



---

# Computational Finance Assignment 3

---

May 17, 2025

**Studenten:**

Xiaoxuan Zhang (15390802)  
Ziyi Xing(15386899)  
Shera Ding(13501380)

**Course:**

Computational Finance  
**Course catalogue number:**  
5284COFI6Y

**Repository:** [https://github.com/HarvestStars/cf\\_assignment1/tree/main/assignment3](https://github.com/HarvestStars/cf_assignment1/tree/main/assignment3)

## 1 Introduction

### 1.1 PDE Method Binary and Barrier Options

This report explores numerical and analytical methods for pricing two classes of exotic options under the Black–Scholes framework: binary (digital) options and knock-out barrier options. Both instruments pose challenges due to discontinuities in payoff or path-dependence, making them ideal testbeds for comparing pricing methodologies.

For binary call options, which pay a fixed amount if the underlying price exceeds a strike  $K$  at maturity, we implement Monte Carlo (MC) simulation, derive the closed-form pricing formula, and numerically solve the pricing PDE using implicit and Crank-Nicolson finite difference methods (FDM) after transforming it into the heat equation. Sensitivity analyses on volatility  $\sigma$  and price  $S$  reveal Delta and Vega characteristics.

For knock-out barrier options, specifically the up-and-out call, the payoff vanishes if the asset ever breaches a barrier  $B$  before maturity. We compare closed-form pricing under continuous monitoring to MC simulation under discrete monitoring. To reconcile the two, we apply the Broadie–Glasserman barrier adjustment technique, correcting for the missed overshoots in discrete sampling. We perform both parameter sensitivity studies and a convergence analysis of the MC method to validate the theoretical  $O(1/\sqrt{M})$  error rate.

### 1.2 Calibration SP500 Implied Volatility using the Heston Model

Stochastic volatility models provide a more flexible framework than Black–Scholes for capturing empirical features of option markets, such as the implied volatility smile and skew. Among these models, the Heston model is widely used in both academia and industry due to its tractable semi-closed pricing formula and parsimonious parameterization.

The rest of report explores the theoretical structure and practical implementation of the Heston model through three main stages:

- We first derive the Heston SDEs in logarithmic coordinates using Itô’s lemma, which leads to a transformed system for  $(X_t, Y_t) := (\log S_t, \log V_t)$ , improving numerical robustness and interpretability.
- Next, we implement a *vanilla option pricing engine* on a strike–maturity grid using both the semi-closed Heston formula and a Monte Carlo approach with Quadratic–Exponential (QE) discretization. From the generated call price surfaces, we invert Black–Scholes to compute the implied volatility surfaces, thereby linking Heston model outputs with market-observable metrics.
- Finally, we conduct *model calibration* to SP500 market implied volatility data. Using least-squares minimization on implied volatilities, we estimate optimal Heston parameters on two separate trading dates. Calibration

diagnostics are performed via residual plots, error metrics, and parameter stability analysis. We also propose a theoretically grounded extension for improved fit.

This integrated approach bridges stochastic calculus, numerical implementation, and empirical calibration, showcasing both the strengths and limitations of the Heston model in practice.

## 2 Method

### 2.1 Binary options

#### 2.1.1 Derivation of the Black-Scholes PDE for Binary Call Option

Under the risk-neutral probability measure  $\mathbb{Q}$ , the asset price  $S_t$  follows the stochastic differential equation:

$$\frac{dS_t}{S_t} = r dt + \sigma dB_t$$

Let  $C_d(t, S)$  denote the price of the binary (digital) call option at time  $t$  when the asset price is  $S$ . The payoff of the binary call option at maturity  $T$  with strike price  $K$  is given by

$$C_d(T, S) = \begin{cases} 1, & \text{if } S \geq K \\ 0, & \text{if } S < K \end{cases} = \mathbf{1}_{[K, \infty)}(S)$$

To derive the PDE satisfied by  $C_d(t, S)$ , we apply Itô's lemma to the function  $C_d(t, S_t)$ :

$$dC_d = \left( \frac{\partial C_d}{\partial t} + rS_t \frac{\partial C_d}{\partial S} + \frac{1}{2} \sigma^2 S_t^2 \frac{\partial^2 C_d}{\partial S^2} \right) dt + \sigma S_t \frac{\partial C_d}{\partial S} dB_t$$

Now, construct a delta-hedged portfolio

$$\Pi_t = C_d(t, S_t) - \Delta_t S_t$$

and choose  $\Delta_t = \frac{\partial C_d}{\partial S}$  to eliminate the stochastic term involving  $dB_t$ . Then, the portfolio evolves as

$$d\Pi_t = \left( \frac{\partial C_d}{\partial t} + rS_t \frac{\partial C_d}{\partial S} + \frac{1}{2} \sigma^2 S_t^2 \frac{\partial^2 C_d}{\partial S^2} - rS_t \frac{\partial C_d}{\partial S} \right) dt$$

Under the no-arbitrage principle, the portfolio must earn the risk-free rate, so:

$$d\Pi_t = r\Pi_t dt = r(C_d - S_t \frac{\partial C_d}{\partial S}) dt$$

Equating both expressions and simplifying, we obtain the Black-Scholes PDE:

$$\boxed{\frac{\partial C_d}{\partial t} + rS \frac{\partial C_d}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C_d}{\partial S^2} = rC_d}$$

with the terminal condition:

$$\boxed{C_d(T, S) = \mathbf{1}_{[K, \infty)}(S)}$$

This PDE has the same structure as the classical Black-Scholes equation for vanilla European options, but the terminal condition is specific to the binary call payoff.

### 2.1.2 Analytical Transformation of the Black-Scholes PDE (Based on Appendix A)

In order to obtain analytical or numerically tractable solutions to the Black-Scholes PDE, we apply a sequence of variable transformations that progressively simplify the equation. The following derivation is based on the steps outlined in Appendix A of the course material.

**Step 1: Time reversal substitution** Let  $\tau := T - t$  denote the time to maturity. Define the transformed function

$$g_\tau(S) := V_t(S),$$

so that

$$\partial_t V_t(S) = -\partial_\tau g_\tau(S).$$

Substituting into the standard Black-Scholes PDE

$$\partial_t V_t + rS\partial_S V_t + \frac{1}{2}\sigma^2 S^2 \partial_{SS}^2 V_t = rV_t,$$

we obtain:

$$-\partial_\tau g_\tau + rS\partial_S g_\tau + \frac{1}{2}\sigma^2 S^2 \partial_{SS}^2 g_\tau = r g_\tau. \quad (14)$$

**Step 2: Remove the nonhomogeneous term** Define a new function by applying an exponential discounting factor:

$$f_\tau(S) := e^{r\tau} g_\tau(S).$$

Differentiating with respect to  $\tau$ :

$$\partial_\tau f_\tau = r e^{r\tau} g_\tau + e^{r\tau} \partial_\tau g_\tau = r f_\tau + e^{r\tau} \partial_\tau g_\tau.$$

Hence,

$$\partial_\tau g_\tau = e^{-r\tau} (\partial_\tau f_\tau - r f_\tau).$$

Substitute into Equation (14) and multiply through by  $e^{r\tau}$ :

$$-\partial_\tau f_\tau + rS\partial_S f_\tau + \frac{1}{2}\sigma^2 S^2 \partial_{SS}^2 f_\tau = 0. \quad (15)$$

**Step 3: Logarithmic substitution** Now define the logarithmic variable  $x := \log(S)$  and the new function

$$\psi_\tau(x) := f_\tau(S).$$

Using the chain rule, we compute:

$$\partial_S f_\tau = \frac{1}{S} \partial_x \psi_\tau, \quad \partial_{SS}^2 f_\tau = \frac{1}{S^2} (\partial_{xx}^2 \psi_\tau - \partial_x \psi_\tau).$$

Substituting into Equation (15) and simplifying, we obtain:

$$-\partial_\tau \psi_\tau + \left(r - \frac{1}{2}\sigma^2\right) \partial_x \psi_\tau + \frac{1}{2}\sigma^2 \partial_{xx}^2 \psi_\tau = 0. \quad (16)$$

**Step 4: Reduction to the heat equation** We start from the transformed Black-Scholes PDE in logarithmic space, as given by Equation (16), and apply an exponential transformation:

$$\psi(\tau, x) = \phi(\tau, x) \cdot e^{\alpha x + \beta \tau}$$

- **Time derivative:**

$$\partial_\tau \psi = \partial_\tau (\phi \cdot e^{\alpha x + \beta \tau}) = e^{\alpha x + \beta \tau} (\partial_\tau \phi + \beta \phi)$$

- **First spatial derivative:**

$$\partial_x \psi = \partial_x (\phi \cdot e^{\alpha x + \beta \tau}) = e^{\alpha x + \beta \tau} (\partial_x \phi + \alpha \phi)$$

- **Second spatial derivative:**

$$\partial_{xx}^2 \psi = \partial_{xx}^2 (\phi \cdot e^{\alpha x + \beta \tau}) = e^{\alpha x + \beta \tau} (\partial_{xx}^2 \phi + 2\alpha \partial_x \phi + \alpha^2 \phi)$$

Substituting into the PDE and factoring out the exponential term, dividing by the common exponential factor and grouping terms we get:

$$\partial_\tau \phi = \left( r - \frac{\sigma^2}{2} + \sigma^2 \alpha \right) \partial_x \phi + \frac{\sigma^2}{2} \partial_{xx}^2 \phi + \left[ \beta - \left( r - \frac{\sigma^2}{2} \right) \alpha - \frac{\sigma^2}{2} \alpha^2 \right] \phi \quad (1)$$

In order to reduce the PDE to the standard heat equation form:

$$\partial_\tau \phi = \frac{\sigma^2}{2} \partial_{xx}^2 \phi$$

we must require that the coefficients of the  $\partial_x \phi$  and  $\phi$  terms vanish:

$$\begin{cases} r - \frac{\sigma^2}{2} + \sigma^2 \alpha = 0 \\ \beta - \left( r - \frac{\sigma^2}{2} \right) \alpha - \frac{\sigma^2}{2} \alpha^2 = 0 \end{cases}$$

From the first condition:

$$\alpha = \frac{1}{2} - \frac{r}{\sigma^2}$$

Substitute into the second:

$$\beta = \left( r - \frac{\sigma^2}{2} \right) \alpha + \frac{1}{2} \alpha^2 \sigma^2.$$

This completes the transformation to the standard heat equation with known coefficients. Finally, arriving at the standard heat equation:

$$\partial_\tau \phi_\tau(x) = \frac{\sigma^2}{2} \partial_{xx}^2 \phi_\tau(x). \quad (17)$$

This transformation enables us to apply analytical methods such as convolution with the Green's function to obtain closed-form solutions.

### 2.1.3 Analytical Solution of the Binary Call Option PDE

We derive the closed-form solution to the binary (digital) call option price  $C_d(t, x)$  by transforming the Black-Scholes PDE obtained in 2.1.1. The derivation relies on the change-of-variable techniques discussed in Appendix A, particularly the steps leading from the Black-Scholes PDE (Equation 13 in the appendix) to the heat equation (Equation 17), details of derivation seen in 2.1.2.

**Step 1: Problem setup** We begin with the Black-Scholes PDE satisfied by the pricing function  $C_d(t, x)$ :

$$\frac{\partial C_d}{\partial t} + rx \frac{\partial C_d}{\partial x} + \frac{1}{2} \sigma^2 x^2 \frac{\partial^2 C_d}{\partial x^2} = rC_d,$$

with the terminal condition

$$C_d(T, x) = \mathbf{1}_{[K, \infty)}(x).$$

**Step 2: Change of variables to transform the PDE** We perform the same transformation as in Appendix A to convert the Black-Scholes PDE into the heat equation:

- Define the time-to-maturity variable  $\tau := T - t$ .
- Introduce a discounting factor:  $C_d(t, x) = e^{-r\tau} f(\tau, x)$ .
- Change variables by setting  $x = e^y$ , and define  $\psi(\tau, y) := f(\tau, e^y)$ .
- Introduce a new function  $\phi(\tau, y)$  via  $\psi(\tau, y) = \phi(\tau, y)e^{\alpha y + \beta\tau}$  with appropriate constants  $\alpha$  and  $\beta$ .

