is due to sampling rather than to time discretisation.

2. Method

2.1. Dataset and preprocessing

The daily data for the SPDR S&P 500 ETF Trust (SPY) was downloaded from the Yahoo Finance historical-data service [11]. The sample is a 12-month sample of 2024 (252 trading-day observations) and 2025 (250 observations), for a total of 502 observations. Every record consists of the date, opening price, highest price, lowest price, closing price and trading volume. Two variables were created using the closing price. The continuously compounded daily return is defined as

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \quad (1)$$

where P_t denotes the closing price on day t . A 21-day realised-volatility series was then computed as the rolling standard deviation of daily log returns, annualised by the square root of the number of trading days in a year:

$$\hat{\sigma}_t = \sqrt{252} \sqrt{\frac{1}{20} \sum_{i=0}^{20} (r_{t-i} - \bar{r}_t)^2} \quad (2)$$

where $\bar{r}_t$ is the mean log return over the same 21-day window. This transformation converts short-window daily fluctuations into an annualised volatility estimate. The complete set of observed

and derived variables is summarised in Table 1. Closing prices in the sample ranged from 467.28 to 690.38 US dollars, and the mean 21-day annualised volatility was 14.27%.

Table 1. Dataset variables

Variable	Type	Description
Date	Observed	Trading date
Open	Observed	First reported transaction price
High / Low	Observed	Daily maximum and minimum prices
Close	Observed	Daily closing price used in return calculations
Volume	Observed	Number of shares traded
Log return	Derived	$\ln(P_t/P_{t-1})$, equation (1)
Volatility	Derived	21-day rolling standard deviation $\times \sqrt{252}$, equation (2)

The two series are plotted for the sample period in Figure 1. The overall pattern of the closing price is upward over the two years, with a pronounced dip in the second quarter of 2025, and the annualised volatility series is generally between 6% and 22% with a spike to over 50% in the same drawdown year. The lower panel shows the cluster of large movements that is exactly what a constant-volatility specification cannot produce.




Figure 1. SPY closing price and 21-day annualised volatility, 2024–2025 (picture credit: original)

The observations were sorted by ascending date and then screened before the model was constructed. This study checked for duplicated dates, missing OHLCV fields, non-positive prices, a daily high below the corresponding low, and opening or closing prices falling outside the reported high–low interval. No repeated dates, missing price fields or irregular price ranges were found. One trading-volume value of 26,054 shares was implausibly small relative to neighbouring observations and was identified by applying the interquartile-range rule to log volume; it was replaced with the median volume of the five preceding and five following trading days, 61,585,765 shares. The first log return and the first 20 rolling-volatility values are undefined by construction and were excluded

from the subsequent calculations. Large positive and negative returns were retained, since they reflect genuine market movements rather than recording errors.

2.2. Heston model and numerical procedure

European call prices were estimated under the Heston stochastic-volatility model [3]. Under the risk-neutral measure, the stock price S_t and the instantaneous variance v_t evolve according to

$$dS_t = (r - q)S_t dt + \sqrt{v_t} S_t dW_t^S \quad (3)$$

$$dv_t = \kappa(\theta - v_t)dt + \xi\sqrt{v_t} dW_t^v \quad (4)$$

$$\text{Corr}(dW_t^S, dW_t^v) = \rho dt \quad (5)$$

where r is the risk-free rate, q the continuous dividend yield, κ the speed of mean reversion, θ the long-run variance, ξ the volatility of variance, and ρ the correlation between stock-price and variance shocks. The baseline experiment used $S_0 = K = 100$, $T = 1$ year, $r = 0.03$, $q = 0$, $v_0 = \theta = 0.04$, $\kappa = 2.0$, $\xi = 0.30$ and $\rho = -0.70$. The negative correlation represents the equity leverage effect, and the initial variance $v_0 = 0.04$ corresponds to an initial volatility of 20%.

The one-year horizon was divided into 252 uniformly spaced time steps. Because an ordinary Euler update can drive the simulated variance below zero, paths were generated using the Full Truncation Euler scheme [9], in which the truncated variance

$$v_t^+ = \max(v_t, 0) \quad (6)$$

is substituted into the drift, the diffusion and the stock-price update at every step. Correlated normal shocks were constructed from two independent standard normal draws. The terminal payoff of path i is $\max(S_T^{(i)} - K, 0)$, and the Monte Carlo estimator is the discounted average of terminal payoffs across all N paths,

$$\hat{P} = e^{-rT} \frac{1}{N} \sum_{i=1}^N \max(S_T^{(i)} - K, 0) \quad (7)$$

following the risk-neutral simulation approach of Boyle [6]. Simulations were carried out at 1,000, 10,000 and 100,000 paths, and each path-count experiment was repeated 20 times under a fixed random seed to ensure reproducibility.

