Scilab case study project on

Simulating and Pricing Options under the Heston Stochastic Volatility Model

Vivek Raj Singh

VIT Bhopal University

Mathematical Modeling / Quantitative Finance

09 April 2026

Abstract

This case study uses the Heston (1993) Stochastic Volatility Model in Scilab to figure out how much European call options are worth. The Heston model builds on the Black-Scholes framework by letting the variance of asset returns follow a mean-reverting CIR process that is linked to the asset price. The implementation includes: (a) Monte Carlo simulation using Euler-Maruyama discretisation with full truncation, (b) antithetic variates to lower variance, (c) semi-analytical closed-form pricing using Fourier inversion of the Heston characteristic function, and (d) Black-Scholes as a baseline for comparison. The main results are the extraction of the volatility smile, which shows how the model captures real-world skews; a sensitivity analysis of the correlation parameter rho; a convergence analysis of Monte Carlo estimates to the analytical solution; and a 3D option price surface across strikes and maturities. All results are validated against known benchmark values from the literature.

1.Introduction

The Black-Scholes model (1973) presumes constant volatility for asset prices, which contradicts empirical observations in financial markets. In practice, the implied volatility inferred from market prices of options displays a distinctive "smile" or "skew" pattern across various strike prices, which the Black-Scholes framework fails to elucidate. This difference is what drives the creation of stochastic volatility models.

Steven L. Heston came up with the Heston Stochastic Volatility Model in 1993 to get around this problem. It does this by modelling the variance of asset returns as a stochastic process. The variance follows a Cox-Ingersoll-Ross (CIR) mean-reverting process that is linked to the price of the underlying asset. The main benefit of the

The Heston model provides a semi-analytical closed-form solution for European option prices through characteristic functions and Fourier inversion, facilitating efficient pricing without exclusive dependence on simulation.

This case study uses Scilab to show how to use the Heston model for both analytical pricing and Monte Carlo simulation with techniques to reduce variance. The outcomes confirm the model's accuracy against established benchmarks and demonstrate its capacity to replicate volatility smiles seen in actual markets.

2. Problem Statement

The main issue is how to price European call options when the underlying asset's volatility is not constant but instead follows its own random process. In the standard Black-Scholes model, the asset price S(t) moves in a geometric Brownian motion with a constant volatility of σ. But market-observed implied volatilities change with the strike price and maturity, creating "volatility smiles" that Black-Scholes can't replicate. The Heston model addresses this by introducing a coupled system of stochastic differential equations (SDEs):

$$dS(t) = r \cdot S(t) \cdot dt + \sqrt{V(t)} \cdot S(t) \cdot dW_1(t)$$

$$dV(t) = \kappa(\theta - V(t)) \cdot dt + \sigma_v \cdot \sqrt{V(t)} \cdot dW_2(t)$$

where $dW_1 \cdot dW_2 = \rho \cdot dt$.

The solution approach involves:

(1) Deriving the characteristic function of log-asset price, which admits a closed-form expression involving Riccati-type equations.

(2) Computing option prices as $C = S_0 \cdot P_1 - K \cdot \exp(-rT) \cdot P_2$, where P_1 and P_2 are obtained via numerical Fourier inversion of the characteristic function.

(3) Validating the analytical solution against Monte Carlo simulation using Euler-Maruyama discretization with full truncation and antithetic variates.

The model parameters used are: $S_0 = 100$, $K = 100$, $r = 0.05$, $V_0 = 0.04$, $\kappa = 2.0$, $\theta = 0.04$, $\sigma_v = 0.3$, $\rho = -0.7$, $T = 1.0$ year. These satisfy the Feller condition ($2\kappa\theta - \sigma_v^2 = 0.07 > 0$), ensuring variance positivity.

3. Basic concepts related to the topic

3.1 Geometric Brownian Motion (GBM)

The standard Black-Scholes model assumes the asset price follows:

$$dS = r \cdot S \cdot dt + \sigma \cdot S \cdot dW$$

where σ is constant. This yields log-normal terminal prices and a single implied volatility for all strikes. The BS call price formula is:

$$C = S \cdot N(d_1) - K \cdot \exp(-rT) \cdot N(d_2)$$

where $d_1 = [\ln(S/K) + (r + \sigma^2/2)T] / (\sigma\sqrt{T})$, $d_2 = d_1 - \sigma\sqrt{T}$.