With suitable choices  $\alpha = -\frac{r - \frac{1}{2}\sigma^2}{\sigma^2}$  and  $\beta = -\frac{1}{2}\alpha^2\sigma^2$ , this sequence of transformations yields the standard heat equation:

$$\frac{\partial \phi}{\partial \tau} = \frac{\sigma^2}{2} \frac{\partial^2 \phi}{\partial y^2}.$$

**Step 3: Solve the heat equation via convolution with the Green's function** The initial condition in  $\phi$ -space becomes

$$\phi(0, y) = \begin{cases} e^{-\alpha y}, & y \geq \log K, \\ 0, & y < \log K. \end{cases}$$

Solving the heat equation using the method of convolution with the Green's function of the heat equation yields:

$$\phi(\tau, y) = e^{-\alpha y + \frac{1}{2}\alpha^2\sigma^2\tau} \Phi(d_-(\tau)),$$

where  $\Phi$  is the standard normal cumulative distribution function and

$$d_-(\tau) := \frac{(r - \frac{1}{2}\sigma^2)\tau + \log(x/K)}{\sigma\sqrt{\tau}}.$$

**Step 4: Recover the original pricing function** Substituting back all changes of variables, we recover the binary call option price:

$$C_d(t, x) = e^{-r(T-t)} \Phi\left(\frac{(r - \frac{1}{2}\sigma^2)(T-t) + \log(x/K)}{\sigma\sqrt{T-t}}\right) = e^{-r(T-t)} \Phi(d_-(T-t)).$$

This is the closed-form analytical solution to the pricing problem for the binary call option under the Black-Scholes model.

### 2.1.4 Monte Carlo, Analytical and Numerical Approaches to Binary Option Pricing

To investigate the pricing and sensitivity of binary (digital) options, we employed three numerical and analytical approaches:

- **Monte Carlo simulation:** We simulated asset price paths under the Black-Scholes risk-neutral measure using geometric Brownian motion. The binary payoff was computed at maturity for each path, and the expected discounted payoff yielded the simulated option price.
- **Closed-form solution:** For benchmarking, we implemented the known analytical pricing formula for binary call options under the Black-Scholes framework. This solution served as the reference surface for evaluating the accuracy of numerical methods.
- **Finite Difference Methods (FDM):** Two PDE-based approaches were used to solve the Black-Scholes PDE transformed into the heat equation via logarithmic and exponential substitutions:
  - The fully implicit (backward Euler) scheme;
  - The Crank-Nicolson scheme (a weighted average of explicit and implicit schemes).

These schemes were applied to the transformed heat equation using backward time-stepping and spatial discretization. The resulting solutions were then inverted to recover the option pricing surfaces  $c(S, t)$ .

After obtaining the pricing surfaces from each method, we compared their accuracy against the closed-form benchmark using 3D absolute error surfaces and 2D heatmaps. Error metrics such as mean absolute error (MAE), mean squared error (MSE), and maximum absolute error were tabulated to quantify differences.

Beyond pricing accuracy, we conducted a sensitivity analysis on the PDE-based methods:

- **Vega analysis:** The effect of volatility changes on pricing was studied by solving the PDEs under varying values of  $\sigma$ . Numerical differences were used to estimate Vega curves.
- **Delta analysis:** The sensitivity of the option price with respect to the underlying asset price  $S$  was computed numerically using central finite differences on the FDM pricing grids, yielding  $\Delta(S)$  curves for fixed time slices.

These combined steps allowed us to evaluate not only the correctness of each numerical method, but also their stability and robustness under changing model parameters.

## 2.2 Knock-out Barrier options

### 2.2.1 Analytical Solution Derivation for Knock-Out Barrier Option

We now derive the closed-form pricing formula for an up-and-out barrier call option using the general path-dependent integral representation introduced in Lemma A.1.

#### Step 1: Define the payoff

The payoff of a knock-out call option is path-dependent, and can be written as

$$C = (S_T - K)^+ \cdot 1_{\{\max_{t \in [0, T]} S_t < B\}},$$

which means the option only pays off if the final price  $S_T > K$  and the barrier  $B$  has not been breached at any time before  $T$ .

Equivalently, we define the payoff function  $\phi$  in the bivariate form:

$$\phi(M_T^0, S_T) = (S_T - K)^+ \cdot 1_{\{M_T^0 < B\}}.$$

#### Step 2: Variable transformation

Following the setup in Lemma A.1, we map the maximum and final value into log space via

$$S_T = S_0 e^{\sigma x}, \quad M_T^0 = S_0 e^{\sigma y},$$

so that the payoff function becomes:

$$\phi(x, y) = (S_0 e^{\sigma x} - K)^+ \cdot 1_{\{y < \frac{1}{\sigma} \log\left(\frac{B}{S_0}\right)\}}.$$

#### Step 3: Plug into the general pricing formula

We now substitute this  $\phi(x, y)$  into the general pricing identity from Lemma A.1:

$$\mathbb{E}[C] = e^{-rT} \cdot \left[ \int_0^\infty \int_y^\infty \phi(S_0 e^{\sigma x}, S_0 e^{\sigma y}) \cdot (2x - y) \cdot f(x, y) dx dy + \dots \right],$$

where  $f(x, y)$  denotes the joint density of the drifted Brownian motion and its running maximum.

Because of the barrier condition, the upper limit of  $y$  is truncated to  $\log(B/S_0)/\sigma$ .

#### Step 4: Derive the closed-form result

By evaluating the above integrals, incorporating the Gaussian kernel and standard normal CDFs, we obtain the closed-form result for the up-and-out call:

$$\begin{aligned}
 C(t, T, S_t, M_t^0) = & S_t \cdot 1_{\{M_t^0 < B\}} \left\{ \Phi \left( \delta_+^{T-t} \left( \frac{S_t}{K} \right) \right) - \Phi \left( \delta_+^{T-t} \left( \frac{S_t}{B} \right) \right) \right\} \\
 & - \left( \frac{B}{S_t} \right)^{1+2r/\sigma^2} S_t \cdot 1_{\{M_t^0 < B\}} \left\{ \Phi \left( \delta_+^{T-t} \left( \frac{B^2}{K S_t} \right) \right) - \Phi \left( \delta_+^{T-t} \left( \frac{B}{S_t} \right) \right) \right\} \\
 & - e^{-r(T-t)} K \cdot 1_{\{M_t^0 < B\}} \left\{ \Phi \left( \delta_-^{T-t} \left( \frac{S_t}{K} \right) \right) - \Phi \left( \delta_-^{T-t} \left( \frac{S_t}{B} \right) \right) \right\} \\
 & + \left( \frac{S_t}{B} \right)^{1-2r/\sigma^2} e^{-r(T-t)} K \cdot 1_{\{M_t^0 < B\}} \left\{ \Phi \left( \delta_-^{T-t} \left( \frac{B^2}{K S_t} \right) \right) - \Phi \left( \delta_-^{T-t} \left( \frac{B}{S_t} \right) \right) \right\},
 \end{aligned}$$

with

$$\delta_{\pm}^{\tau}(z) = \frac{1}{\sigma\sqrt{\tau}} \left( \log z + \left( r \pm \frac{\sigma^2}{2} \right) \tau \right).$$

This result demonstrates that the knock-out barrier option can be priced explicitly through the probabilistic maximum representation and careful variable transformation of the standard GBM process.

### 2.2.2 Monte Carlo and Analytical Approaches to Barrier Option Pricing

**Adjusted Barrier Approach** To account for the fact that the MC simulation monitors the barrier at discrete time intervals (rather than continuously), we adopt the adjusted barrier technique proposed by Broadie et al., where the original barrier level  $H$  is shifted upward. The closed-form solution is evaluated using the following adjusted barrier:

$$C_m(H) \approx C \left( H \cdot e^{+\beta_1 \sigma \sqrt{T/m}} \right), \quad \text{with } \beta_1 \approx 0.5826$$

This correction compensates for the tendency of discretely monitored MC paths to underestimate the frequency of barrier breaches, especially between monitoring points.

**Sensitivity Analysis.** We first conduct a sensitivity analysis by varying key model and option parameters: the strike price  $K$ , barrier level  $B$ , time to maturity  $T$ , volatility  $\sigma$ , and interest rate  $r$ . For each parameter, we compute and compare the option prices obtained from:

- Standard Monte Carlo simulations using  $M$  discrete time steps and  $N$  sample paths;
- The adjusted closed-form formula with barrier  $H \cdot e^{+\beta_1 \sigma \sqrt{T/M}}$ .

This allows us to assess how robust the MC method is across different market and contract conditions, and how closely it tracks the analytical benchmark under parameter shifts.

**Convergence Analysis.** We then evaluate the convergence rate of the MC method with respect to the number of time steps  $M$ . According to theory, the adjusted MC estimator satisfies the following asymptotic behavior:

$$|C_m(H) - C(H \cdot e^{+\beta_1 \sigma \sqrt{T/m}})| = O \left( \frac{1}{\sqrt{m}} \right)$$

Taking logarithms on both sides yields:

$$\log_{10}(\text{error}) = -\frac{1}{2} \log_{10}(m) + \text{const.}$$

This implies that plotting the log absolute error against  $\log_{10}(m)$  should produce a straight line with slope approximately  $-0.5$ . To verify this, we compute the absolute pricing error across a range of  $m$  values, averaging results over multiple simulations to reduce variance. We then fit a linear regression to the  $\log_{10}(\text{error})$  vs.  $\log_{10}(m)$  data to confirm whether the observed slope is consistent with the theoretical convergence rate.

### 2.2.3 Implicit FDM vs Closed-Form Barrier Option Pricing and Sensitivity Analysis

We employ an implicit finite difference scheme (FDM) to numerically solve the pricing problem of an up-and-out barrier call option. To verify the correctness and consistency of the implementation, we first compare the option price surface computed via FDM with the analytical closed-form solution derived from barrier option theory. This comparison serves as a benchmark to validate the numerical approach.

The computational framework follows a log-space transformation of the Black-Scholes PDE into a standard heat equation form, allowing stable and efficient implicit time-stepping. Boundary and initial conditions are appropriately set according to the transformed payoff structure, incorporating the barrier constraint directly via Dirichlet conditions in the finite difference grid.

The methodology proceeds in three main stages:

- **Price Surface Comparison:** We generate the full option price surface  $V(S, t)$  using both the closed-form formula and the implicit FDM scheme. These surfaces are visualized over a range of asset prices  $S$  and times  $t$ , and visually compared to ensure that the FDM result converges to the analytical benchmark within a reasonable tolerance.
- **Sensitivity Analysis:** Next, we investigate the sensitivity of the option price with respect to three key parameters: the strike price  $K$ , the barrier level  $B$ , and the volatility  $\sigma$ . For each parameter, we vary its value across a representative range while keeping others fixed. The resulting option prices computed by both FDM and the closed-form method are plotted for comparison. This highlights how each parameter influences pricing behavior and whether the numerical scheme correctly reflects the underlying financial sensitivity.
- **Delta Curve Extraction:** Finally, we compute the option delta  $\Delta = \frac{\partial V}{\partial S}$  as a function of the underlying asset price  $S$ . For the FDM result,  $\Delta$  is obtained via central finite differences on the computed price grid at  $t = 0$ . For the closed-form benchmark, we apply numerical differentiation to the analytical price formula. The resulting delta curves are compared, emphasizing the barrier effect on hedging sensitivity near the barrier level.

This numerical methodology enables a robust and interpretable study of the barrier option structure, while also confirming that the implicit discretization aligns well with known analytical behavior.