2.3. Evaluation metrics

Model performance was assessed against the semi-analytical Heston price computed from the same parameter set, evaluated using the numerically stable formulation of the characteristic function in order to avoid the branch-cut discontinuity documented in the literature [10]. For each repetition j , the absolute error was $|\hat{P}_j - P_{\text{ref}}|$, the relative error $|\hat{P}_j - P_{\text{ref}}|/P_{\text{ref}}$, and the pricing bias $\hat{P}_j - P_{\text{ref}}$. The root mean square error was taken over the $R = 20$ repetitions as

$$RMSE = \sqrt{\frac{1}{R} \sum_{j=1}^R (\hat{P}_j - P_{\text{ref}})^2} \quad (8)$$

Monte Carlo uncertainty was reported as the payoff-based standard error together with the 95% confidence interval $\hat{P} \pm 1.96 \text{ SE}$, following standard estimator theory [8]. A Black–Scholes special-case test, which fixes the variance and sets $\xi = 0$, was used to verify that the simulation code reproduces the closed-form solution [1]. Because the target variable is a continuous option price rather than a binary outcome, classification measures such as the Brier score were not applicable.

3. Results and discussion

3.1. Convergence and pricing accuracy

Table 2 and Figure 2 present the pricing results of the 20 independent repetitions at each simulation size. The semi-analytical Heston price for the baseline parameter set is 9.2425, and this value serves as the reference throughout; every confidence interval reported below is constructed from the mean payoff-based standard error as $\hat{P} \pm 1.96 \text{ SE}$. With 1,000 paths, the mean Monte Carlo estimate was 9.2595, the average standard error 0.3698 and the root mean square error (RMSE) 0.3342. Raising the path count to 10,000 reduced the standard error to 0.1162 and the RMSE to 0.1097. At 100,000 paths the mean price was 9.2548, the standard error fell further to 0.0367 and the RMSE to 0.0394. Uncertainty therefore contracted by approximately a factor of $\sqrt{10}$ for each tenfold increase in the path count, a pattern consistent with the expected sampling behaviour of Monte Carlo estimators [6] and with general computational analysis of simulation-based pricing [7].

Table 2. Heston Monte Carlo convergence results across 20 repetitions

Paths	Mean price	Mean SE	95% CI	RMSE	Relative bias
1,000	9.2595	0.3698	[8.5348, 9.9842]	0.3342	0.184%
10,000	9.2650	0.1162	[9.0373, 9.4927]	0.1097	0.244%
100,000	9.2548	0.0367	[9.1829, 9.3268]	0.0394	0.133%




Figure 2. Monte Carlo convergence under the Heston model (picture credit: original)

The mean prices did not approach the benchmark monotonically: the 10,000-path estimate lay further from the reference value than the 1,000-path estimate. This is not evidence of deterioration. A larger path sample reduces random variation, but individual realised estimates need not all move in the same direction. The precision of the estimator is better read from the width of the confidence interval, which narrowed from 1.4494 at 1,000 paths to 0.1439 at 100,000 paths, and all three

intervals contained the semi-analytical value. Relative bias remained below 0.25% in every experiment. The combined reduction in RMSE, across-run dispersion and confidence-interval width therefore indicates genuine convergence rather than an accidental match in a single run.

3.2. Black–Scholes validation

To test the implementation, a stand-alone special case experiment was set up. By setting the volatility of volatility parameter to zero and fixing the variance at 0.04, the simulated process turns into the constant-volatility case of the Black–Scholes model [1]. As Table 3 shows, the exact call price was 9.4134 and the 100,000-path Monte Carlo mean was 9.4066, an absolute error of 0.0068 and a relative error of 0.073%. Therefore, the exact value is also contained in the Monte Carlo 95% confidence interval [9.3191, 9.4940]. This close agreement verifies the validity of the correlated-shock generation, terminal-payoff calculation, discounting and averaging. The test should not be interpreted as saying that the Black–Scholes and Heston models give the same price in general; the validation explicitly assumes that there is no stochastic variance, while the complete Heston model allows for variance to be mean-reverting, and correlation between asset returns and variance [3].