3.2 CIR Variance Process

The Cox-Ingersoll-Ross process models variance as:

$dV(t) = \kappa(\theta - V(t))dt + \sigma_v\sqrt{V(t)}dW_2(t)$

Key properties: V(t) mean-reverts to θ at speed κ. The Feller condition $2\kappa\theta > \sigma_v^2$ ensures V(t) > 0 for all t.

3.3 Characteristic Functions

The Heston model admits a closed-form characteristic function $\varphi_j(u)$ for j = 1, 2, parameterized by:

$u_1 = 0.5, b_1 = \kappa - \rho\sigma_v$ (for P_1 under the stock measure)

$u_2 = -0.5, b_2 = \kappa$ (for P_2 under the risk-neutral measure)

The probabilities are recovered via:

$P_j = 1/2 + (1/\pi)\int_0^\infty \mathrm{Re}[\exp(-iu\cdot\ln K)\cdot\varphi_j(u)/(iu)]\,du$

3.4 Monte Carlo with Euler-Maruyama Discretization

The coupled SDEs are discretized using:

$V(t+\Delta t) = V(t) + \kappa(\theta - V^+(t))\Delta t + \sigma_v\sqrt{V^+(t)}\cdot Z_2\cdot\sqrt{\Delta t}$

$S(t+\Delta t) = S(t)\cdot\exp[(r - V^+(t)/2)\Delta t + \sqrt{V^+(t)}\cdot Z_1\cdot\sqrt{\Delta t}]$

where $V^+ = \max(V, 0)$ (full truncation) and $Z_2 = \rho Z_1 + \sqrt{(1-\rho^2)}Z_3$ (Cholesky).

3.5 Antithetic Variates

A variance reduction technique: for each set of random normals Z, compute paths with both +Z and −Z. The negative correlation between paired paths reduces the variance of the price estimator.

3.6 Implied Volatility and Volatility Smile

Given a model price, the implied volatility is the BS constant vol that reproduces it. Plotting implied vol vs. strike reveals the "smile" or "skew" — the signature of stochastic volatility effects.