## 2.3 Calibration SP500 Implied Volatility using the Heston Model

### 2.3.1 Dynamics in logarithmic coordinates

(a) **SDEs for  $X_t = \ln S_t$  and  $Y_t = \ln V_t$ .** Under the risk-neutral Heston model

$$\begin{aligned} dS_t &= r S_t dt + S_t \sqrt{V_t} dB_t, \\ dV_t &= \kappa(\theta - V_t) dt + \sigma \sqrt{V_t} dW_t, \\ d\langle B, W \rangle_t &= \rho dt, \end{aligned}$$

apply Itô's formula to  $X_t = \ln S_t$  and  $Y_t = \ln V_t$ :

$$dX_t = \frac{1}{S_t} dS_t - \frac{1}{2} \frac{1}{S_t^2} (dS_t)^2 = \left( r - \frac{1}{2} V_t \right) dt + \sqrt{V_t} dB_t$$

$$dY_t = \frac{1}{V_t} dV_t - \frac{1}{2} \frac{1}{V_t^2} (dV_t)^2 = \left[ \kappa(\theta e^{-Y_t} - 1) - \frac{1}{2} \sigma^2 e^{-Y_t} \right] dt + \sigma e^{-Y_t/2} dW_t$$

Equivalently, in terms of  $V_t$ ,

$$dY_t = \left[ \frac{\kappa\theta - \frac{1}{2}\sigma^2}{V_t} - \kappa \right] dt + \frac{\sigma}{\sqrt{V_t}} dW_t.$$

(b) **Why the log-variance coordinate is convenient.**

- **Positivity handled automatically:**  $V_t \geq 0$  corresponds to  $Y_t \in \mathbb{R}$ , eliminating boundary constraints in simulations and optimisation.

- **Numerical stability:** The diffusion coefficient  $\sigma e^{-Y_t/2}$  varies smoothly, reducing discretisation error relative to  $\sqrt{V_t}$ .
- **Analytic compatibility:** Many semi-closed formulas for Heston (e.g. characteristic functions, FFT pricing) are expressed in  $\ln S$  and  $\ln V$ , so the log coordinate simplifies subsequent transforms.

### 2.3.2 Vanilla option pricing experiments

We begin by fixing a Heston parameter vector  $\Theta = (\kappa, \theta, \rho, \sigma, v_0)$  and a continuously-compounded risk-free rate  $r$ . A rectangular grid  $\mathcal{G} = \{(K_i, T_j)\}_{i=1, \dots, n_K}^{j=1, \dots, n_T}$  of strikes and maturities is then constructed. For every node we compute

$$C_{ij}^{\text{model}} = e^{-rT_j} \mathbb{E}[(S_{T_j} - K_i)^+],$$

using **two independent engines:**

**Monte-Carlo engine (Quadratic-Exponential scheme)** The variance process is advanced over the step  $\Delta t = T_j/N$  by the Quadratic-Exponential (QE) method of Andersen (Andersen, 2006).

**Semi-closed Heston formula** The analytic benchmark employs the well-known representation

$$C^{\text{CF}}(S_0, K, T) = S_0 P_1 - K e^{-rT} P_2,$$

$$P_1 = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re\left(e^{-iu \ln K} \frac{\phi(u-i; T)}{iu F}\right) du, \quad F = S_0 e^{rT}.$$

$$P_2 = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re\left(e^{-iu \ln K} \frac{\phi(u; T)}{iu}\right) du,$$

where the characteristic function of  $\ln S_T$  is

$$\phi(u; T) = \exp\left\{iu(\ln S_0 + rT) + \frac{\kappa\theta}{\sigma^2}\left[(\xi - d)T - 2\ln\left(\frac{1-g_2 e^{-dT}}{1-g_2}\right)\right] + \frac{v_0}{\sigma^2}(\xi - d) \frac{1 - e^{-dT}}{1 - g_2 e^{-dT}}\right\},$$

with  $\xi = \kappa - i\rho\sigma u$ ,  $d = \sqrt{\xi^2 + \sigma^2(u^2 + iu)}$ ,  $g_2 = (\xi - d)/(\xi + d)$ . We compute the oscillatory integrals using the adaptive Gauss-Lobatto routine of Cui et al. (2017). The two call-price surfaces  $\{C_{ij}^{\text{MC}}\}$  and  $\{C_{ij}^{\text{CF}}\}$  are plotted side-by-side for visual comparison.

**Extraction of implied volatility** For any obtained price we determine the Black-Scholes implied volatility  $\sigma_{ij}^{\text{imp}}$  as the unique root of

$$C^{\text{BS}}(S_0, K_i, T_j, r, \sigma) - C_{ij}^{\text{model}} = 0,$$

where  $C^{\text{BS}} = S_0 N(d_1) - K_i e^{-rT_j} N(d_2)$  and  $d_{1,2} = \frac{\ln(S_0/K_i) + (r \pm \frac{1}{2}\sigma^2)T_j}{\sigma\sqrt{T_j}}$ . Monotonicity of  $C^{\text{BS}}$  in  $\sigma$  allows us to use Brent's method, resulting in two implied-volatility surfaces  $\{\sigma_{ij}^{\text{MC}}\}$  and  $\{\sigma_{ij}^{\text{CF}}\}$ .

### 2.3.3 Calibration to S&P 500 implied volatilities

The final task is to adjust the five Heston parameters  $\Theta = (\kappa, \theta, \rho, \sigma, v_0)$  so that the model reproduces the *market* implied-volatility (IV) surface of S&P 500 options as closely as possible. The workflow is organised as follows.

- (1) **Generate model IVs.** For any candidate parameter set  $\Theta$  we evaluate the semi-closed Heston price  $C_{ij}^{\text{CF}}(\Theta)$  on the grid, followed by a one-dimensional root-finder to obtain the corresponding model volatilities  $\sigma_{ij}^{\text{Heston}}(\Theta)$ .
- (2) **Define the calibration objective.** We minimise a weighted root-mean-square error (RMSE) on implied vols:

$$J(\Theta) = \sqrt{\frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N w_{ij} \left(\sigma_{ij}^{\text{Heston}}(\Theta) - \sigma_{ij}^{\text{mkt}}\right)^2}, \quad w_{ij} = 1 \text{ or } w_{ij} = \left(\sigma_{ij}^{\text{mkt}}\right)^{-2}.$$

The objective is evaluated under positivity constraints  $\kappa, \theta, \sigma, v_0 > 0$  and  $\rho \in (-1, 1)$ .

- (3) **Optimisation strategy.** A two-stage procedure is adopted: (i) a global search via Differential Evolution for a robust starting point; (ii) a local, gradient-based L-BFGS-B refinement with analytical gradients supplied by reverse-mode AD.
- (4) **Multiple-date calibration.** The routine is run independently on *at least two* trade dates (e.g. 2023-10-31 and 2023-11-02) to assess parameter stability. Optionally, a joint-fit variant with a piece-wise-constant  $\Theta(t)$  can be explored.
- (5) **Diagnostics.**
  - *Visual fit:* heat-maps / wireframes of model vs. market IV surfaces.
  - *Error metrics:* RMSE, mean absolute error (MAE), and maximum absolute error across the grid.
  - *Stability analysis:* side-by-side comparison of calibrated  $\Theta$  across dates, highlighting where the fit deteriorates (short maturities, deep OTM, etc.).
  - *Model improvements:* discuss possible extensions such as time-dependent parameters, stochastic-jump volatility, or a SABR-type volatility-of-vol.

This structured pipeline leverages the pricing and IV-extraction modules developed in the previous sections and provides a quantitative assessment of the Heston model's ability to replicate the observed S&P 500 volatility smile.

## 3 Results

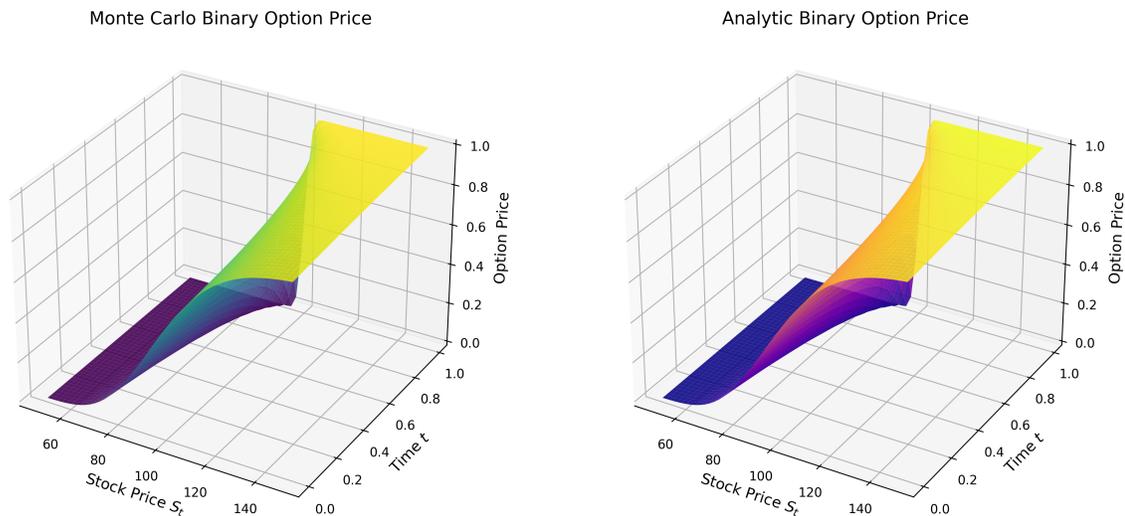
### 3.1 Binary options

#### 3.1.1 Pricing Results via Monte Carlo and Analytical Comparison

We validate the Monte Carlo (MC) method for pricing binary call options by comparing it against the analytical Black-Scholes solution under the following parameters:

- $S_t \in [50, 150]$ ,  $t \in [0.001, 0.99]$
- $K = 100$ ,  $T = 1.0$ ,  $r = 0.05$ ,  $\sigma = 0.2$

Figure 1 shows the option price surfaces from the MC simulation and the closed-form solution. Both exhibit the sharp transition in price near the strike as maturity approaches, characteristic of digital options.

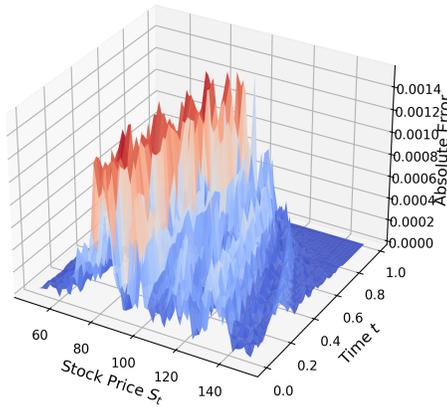


**Figure 1:** Binary option pricing surfaces. Left: Monte Carlo. Right: analytical solution.

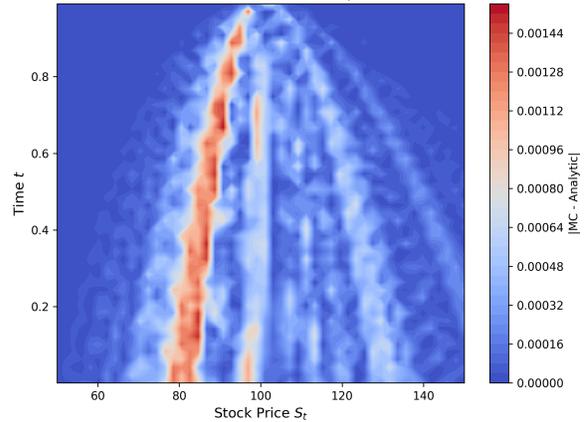
To assess accuracy, we compute the absolute error between the MC and analytical surfaces. As shown in Figure 2, errors are localized near  $S = K$  and  $t \approx T$ , where the payoff discontinuity dominates.