Table 3. Black–Scholes special-case validation

Method	Price	Mean SE	95% CI	Relative error
Black–Scholes closed form	9.4134	—	—	—
Monte Carlo (100,000 paths)	9.4066	0.0446	[9.3191, 9.4940]	0.073%

3.3. Discussion

The results of the convergence are in line with the theoretical expectations. The remaining differences in pricing are due to time-discretisation error and random sampling error, while the other paths further reduced sampling noise. The Full Truncation Euler scheme avoids negative simulated variance within the square-root, and is effective compared to other Euler corrections [9] but it is still a discrete approximation to a continuous-time process. More paths do not correct the bias of the grid of 252 fixed time steps. It can therefore be inferred that the small positive bias at 100,000 paths could be due to the discretisation component, though small when compared to the option value.

There are several restrictions that limit the scope of the experiment. First, the results were limited to a single at-the-money European call contract, with one maturity and one set of parameters, and comparable accuracy for deep in-the-money or out-of-the-money contracts, longer maturities, or more extreme volatility-of-variance parameter settings has not yet been demonstrated. Second, the parameters were not fitted to a cross section of actually observed option prices, but were tested numerically for accuracy. Third, 20 repetitions are enough to show the convergence pattern but are only a very rough indicator of a convergence in the tail behaviour or confidence-interval coverage.

In future work the time grid would be changed in combination with the count on the path to distinguish discretisation bias from sampling error. An antithetic and control variate can be used to decrease the variance without an increase in the computation time [8]. There are alternative models that can be simulated efficiently, such as Andersen's quadratic-exponential scheme for Heston-type variance processes [12] or exact simulation methods, which have no discrete bias, but are more expensive per path [13]. Multilevel Monte Carlo is based on the idea of many cheap coarse paths and fewer fine paths to achieve a desired level of accuracy at reduced cost [7]. Finally, the calibration of different strikes and maturities would give an indication as to whether this extra

stochastic-volatility structure is useful in practice for pricing or simply a replication of the theoretical benchmark.

4. Conclusion

This study assessed the accuracy of Monte Carlo simulation for pricing European call options under the Heston model using the Full Truncation Euler scheme. Across 20 repetitions at each path count, the mean standard error fell from 0.3698 to 0.0367 and the RMSE from 0.3342 to 0.0394 as the sample grew from 1,000 to 100,000 paths, contracting at approximately the square-root-of-N rate. Every confidence interval contained the semi-analytical benchmark of 9.2425, relative bias never exceeded 0.25%, and the Black–Scholes special case reproduced the closed-form price to within 0.073%. The estimator is therefore accurate and well behaved at this scale. The small residual bias at the largest path count nevertheless points to time-discretisation error, which finer grids or higher-order schemes, rather than additional paths, would be required to remove.

References

- [1] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654. <https://doi.org/10.1086/260062>
- [2] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007. <https://doi.org/10.2307/1912773>
- [3] Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2), 327–343. <https://doi.org/10.1093/rfs/6.2.327>
- [4] Metropolis, N., & Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, 44(247), 335–341. <https://doi.org/10.1080/01621459.1949.10483310>
- [5] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6), 1087–1092. <https://doi.org/10.1063/1.1699114>
- [6] Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics*, 4(3), 323–338. [https://doi.org/10.1016/0304-405X\(77\)90005-8](https://doi.org/10.1016/0304-405X(77)90005-8)
- [7] Giles, M. B. (2008). Multilevel Monte Carlo path simulation. *Operations Research*, 56(3), 607–617. <https://doi.org/10.1287/opre.1070.0496>
- [8] Glasserman, P. (2004). *Monte Carlo methods in financial engineering*. Springer.
- [9] Lord, R., Koekkoek, R., & van Dijk, D. (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2), 177–194. <https://doi.org/10.1080/14697680802392496>
- [10] Albrecher, H., Mayer, P., Schoutens, W., & Tistaert, J. (2007). The little Heston trap. *Wilmott Magazine*, (1), 83–92.
- [11] Yahoo Finance. (2026). SPDR S&P 500 ETF Trust (SPY): Historical data. Retrieved August 6, 2026, from <https://finance.yahoo.com/quote/SPY/history/>
- [12] Andersen, L. B. G. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance*, 11(3), 1–42. <https://doi.org/10.21314/JCF.2008.189>
- [13] Broadie, M., & Kaya, Ö. (2006). Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research*, 54(2), 217–231. <https://doi.org/10.1287/opre.1050.0247>