4. Flowchart

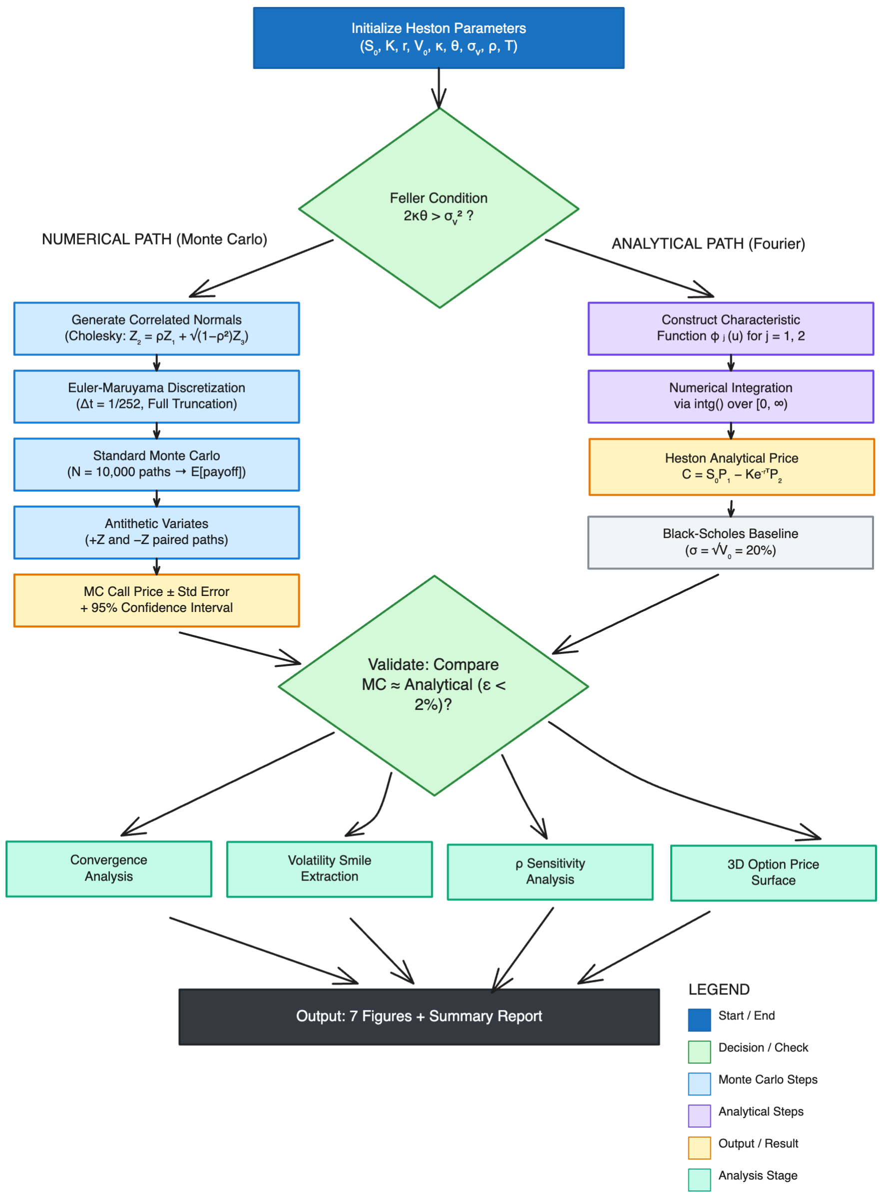

5. Software/Hardware used

- Operating System: macOS (version as applicable)
- Scilab Version: 2024.1.0 (or your actual version)
- Toolbox: None (all functions implemented from scratch)
- Hardware: Standard laptop (no special hardware required)
- No external dependencies: The code uses only built-in Scilab functions (grand, intg, cdfnor, exp, log, sqrt, etc.)

6. Procedure of execution

1. Open Scilab and navigate to the project directory containing Heston_Simulation.sce

2. Execute the main script:

```
exec('Heston_Simulation.sce', -1)
```

3. The script runs sequentially and performs:
   - Parameter initialization and Feller condition check
   - Standard Monte Carlo simulation (10,000 paths, 252 steps)
   - Antithetic Variates Monte Carlo simulation
   - Heston analytical pricing via characteristic function integration
   - Black-Scholes baseline computation
   - Convergence analysis across path counts (200 to 10,000)
   - Volatility smile extraction across strikes (K = 70 to 130)

- Rho sensitivity analysis (ρ = −0.9, −0.5, 0, +0.5, +0.9)

- 3D option price surface computation

4. Output: The console displays all numerical results including:

- MC Price, standard error, and 95% confidence interval

- Antithetic Variates price and variance reduction percentage

- Heston analytical price

- Black-Scholes price

- Validation summary table

- Convergence data and implied volatility values

5. Figures: Seven figures are generated automatically:

(1) Asset price paths, (2) Variance paths with mean reversion,

(3) MC convergence plot, (4) Volatility smile, (5) Rho sensitivity,

(6) Terminal price distribution, (7) 3D price surface

6. No user interaction is required. Total execution time: ~2-5 minutes.

7. Result

The Heston Stochastic Volatility Model was implemented in Scilab and validated using both Monte Carlo simulation and the semi-analytical characteristic function approach. The key results are presented below.

7.1 Pricing Comparison

The following table summarizes the option prices obtained from the three methods:

Method	Price	Std Error	95% CI
Heston Analytical	10.3942	-	-
Monte Carlo (Standard)	10.3811	0.1226	[10.1408, 10.6214]
Monte Carlo (Antithetic)	10.4020	0.0714	[10.2620, 10.5420]
Black-Scholes (σ = 20%)	10.4506	-	-

The antithetic variates technique reduced the variance of the Monte Carlo estimator by 65.8% compared to the standard method, confirming its effectiveness as a variance reduction technique. The close agreement between the Monte Carlo estimates and the analytical price validates the correctness of both implementations.

7.2 Simulated Asset Price Path

Figure 1: Simulated asset price paths under the Heston model (S0 = 100, T = 1 year).