Quantitative error statistics are summarized below:

Absolute Error Surface (|MC - Analytic|)



Absolute Error Heatmap



**Figure 2:** Left: Absolute error surface  $|C_{MC} - C_{analytic}|$ . Right: corresponding heatmap.

- Mean Absolute Error (MAE): 0.000231
- Mean Relative Error: 5.4793%
- Max Absolute Error: 0.001553
- Mean Squared Error (MSE): 0.000000

These results demonstrate that the Monte Carlo method produces accurate prices on average, with discrepancies concentrated near the strike where the payoff is non-smooth.

### 3.1.2 Comparison of FDM Schemes and Sensitivity Analysis

Based on the parameters shown in 3.1.1, we then compare the performance of the Implicit and Crank–Nicolson (CN) finite difference schemes in pricing binary options, and further analyze their sensitivity to model parameters such as volatility  $\sigma$  and underlying price  $S$ .

**Dirichlet Boundary and Initial Conditions in the FDM Workflow.** In our numerical implementation of the Black-Scholes PDE for binary options, we perform a standard sequence of transformations to reduce the PDE to the heat equation. Let  $\tau = T - t$ ,  $x = \log(S/K)$ , and define the transformed function  $\phi_\tau(x)$  satisfying

$$\partial_\tau \phi_\tau(x) = \frac{\sigma^2}{2} \partial_{xx}^2 \phi_\tau(x),$$

which is the canonical heat equation. The spatial domain  $x \in [x_{\min}, x_{\max}]$  corresponds to the log-transformed asset prices  $S \in [50, 150]$ , with  $K = 100$ . Hence,

$$x_{\min} = \log(50/100) = \log(0.5) \approx -0.6931, \quad x_{\max} = \log(150/100) = \log(1.5) \approx 0.4055.$$

We impose Dirichlet boundary conditions by fixing  $\phi_\tau(x_{\min}) = \phi_0(x_{\min})$  and  $\phi_\tau(x_{\max}) = \phi_0(x_{\max})$  for all  $\tau \in [0, T]$ , where  $\phi_0(x)$  is derived from the initial condition of the digital option payoff. Specifically, the original payoff at maturity is

$$V_T(S) = \begin{cases} 1 & \text{if } S \geq K, \\ 0 & \text{if } S < K, \end{cases}$$

which translates to  $\psi_0(x) = 1_{x \geq 0}$ , and hence

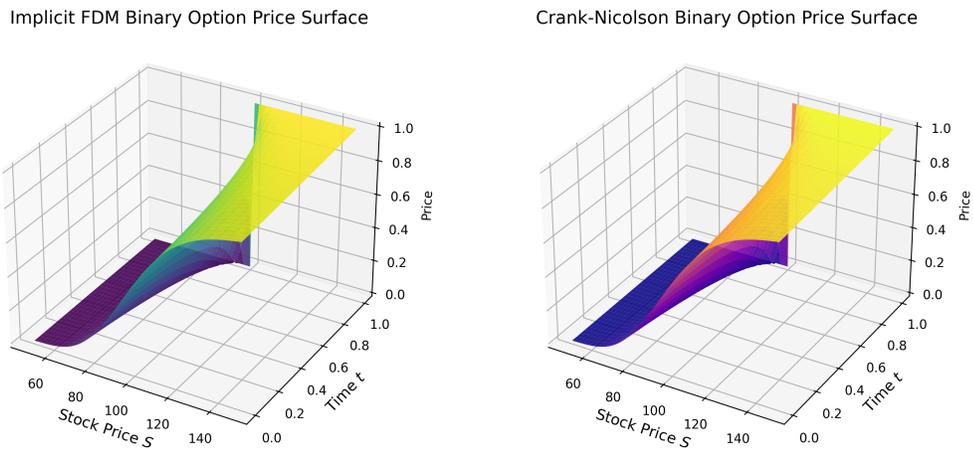
$$\phi_0(x) = \psi_0(x) \cdot e^{-\alpha x} = \begin{cases} e^{-\alpha x} & \text{if } x \geq 0, \\ 0 & \text{if } x < 0. \end{cases}$$

With  $\alpha = \frac{2r}{\sigma^2} \approx \frac{2 \cdot 0.05}{0.04} = 2.5$ , the right-end boundary value evaluates to

$$\phi_\tau(x_{\max}) = \phi_0(x_{\max}) = \exp(-2.5 \cdot \log(1.5)) = \frac{1}{1.5^{2.5}} \approx 0.3628, \quad \text{and} \quad \phi_\tau(x_{\min}) = 0.$$

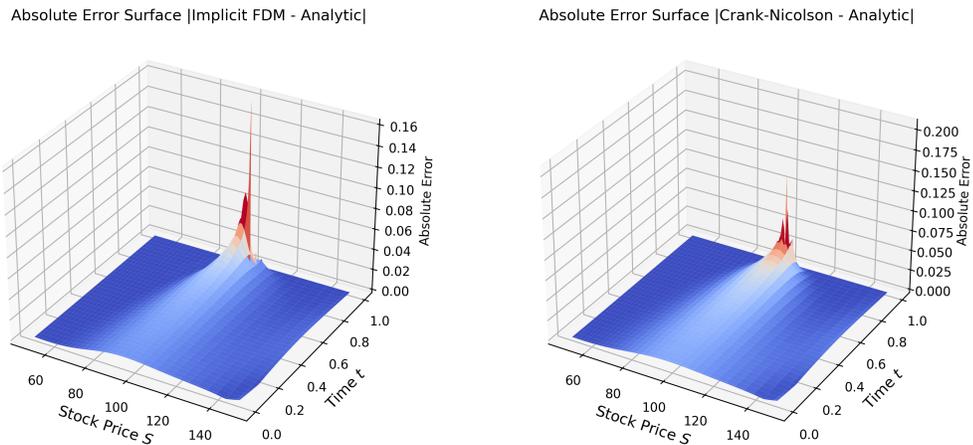
This setup ensures that the heat equation is solved on a finite domain with a well-defined initial condition and fixed Dirichlet boundaries. These are enforced in both the backward Euler and Crank-Nicolson solvers, and ensure numerical stability while respecting the theoretical solution structure of the transformed binary option problem.

**Pricing and Error Surface Comparison.** Figure 3 shows the option price surfaces generated by the two numerical methods. Both surfaces agree well with the expected binary option shape, featuring a sharp transition near  $S = K$ .



**Figure 3:** Binary option price surfaces using Implicit (left) and Crank-Nicolson (right) finite difference methods.

Figure 4 presents the corresponding absolute error surfaces against the analytical solution. Although the CN scheme produces a smoother surface overall, a sharp peak near  $(S = 100, t \rightarrow T)$  yields a larger maximum error, as confirmed below.



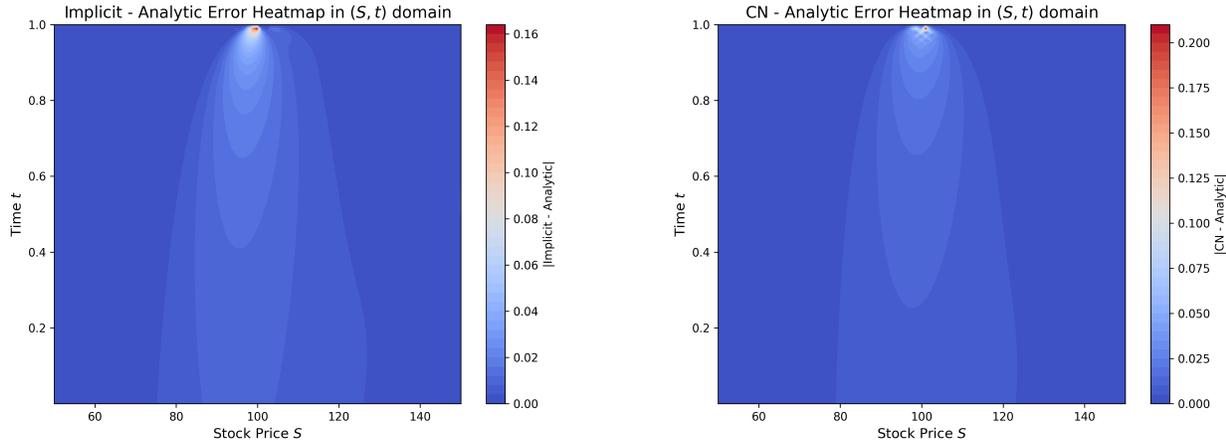
**Figure 4:** Absolute error surfaces: Implicit (left) and CN (right) vs. analytical pricing. With  $K = 100$

### Error Statistics.

- **Implicit Scheme:** MAE = 0.004128, Max Error = 0.1616, MSE = 0.000053
- **CN Scheme:** MAE = 0.004115, Max Error = 0.207975, MSE = 0.000051

These confirm the superior average performance of CN, but with potential instability near the strike as  $t \rightarrow T$ .

To better localize the error spike, we plot the heatmaps in Figure 5. The CN heatmap clearly exposes a small localized region with elevated error amplitude.



**Figure 5:** Error heatmaps (log-scaled). Left: Implicit vs. analytic. Right: CN vs. analytic. With  $K = 100$

**Sensitivity to Option Prices in parameters space.** Figure 6 shows the sensitivity of the binary option price at  $S_0 = 100$  to four key parameters: strike price  $K$ , time to maturity  $T$ , volatility  $\sigma$ , and risk-free rate  $r$ . Across all subfigures, we observe that the results from the implicit FDM and the Crank–Nicolson (CN) method nearly coincide, highlighting the numerical consistency between these two finite difference schemes.

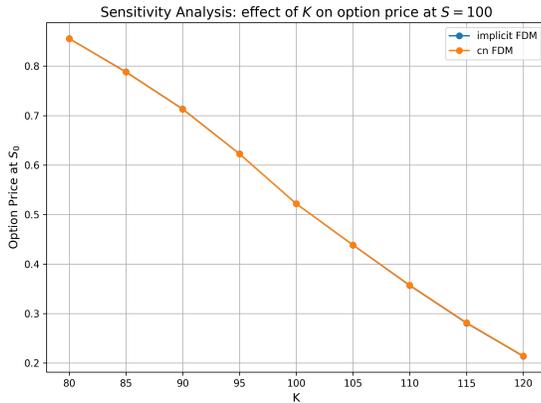
Panel (a) demonstrates that as  $K$  increases, the option price decreases almost linearly, since a higher strike reduces the probability of finishing in-the-money. In contrast, panel (b) reveals that the option price increases approximately linearly with  $r$ , reflecting the stronger discounting effect favoring binary options due to their all-or-nothing payoff. Panel (c) shows that increasing  $\sigma$  decreases the option price, which may seem counterintuitive at first. However, for digital options, higher volatility leads to a flatter price profile and reduces the steepness around the strike, thereby reducing the overall probability density concentrated near the threshold. Finally, panel (d) presents a concave response to  $T$ : longer maturities increase the chance of reaching the strike, but the marginal benefit of additional time diminishes, leading to a plateau in the option value as  $T$  grows.

**Sensitivity to Volatility (Vega).** From Figure 7, we observe that binary option prices increase with volatility  $\sigma$  in a non-linear fashion. Under both implicit and CN schemes, the slope of the price curve around  $S = K = 100$  becomes less steep as  $\sigma$  increases, reflecting the smoothing effect of volatility on the digital option's discontinuous payoff structure.

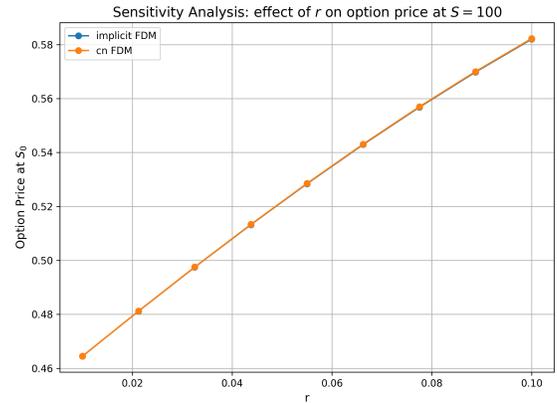
However, the CN scheme exhibits a distinct numerical instability when  $\sigma = 0.9$ , notably at  $S = K = 100$ , where the curve sharply oscillates. This artifact arises due to the higher sensitivity of CN's centered time-stepping to boundary and discretization errors in regions of steep gradients. When volatility is large, the transition band of the digital option broadens and becomes highly sensitive to grid resolution and rounding errors, which may trigger amplified instability during matrix inversion in CN updates. In contrast, the fully implicit method, being backward in time and unconditionally stable, remains robust even at extreme volatility levels.