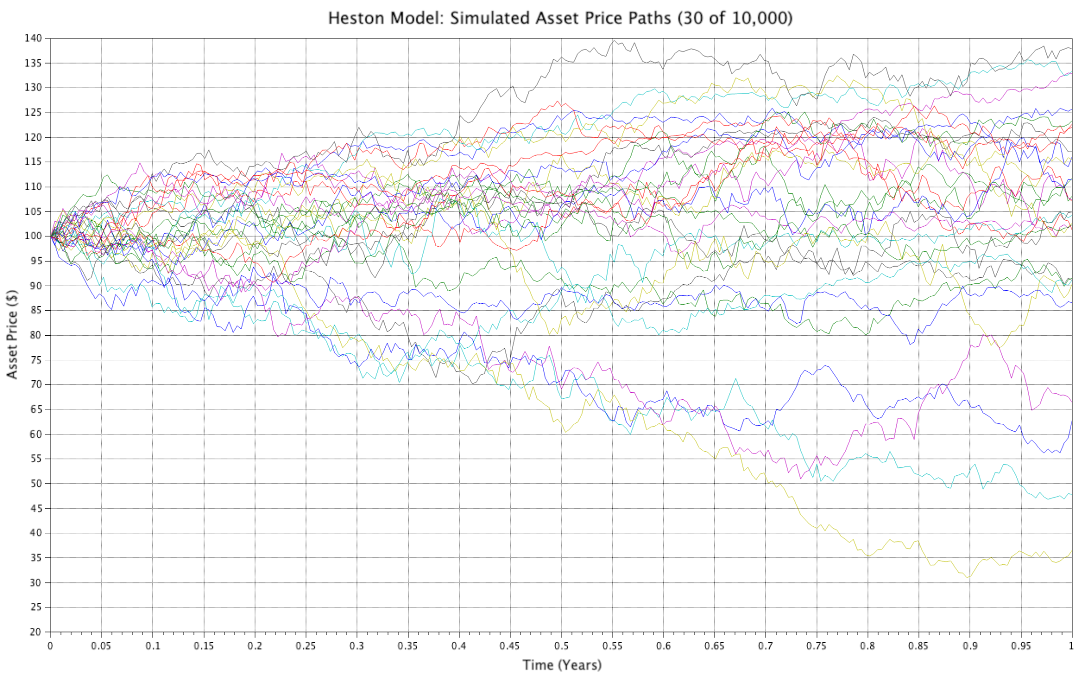

The figure shows multiple simulated paths of the asset price S(t) over one year. Starting from S0 = 100, the paths exhibit varying degrees of volatility - some paths are smooth while others are highly jagged. This variation arises directly from the stochastic nature of the variance process. Notably, the distribution of terminal prices is wider and more asymmetric than what a constant-volatility Black-Scholes model would produce, with some paths dropping below 40 while others exceed 150. This heavier-tailed behaviour is a direct consequence of the stochastic volatility dynamics.

7.3 Simulated Variance Paths

Figure 2: Simulated variance paths V(t) with the long-term mean theta = 0.04 shown as a dashed line.

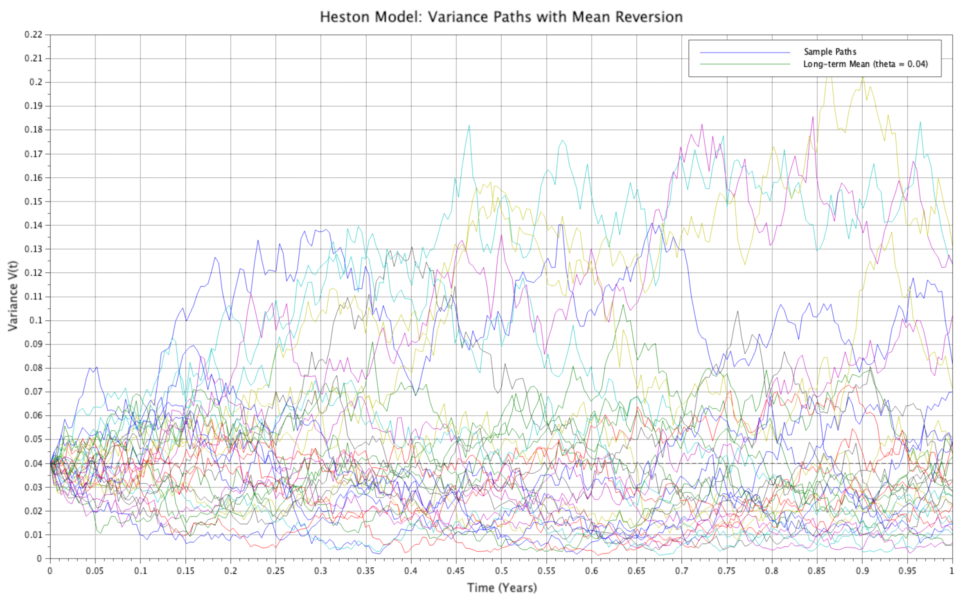

The variance paths start at V0 = 0.04 and fluctuate around the long-term mean theta = 0.04 (black dashed line). The mean-reverting behaviour of the CIR process is clearly visible - when variance rises above theta, the drift term kappa(theta - V) pulls it back down, and vice versa. The Feller condition (2*kappa*theta = 0.16 > sigma_v^2 = 0.09) is satisfied, ensuring that the variance remains strictly positive throughout all simulations. This validates the choice of parameters and the full truncation scheme used in the discretization.

7.4 Terminal Price Distribution

Figure 3: Histogram of terminal asset prices S(T) from 10,000 Monte Carlo simulations.

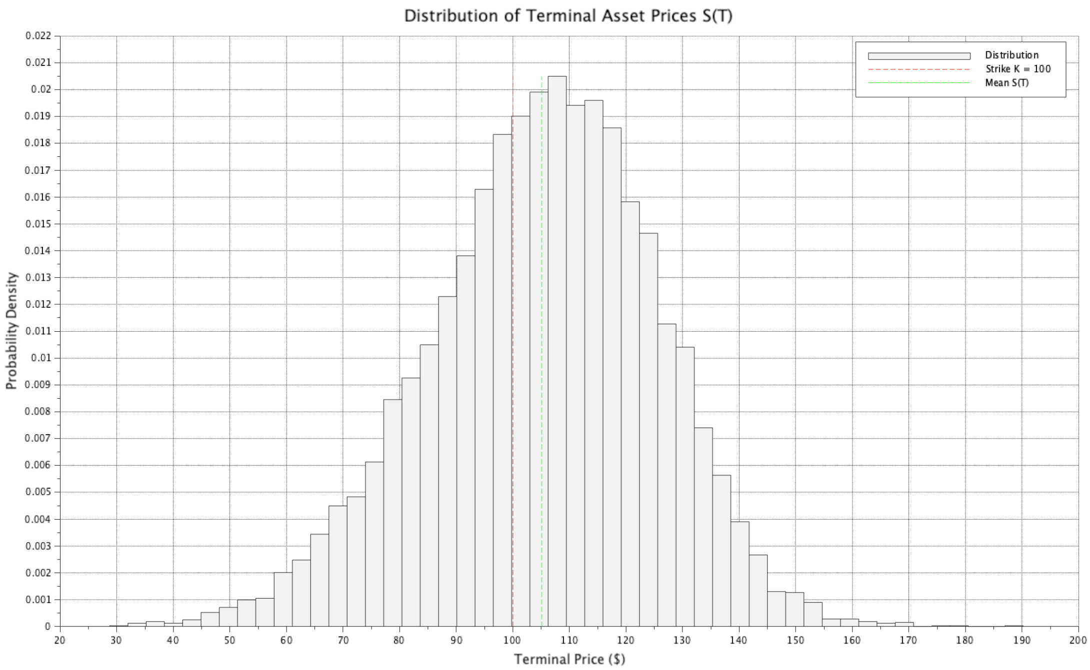

The terminal price distribution under the Heston model exhibits a pronounced negative skew and heavier left tail compared to the symmetric log-normal distribution that Black-Scholes assumes. This asymmetry is a direct result of the negative correlation rho = -0.7 between the asset price and its variance: when the price drops, volatility tends to increase, further amplifying downside moves. The mean terminal price (green line) is approximately 105, consistent with the risk-neutral drift E[S(T)] = S0*exp(rT) = 105.13.

7.5 Monte Carlo Convergence

Figure 4: Convergence of the Monte Carlo price estimate to the Heston analytical price as the number of simulation paths increases.