Figure 8 illustrates the numerical Vega. As seen in the plots, for small volatility (e.g.,  $\sigma = 0.1$ ), the Vega curve is highly non-linear with strong oscillatory behavior around the strike  $S = K = 100$ , reflecting high sensitivity to volatility changes in that region. As volatility increases, the Vega curve flattens significantly, indicating that the option becomes less sensitive to changes in  $\sigma$  when volatility is already high.

Additionally, the CN method exhibits noticeable numerical artifacts near  $S = 100$ , particularly when  $\sigma$  is large, such as  $\sigma = 0.9$ . These localized spikes are not present in the implicit scheme, which suggests that the CN method, while smoother in general, can be more sensitive to grid resolution or stability issues near discontinuities in the payoff derivative.



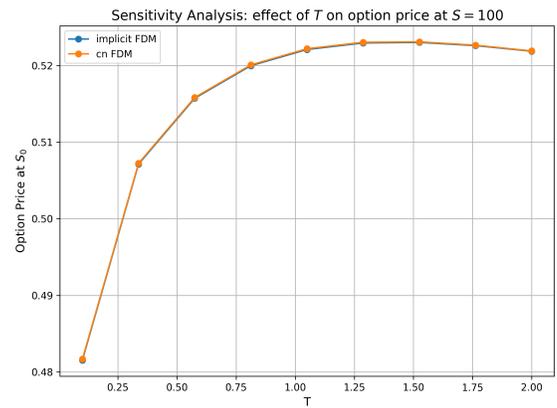
(a) Effect of strike price  $K$



(b) Effect of interest rate  $r$

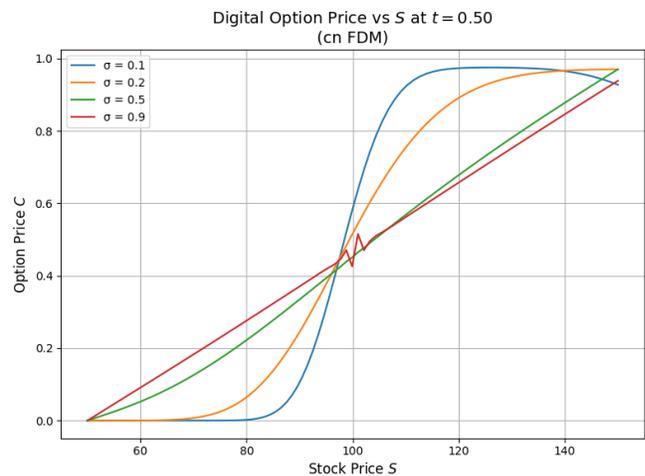
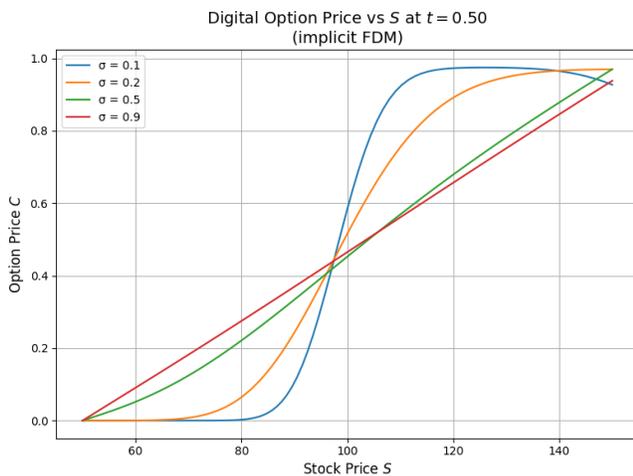


(c) Effect of volatility  $\sigma$



(d) Effect of maturity  $T$

**Figure 6:** Sensitivity analysis of binary option prices with respect to model and option parameters, evaluated at  $S_0 = 100$ . Both implicit and Crank–Nicolson FDM schemes yield nearly identical prices across all parameter regimes.



**Figure 7:** Sensitivity of binary option prices to different  $\sigma$  under implicit (left) and CN (right) FDM. With  $K = 100$

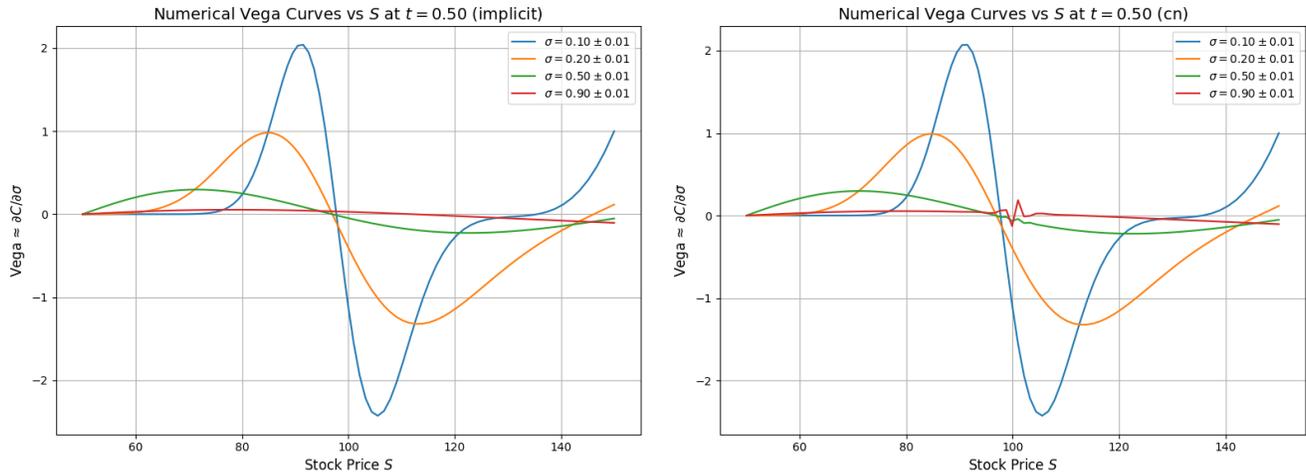


Figure 8: Numerical Vega  $\partial C/\partial\sigma$  using implicit (left) and CN (right) methods. With  $K = 100$

**Sensitivity to Underlying (Delta).** Figure 9 displays the estimated Delta  $\partial C/\partial S$  as a function of stock price  $S$  at fixed time  $t = 0.50$ , using both implicit and Crank-Nicolson (CN) schemes under varying volatilities  $\sigma \in \{0.1, 0.2, 0.5, 0.9\}$ . Several key observations can be made:

- For both schemes, Delta exhibits a peaked shape centered near the strike price  $K = 100$ , reflecting the sharp sensitivity of binary options to underlying price movements when  $S \approx K$ .
- As volatility increases, the peak of the Delta flattens. This aligns with intuition: higher volatility blurs the transition region where the option payoff changes, reducing local sensitivity.
- For the implicit scheme, the Delta remains smooth and well-behaved across all tested  $\sigma$  values.
- In contrast, the CN scheme exhibits oscillations and numerical instability when  $\sigma$  becomes large (e.g.,  $\sigma = 0.9$ ). These wiggles occur around  $S = 100$ , likely due to the discretization error magnified by the scheme's semi-implicit nature in the presence of a discontinuous payoff and high curvature in the pricing surface.

Such differences underscore the practical trade-off between accuracy and stability in finite difference schemes, particularly for options with discontinuous payoffs. While CN provides second-order accuracy in theory, implicit schemes may offer better numerical robustness under extreme parameter settings.

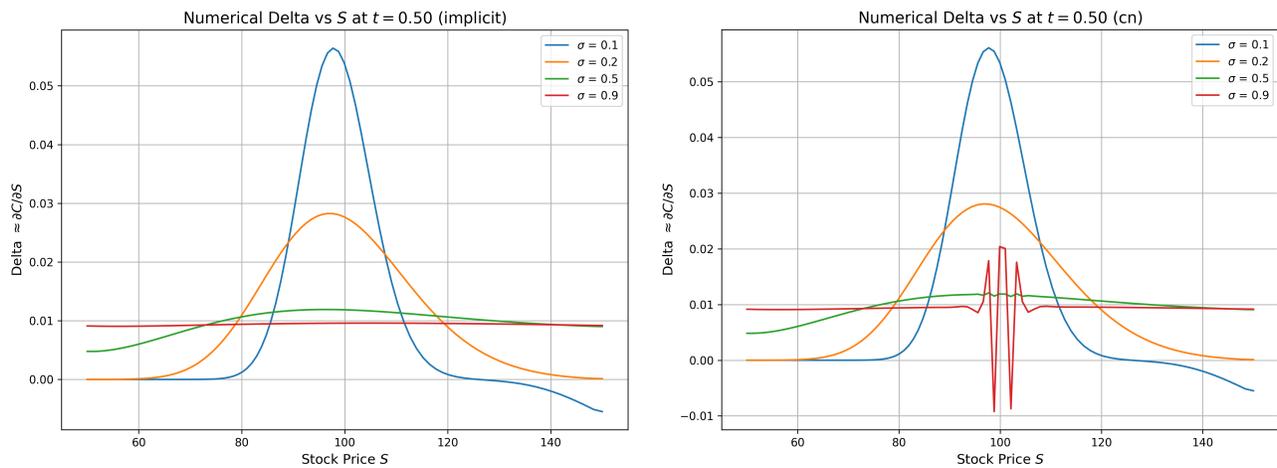


Figure 9: Estimated Delta  $\partial C/\partial S$  for varying volatility  $\sigma$  at  $t = 0.50$ , using implicit (left) and Crank-Nicolson (right) FDM.

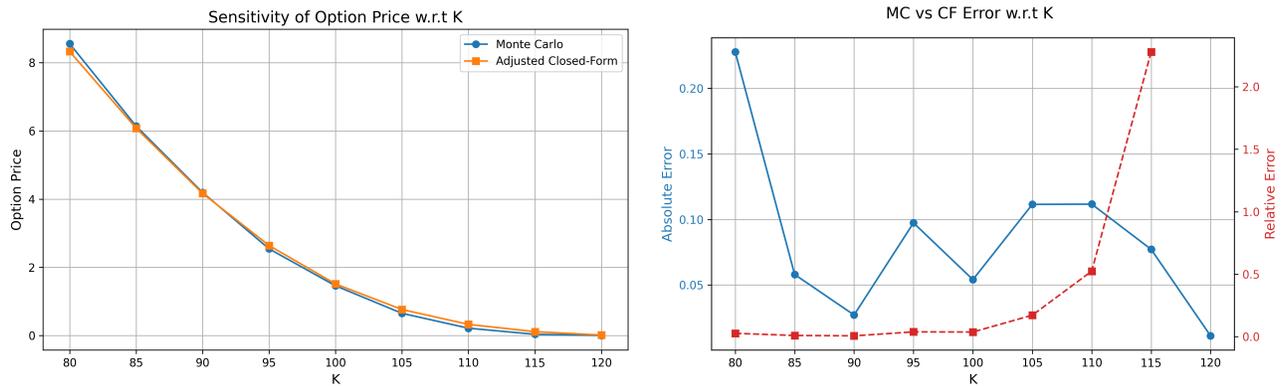
### 3.2 Knock-out Barrier options

#### 3.2.1 Pricing Sensitivity and Convergence via Monte Carlo and Analytical Comparison

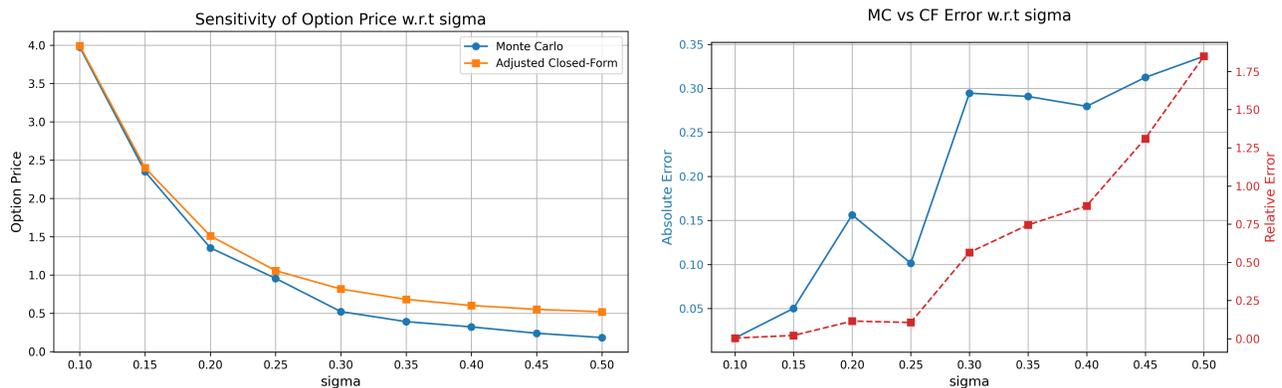
**Sensitivity Analysis.** We first evaluate the sensitivity of the option pricing results with respect to three key parameters: strike price  $K$ , volatility  $\sigma$ , and barrier level  $B$ , using the following baseline setup:

- Spot price  $S_0 = 100$ , strike  $K = 100$ , barrier  $B = 120$ , maturity  $T = 1.0$
- Volatility  $\sigma = 0.2$ , interest rate  $r = 0.05$
- Monte Carlo parameters:  $M = 100$  time steps,  $N = 100,000$  paths
- Barrier correction factor:  $\beta_1 = 0.5826$

Figures 10 to 12 display comparisons between Monte Carlo and adjusted closed-form prices across different values of  $K$ ,  $\sigma$ , and  $B$ , respectively. Each set contains two plots: one showing the option price curves, and the other showing the corresponding absolute and relative error between the two methods.



**Figure 10:** Sensitivity of option price and error comparison w.r.t. strike price  $K$ .



**Figure 11:** Sensitivity of option price and error comparison w.r.t. volatility  $\sigma$ .

Across all three cases, we observe that the Monte Carlo estimates closely follow the adjusted analytical values. The absolute error generally remains below 0.2, and the relative error increases when the option becomes deep out-of-the-money (i.e., when the payoff approaches zero). This is especially noticeable in the right-hand plots, where relative error spikes at extreme parameter values due to the denominator shrinking.

Overall, the results validate the robustness of the adjusted closed-form method in approximating the discretely monitored MC estimator across varying option features and market conditions.

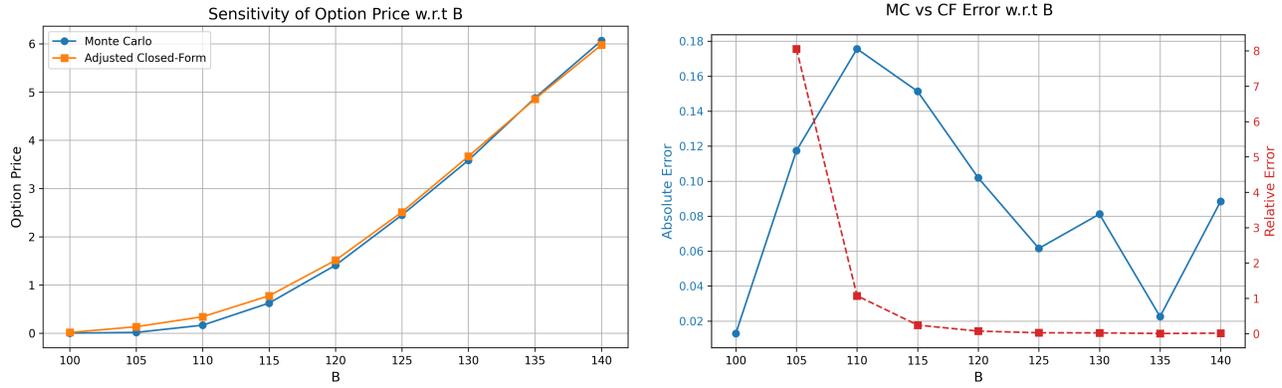


Figure 12: Sensitivity of option price and error comparison w.r.t. barrier level  $B$ .

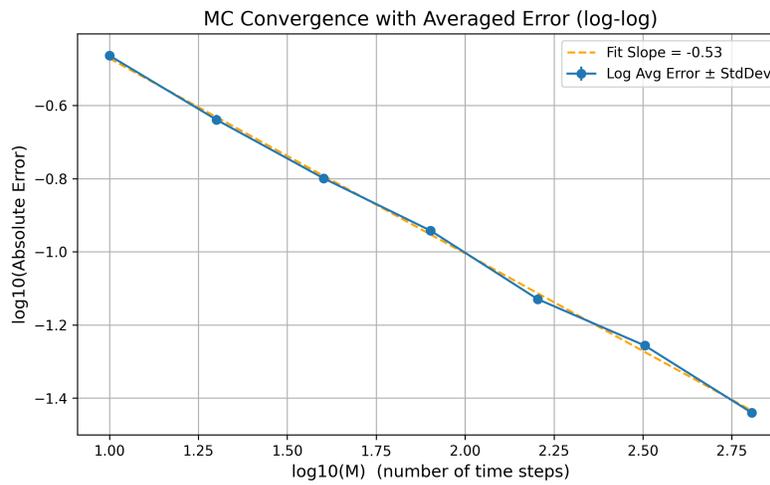


Figure 13: Log-log convergence analysis of Monte Carlo error vs. number of time steps  $M$ . Slope  $\approx -0.53$ .

**Convergence Analysis.** Next, we study the convergence of the MC estimator with increasing time resolution  $M$ , under the following fixed configuration:

- $S_0 = 100, K = 100, B = 120, T = 1.0, r = 0.05, \sigma = 0.2$
- Barrier correction:  $H_{\text{adj}} = H \cdot e^{\beta_1 \sigma \sqrt{T/M}}$
- Sample size  $N = 100,000$ , repeated 5 times for each  $M$

Figure 13 shows the log-log plot of the average absolute error versus  $\log_{10}(M)$ , along with a linear regression fit. The resulting slope is approximately  $-0.53$ , which is in close agreement with the theoretical convergence rate of  $O(1/\sqrt{m})$  discussed in the Method section.

This result provides strong empirical confirmation of the effectiveness of the adjusted barrier correction. It also demonstrates that the MC method achieves the expected convergence behavior when pricing discretely monitored knock-out barrier options.

### 3.2.2 Implicit FDM vs Closed-Form Barrier Option Pricing and Sensitivity Analysis

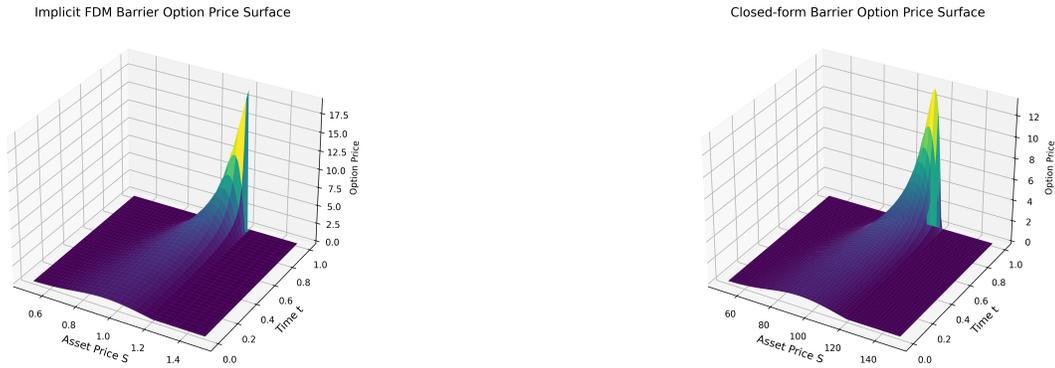
To validate the accuracy of our implicit finite difference method (FDM) for pricing up-and-out barrier call options, we first compare its pricing surface with that of the analytical closed-form solution.

For this comparison, the following parameter values are used:

- Spot price  $S_0 = 100$ , strike price  $K = 100$ , barrier level  $B = 120$

- Maturity  $T = 1.0$ , interest rate  $r = 0.05$ , volatility  $\sigma = 0.25$
- Asset price range  $S \in [50, 150]$ , grid resolution  $N = 100$  (space),  $M = 100$  (time)

Figure 14 shows the resulting 3D surfaces of option prices  $C(S, t)$  as functions of asset price and time.



(a) Implicit FDM surface ( $S_0 = 100, K = 100, B = 120, \sigma = 0.25$ )

(b) Closed-form solution surface (same parameters)

**Figure 14:** Comparison of pricing surfaces  $C(S, t)$  between implicit FDM and closed-form under identical parameter setup.

We observe strong visual agreement between the two surfaces, particularly in the interior region before the barrier  $S < B$ . This validates the correctness of our FDM discretization and implementation for pricing the barrier option.

Next, we conduct a sensitivity analysis over three key parameters: the barrier level  $B$ , strike price  $K$ , and volatility  $\sigma$ . In each case, we vary one parameter while keeping others fixed at:

$$T = 1.0, \quad r = 0.05, \quad S_0 = 100, \quad S \in [20, 140], \quad N = M = 200.$$

The specific parameter ranges explored are:

- $K \in [80, 120]$ , in 9 equally spaced values
- $B \in [105, 140]$ , in 8 equally spaced values
- $\sigma \in [0.1, 0.5]$ , in 9 equally spaced values