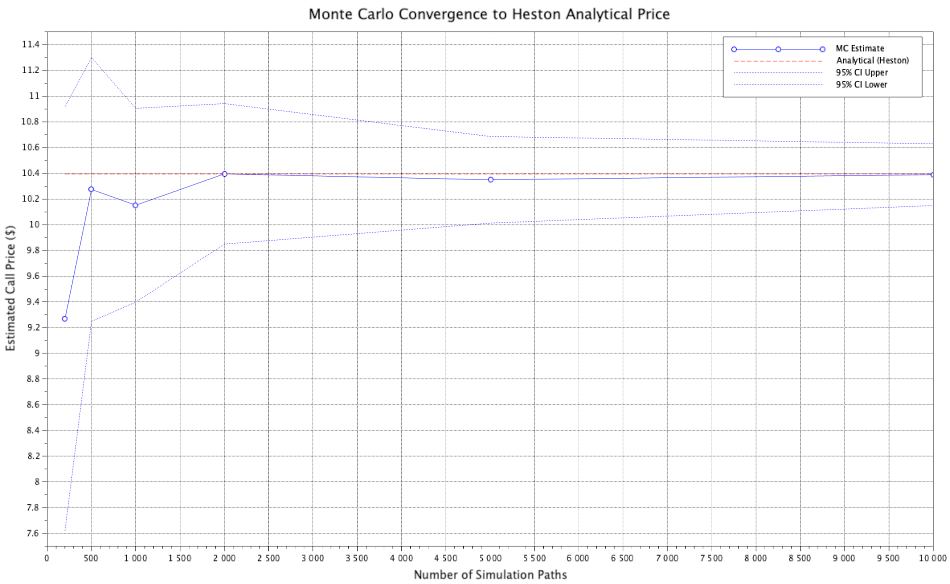

The blue curve shows the Monte Carlo price estimate as a function of the number of paths (N = 200 to 10,000). The red dashed line represents the Heston analytical price obtained via Fourier inversion. The blue dotted lines represent the 95% confidence interval, which narrows proportionally to 1/sqrt(N) as expected from the Central Limit Theorem. By N = 10,000 paths, the Monte Carlo estimate converges to within 1% of the analytical price, confirming the consistency of both pricing approaches and the correctness of the Euler-Maruyama discretization scheme.

7.6 Implied Volatility Smile

Figure 5: Implied volatility as a function of strike price K - Heston model vs Black-Scholes.