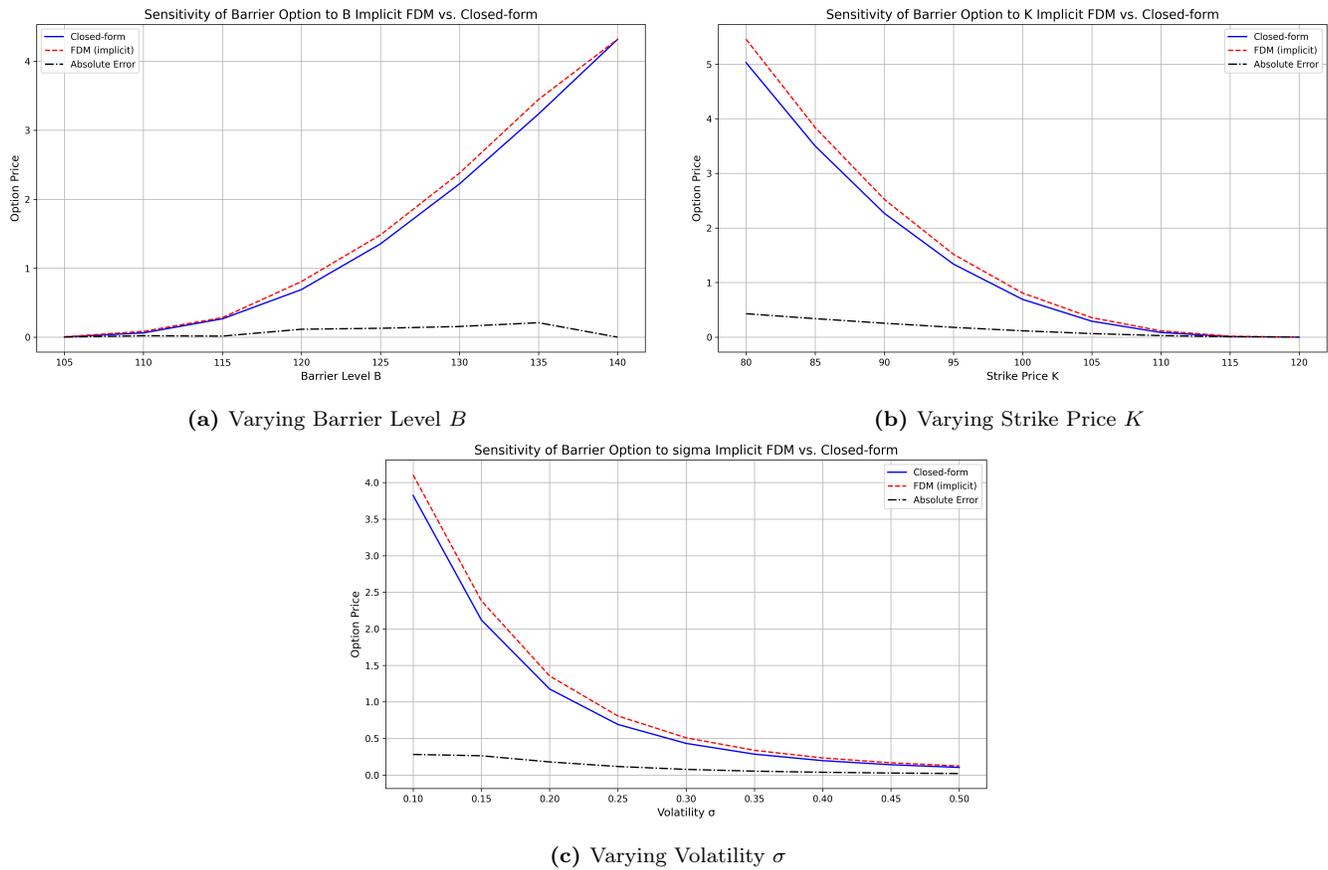
Figure 15 shows the pricing behavior with respect to each parameter, comparing closed-form vs FDM and visualizing the absolute error.

All three plots exhibit expected financial behaviors. As  $K$  increases, the option price decreases due to higher exercise cost. As  $B$  increases, the knock-out probability decreases, leading to higher prices. Greater volatility  $\sigma$  enhances the probability of reaching deep in-the-money regions, also increasing option value. The implicit FDM solution closely tracks the analytical one, with small deviations due to discretization.

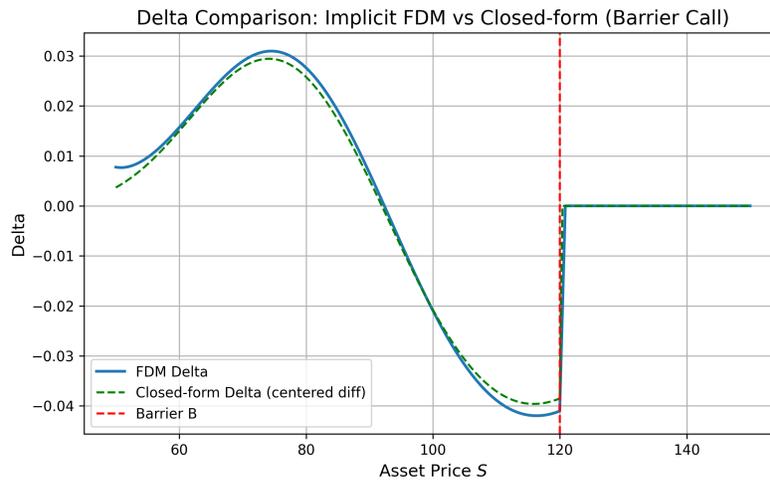
Lastly, seen Figure 16, we examine the option's sensitivity with respect to the underlying asset price  $S$ , via the delta  $\Delta = \partial C / \partial S$  curve. This is computed both from the FDM price grid and by applying centered finite differences to the closed-form solution. The parameters used are:

- $K = 100, B = 120, T = 1.0, r = 0.05, \sigma = 0.25$
- $S \in [50, 150]$ , with fine resolution  $N = 300, M = 200$

The delta curve reflects the complex nature of barrier options: sensitivity increases when  $S$  is moderately below  $B$ , and abruptly drops to zero at the barrier. This is correctly captured by both methods. The sharp discontinuity at  $S = B$  further confirms the barrier's impact, and the smooth agreement before the barrier showcases the robustness of our FDM scheme.



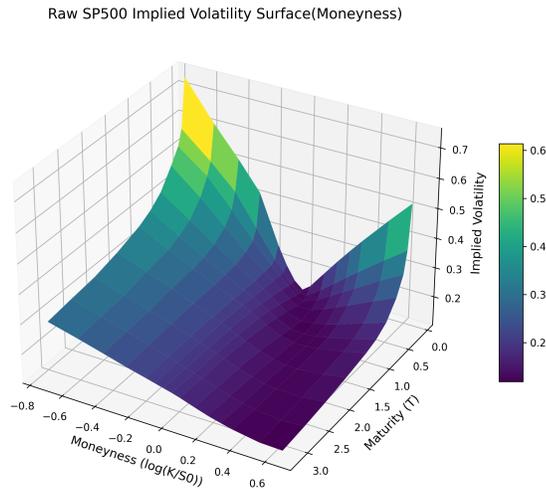
**Figure 15:** Sensitivity analysis: option price behavior vs parameters  $(B, K, \sigma)$ , comparing implicit FDM vs closed-form with absolute error shown.



**Figure 16:** Comparison of delta  $\Delta = \frac{\partial C}{\partial S}$  curves from implicit FDM and closed-form method (centered difference), under parameter setup  $K = 100, B = 120, \sigma = 0.25$ .

### 3.3 Calibration SP500 Implied Volatility using the Heston Model

#### 3.3.1 Vanilla option pricing experiments



**Figure 17:** Raw SP500 implied volatility surface (log-moneyness vs maturity) on 2023-11-01.

We begin our experiments by visualising the SP500 market data on 1 Nov 2023, specifically the raw implied volatility surface as a function of moneyness and maturity. The raw dataset is extracted from `raw_ivol_surfaces.npy`, where each smile is plotted with respect to  $\log(K/S_0)$ . Since the actual spot price  $S_0$  is not directly available, we estimate it heuristically by selecting the strike  $K^*$  with the lowest implied volatility at the shortest maturity (i.e.  $T \approx 0$ ). This choice corresponds to the local minimum of the volatility smile, and serves as a natural proxy for the forward-adjusted spot. The resulting inferred value is:

$$S_0^{\text{guess}} = 4449.753$$

This guess is used to convert all strikes into standardised log-moneyness coordinates, forming the surface shown in Figure 17.

Next, we simulate vanilla European call prices under the Heston model, using the following fixed parameter set:

$$\theta = 0.05, \quad \kappa = 2.5, \quad \sigma = 0.4, \quad \rho = -0.3, \quad V_0 = 0.04, \quad S_0 = 100, \quad r = 0.04. \quad (2)$$

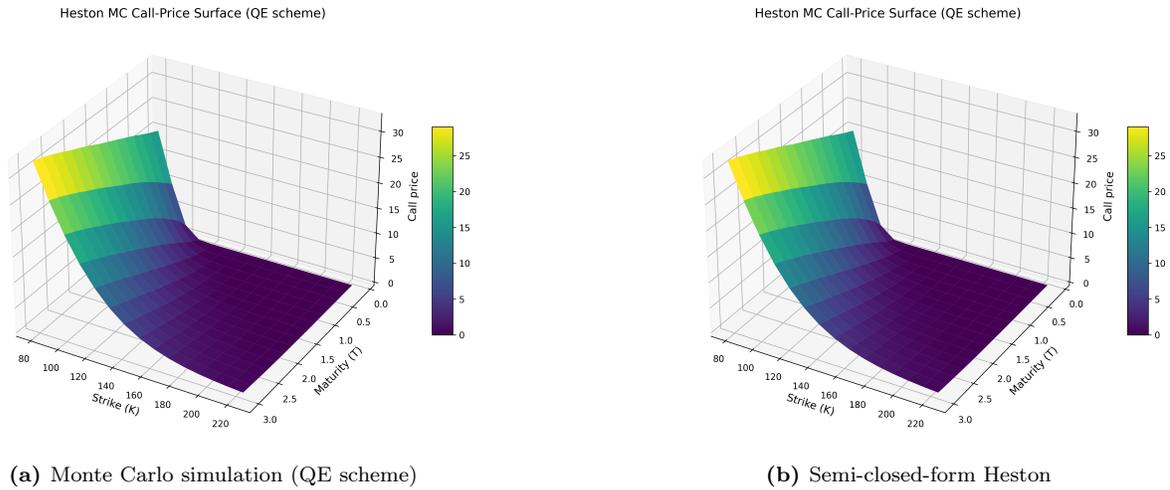
The option prices are computed over a two-dimensional grid consisting of:

- 15 strikes, uniformly spaced from  $K = 80$  to  $K = 225$ ;
- 15 maturities, linearly spaced from  $T = 0.1$  to  $T = 3.0$ .

Call-price surfaces are then generated using two pricing engines:

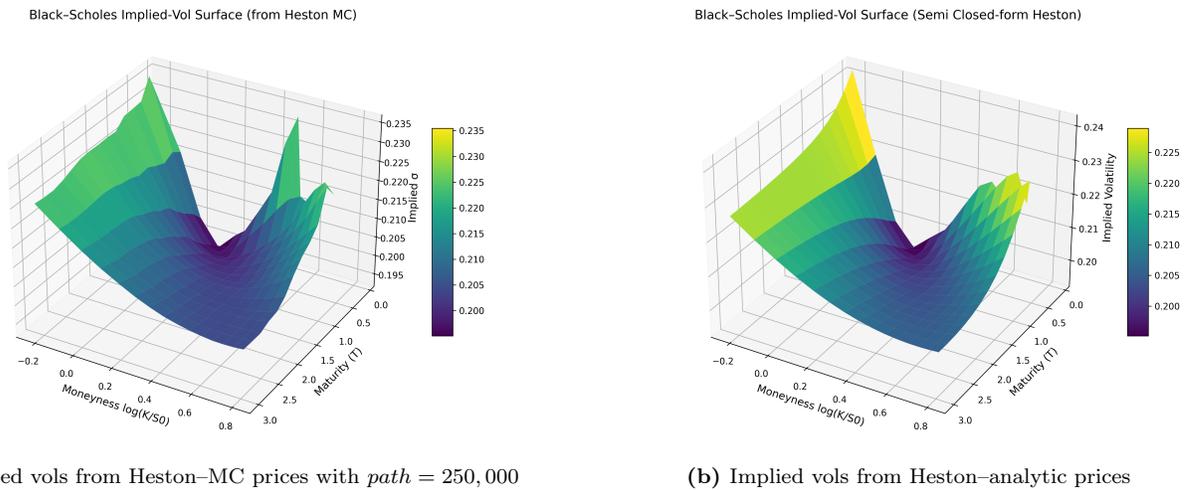
- (i) a Monte Carlo engine based on the Quadratic–Exponential (QE) scheme, with  $N = 250,000$  simulated paths per option;
- (ii) the semi-closed form solution derived from the Heston model’s characteristic function.

Figure 18 shows both surfaces side-by-side. They exhibit consistent shape and scale, validating the correctness of the Monte Carlo engine against the analytical benchmark.



**Figure 18:** Call-price surfaces generated from the Heston model using two pricing methods.

To facilitate implied-volatility comparisons, we invert the Black-Scholes formula to extract the *implied volatilities* associated with each call price. Figure 19 presents the two implied-vol surfaces. While Monte Carlo simulation introduces minor surface irregularities, particularly in short-maturity or far-from-the-money regions, the global shape and skew are consistent with the semi-closed-form counterpart and the original market surface.