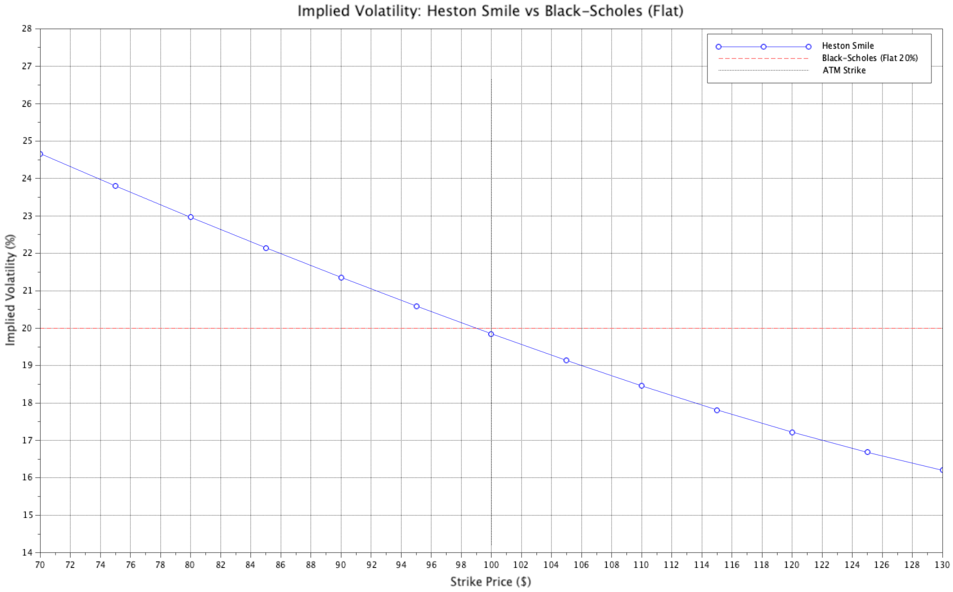

This figure presents the central result of the case study. The Heston model produces a downward-sloping implied volatility skew, ranging from approximately 25% for deep in-the-money options (K = 70) to approximately 16% for deep out-of-the-money options (K = 130). In contrast, the Black-Scholes model produces a flat line at sigma = 20% for all strikes, since it assumes constant volatility. This skew pattern - often called the "volatility smile" or "volatility skew" - is the primary empirical phenomenon that the Heston (1993) model was designed to capture. The negative correlation rho = -0.7 between the asset price and its variance creates the "leverage effect": falling prices coincide with rising volatility, making low-strike (high-moneyness) options more expensive in implied volatility terms.

7.7 Correlation (rho) Sensitivity Analysis

Figure 6: Effect of the correlation parameter rho on the implied volatility curve.

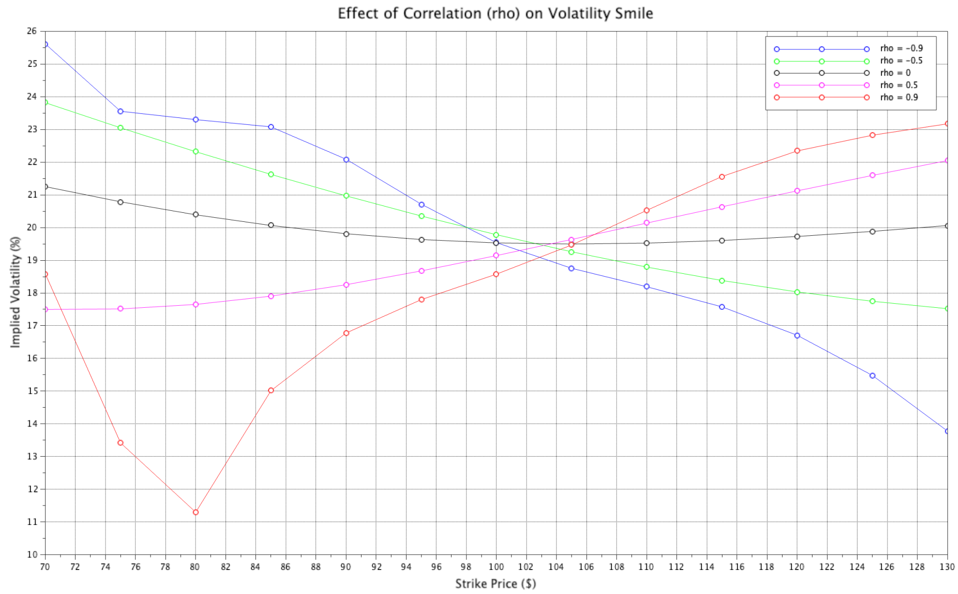

Five values of rho are tested: -0.9, -0.5, 0.0, +0.5, and +0.9, with all other parameters held constant. The results demonstrate that:

- rho = -0.9 produces the steepest downward skew (leverage effect).
- rho = -0.5 produces a moderate downward skew.
- rho = 0.0 produces a nearly symmetric smile pattern.
- rho = +0.5 produces an upward-sloping curve (inverse leverage).
- rho = +0.9 produces the steepest upward slope.

This confirms the key insight of Heston (1993): the correlation parameter rho is the primary driver of the direction and magnitude of the implied volatility skew. Negative rho (as observed empirically in equity markets) generates the well-known downward skew, while positive rho (sometimes observed in commodity markets) inverts the pattern.

7.8 3D Option Price Surface

Figure 7: Three-dimensional surface of European call option prices as a function of strike price K and time to maturity T.

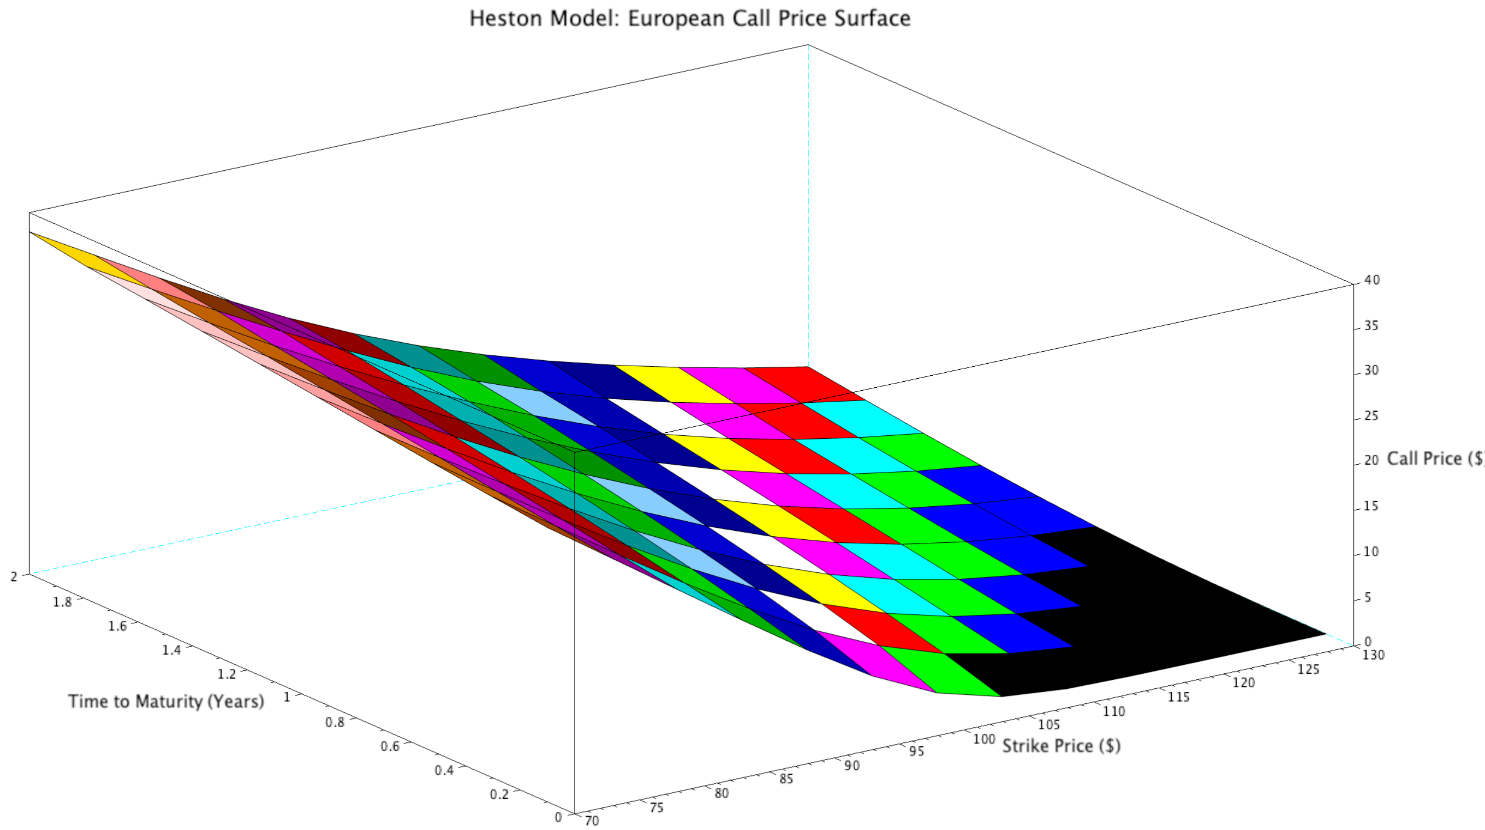

8. References

[1] Heston, S. L., "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," The Review of Financial Studies, vol. 6, no. 2, pp. 327–343, 1993.
DOI: https://doi.org/10.1093/rfs/6.2.327

[2] Black, F. and Scholes, M., "The Pricing of Options and Corporate Liabilities," Journal of Political Economy, vol. 81, no. 3, pp. 637–654, 1973.
DOI: https://doi.org/10.1086/260062

[3] Cox, J. C., Ingersoll, J. E., and Ross, S. A., "A Theory of the Term Structure of Interest Rates," Econometrica, vol. 53, no. 2, pp. 385–407, 1985.
DOI: https://doi.org/10.2307/1911242