**Figure 19:** Black-Scholes implied volatility surfaces extracted from Heston-generated call prices.

### 3.3.2 Calibration to SP500 implied volatilities

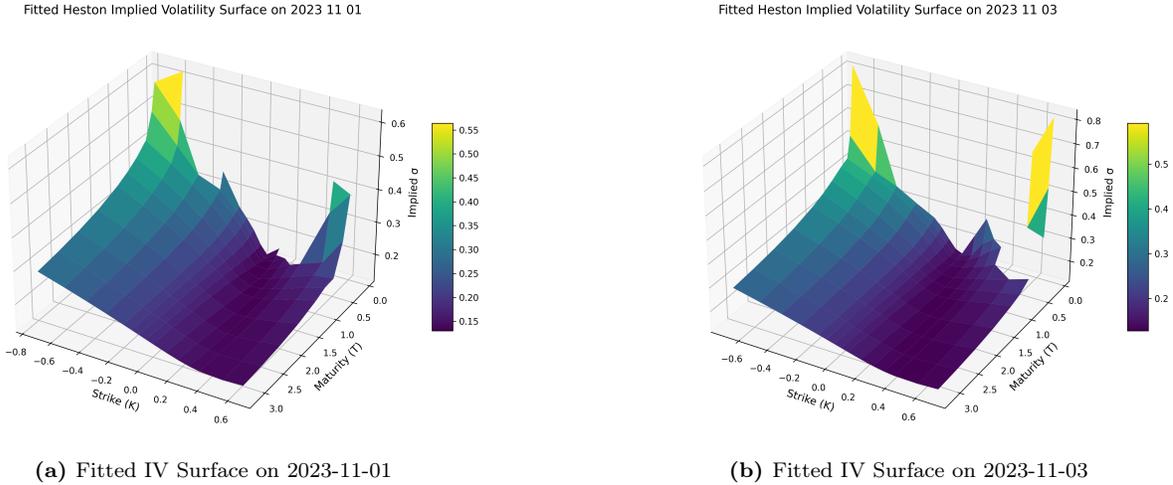
To assess the Heston model’s ability to fit market-observed implied volatility (IV) surfaces, we perform calibration on two SP500 trading dates: **2023-11-01** and **2023-11-03**. The calibration objective is defined as the root-mean-square error (RMSE) between model-generated implied volatilities and the market-observed values, uniformly weighted.

**Calibration Methodology.** We employ a least-squares optimization using `scipy.optimize.minimize` with the L-BFGS-B algorithm. The following initial guess and parameter bounds are used:

$$\begin{aligned} \text{Initial guess: } & [\kappa, \theta, \sigma, \rho, V_0] = [2.5, 0.05, 0.4, -0.3, 0.04] \\ \text{Bounds: } & \kappa \in [0.1, 10], \quad \theta \in [0.01, 0.5], \quad \sigma \in [0.05, 1.0], \quad \rho \in [-0.99, 0.99], \quad V_0 \in [0.001, 0.2] \end{aligned}$$

For each date, the market strike and maturity grid is standardized using a guessed spot price (estimated from short-tenor ATM volatility minimum). The model is run under fixed  $S_0 = 100$ , and moneyness alignment ensures parameter consistency between model and market.

**Fitted IV Surfaces.** The fitted implied volatility surfaces obtained using the optimal parameters are shown below:

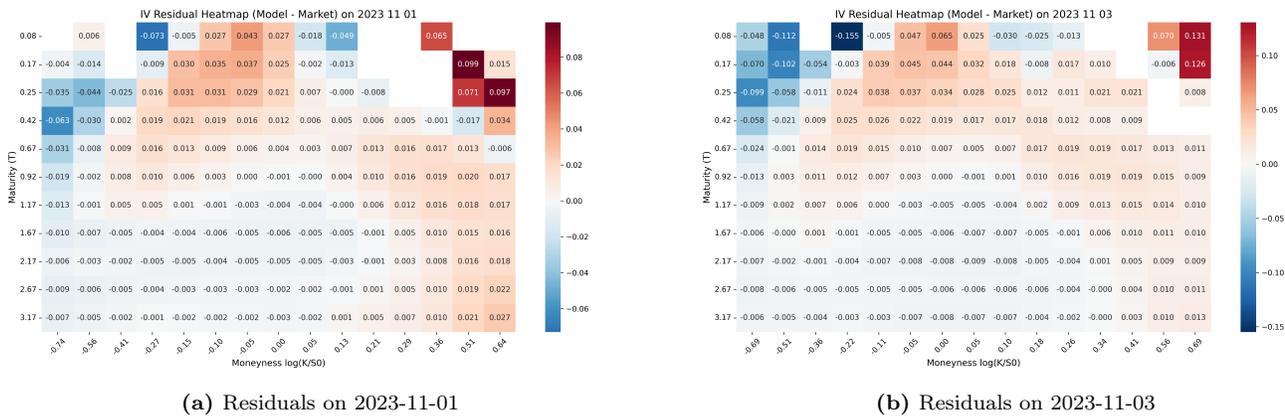


**Figure 20:** Heston model-fitted implied volatility surfaces on two SP500 dates.

**Calibration Results.** The resulting parameter estimates and error diagnostics are summarized below:

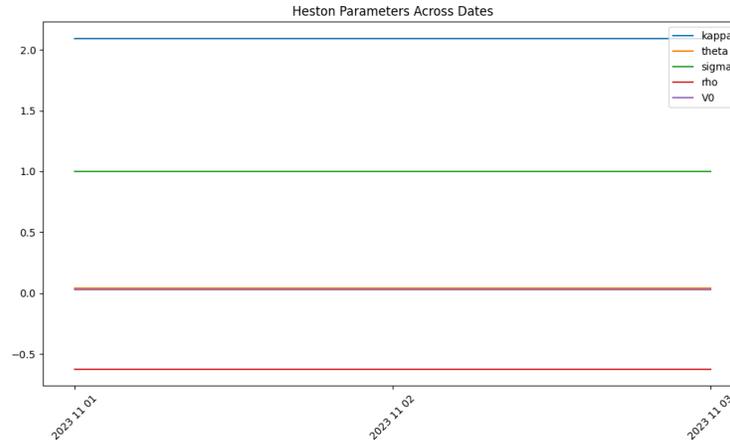
- **2023-11-01:**
  - Optimal parameters:  $\kappa = 2.094$ ,  $\theta = 0.0439$ ,  $\sigma = 1.000$ ,  $\rho = -0.625$ ,  $V_0 = 0.028$
  - RMSE = 0.0214, MAE = 0.0135, Max Error = 0.0993
- **2023-11-03:**
  - Optimal parameters:  $\kappa = 2.002$ ,  $\theta = 0.0414$ ,  $\sigma = 1.000$ ,  $\rho = -0.681$ ,  $V_0 = 0.0399$
  - RMSE = 0.0305, MAE = 0.0175, Max Error = 0.1548

**Residual Diagnostics.** To identify regions of model under/overfitting, we visualize the residuals between model-generated and market-observed IVs.



**Figure 21:** Heatmaps of model-market IV residuals. Positive values indicate model overestimation.

**Parameter Stability and Fit Deterioration.** To examine model robustness across time, we performed Heston parameter calibration on three distinct trade dates. The resulting optimal parameters are visualized in Figure 22. The key parameters —  $\kappa$ ,  $\rho$ , and  $\sigma$  — exhibit remarkable stability across all dates, with only minor variations. This consistency suggests that the model captures persistent market dynamics. In contrast, the mean-reversion level  $\theta$  and the initial variance  $V_0$  display slight fluctuations, reflecting adaptations to shifting volatility regimes.



**Figure 22:** Optimal Heston parameters calibrated across dates 2023-11-01, 2023-11-02, and 2023-11-03.

The residual heatmaps (see Figure 21) reveal characteristic patterns of model mismatch. While the Heston model accurately captures the implied volatility smile around at-the-money strikes and mid-term maturities, the following issues arise:

- **Short maturities.** For options with very short time-to-maturity ( $T < 0.25$ ), the model tends to *underfit* the sharp skew observed in deep OTM options. This is especially visible in the top-left and top-right corners of the residual heatmaps.
- **Extreme strikes.** Maximum fitting errors concentrate at the edges of the moneyness range, both in deep in-the-money and deep out-of-the-money regions. These discrepancies highlight the Heston model's limitations in extrapolating extreme tail behavior.
- **Temporal variation.** Comparing residuals on November 1 and 3, the model exhibits slightly larger errors on November 3, with the RMSE increasing from 0.021 to 0.031. This degradation may stem from increased market volatility or irregularities in the observed smile shape.

**Model Extensions.** While the classical Heston model successfully captures the volatility smile and term structure of implied volatilities under normal market conditions, it struggles to fit short-term skews and extreme strike behaviors, as observed in our residual diagnostics. One natural extension is to incorporate jump components into the underlying asset price process.

A prominent example is the *Heston-Jump* model, which augments the standard stochastic volatility dynamics with a compound Poisson process to capture discontinuous price movements. This approach is supported by the work of Duffie, Pan, and Singleton Duffie et al., 2000, who demonstrate that combining stochastic volatility with jump risk significantly improves fit to short-maturity option prices and extreme-strike implied volatilities.

Such jump-diffusion models allow greater flexibility in capturing the heavy tails and sudden shifts observed in empirical returns, thereby reducing the persistent misfit observed at the boundaries of our implied volatility surface. Moreover, the characteristic function framework of the original Heston model naturally extends to this jump setting, preserving numerical tractability for calibration and pricing.

## 4 Conclusion

### 4.1 PDE Methods: Binary and Barrier Options

The analytical solutions for both binary and barrier options offer reliable benchmarks, with the numerical methods showing good agreement across most of the domain. For binary options, the FDM schemes capture the correct surface shape, with Crank-Nicolson producing smoother results, albeit with sensitivity near payoff discontinuities.

In the barrier case, MC simulations closely match the closed-form values, particularly when the barrier correction is applied. Pricing errors are small except near vanishing payoffs. Convergence tests confirm the expected rate of

$O(1/\sqrt{M})$  with increasing time steps.

Together, these experiments highlight the strengths and trade-offs of numerical techniques when applied to discontinuous or path-dependent payoffs, reinforcing the need for careful discretization, boundary handling, and theoretical consistency in financial option pricing.

## 4.2 Calibration SP500 Implied Volatility using the Heston Model

We have systematically implemented the Heston stochastic volatility model across both theoretical and empirical tasks. In the vanilla pricing experiments, we confirmed that the semi-closed formula and Monte Carlo methods yield consistent call-price surfaces. Using inverse Black–Scholes transformation, we recovered implied volatility surfaces that qualitatively match the stylized smile and skew observed in markets.

For the calibration task, we optimized Heston parameters over two SP500 trading dates (2023-11-01 and 2023-11-03). The calibrated models achieved reasonably low RMSE values (0.0214 and 0.0305), and residual diagnostics showed that the model fits are most accurate in the near-the-money and mid-maturity regions. However, performance deteriorates at short tenors and extreme moneyness—a known limitation of the standard Heston framework.

Parameter trajectories across dates were found to be relatively stable, particularly for the key long-term mean  $\theta$ , speed of mean reversion  $\kappa$ , and correlation  $\rho$ . This suggests robustness in capturing structural volatility features. Nevertheless, higher errors in certain regions motivate model extensions such as time-dependent coefficients or jump components, which can better account for observed skews and short-term volatility bursts.

This part of study demonstrates the Heston model's practical effectiveness while highlighting calibration challenges and the necessity of model refinements for high-fidelity fitting.

## 5 References

- Andersen, L. B. G. (2006, November). *Efficient simulation of the heston stochastic volatility model* [Working paper, Bank of America]. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=946405](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=946405)
- Cui, Y., del Baño Rollin, S., & Germano, G. (2017). Full and fast calibration of the heston stochastic volatility model. *Quantitative Finance*, 17(2), 149–167. <https://doi.org/10.1080/14697688.2016.1198680>
- Duffie, D., Pan, J., & Singleton, K. J. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6), 1343–1376.