Dipartimento di Economia e Finanza

Indirizzo Banche ed Intermediari Finanziari

Option Pricing: Heston Model Calibration

Relatore:

Prof. Hlafo Alfie Mimun

Prof. Guido Germano

Correlatore:

Prof. Sara Biagini

Candidato:

Andrea Pezzella

785231

Anno Accademico 2024/2025

Alla mia Famiglia,

per avermi insegnato ad essere.

Per aspera ad astra

Contents

Introduction

Option pricing remains among the most critical topics in quantitative finance. Options are financial derivatives whose value depends on the behavior of an underlying asset. They are widely used for purposes such as speculation, portfolio insurance, and risk management. As such, constructing accurate and robust pricing models for these instruments is important.

The classical model developed by Black and Scholes in 1973 marked a turning point in option pricing theory. By assuming constant volatility and a lognormal distribution for asset prices, the Black-Scholes framework allows for the derivation of a closed-form solution for European options. Despite its elegance and analytical tractability, the model fails to capture several well-documented empirical features observed in real markets.

One of the most notable shortcomings of the Black-Scholes model is its inability to reproduce the so-called "volatility smile" and "volatility skew", that are patterns observed in implied volatilities across different strikes and maturities. In practice, implied volatility is not flat, as the Black-Scholes model would suggest, but instead varies systematically.

To address these inconsistencies, more sophisticated models have been proposed. One of the most prominent alternatives is the Heston model, introduced by Steven Heston in 1993. Unlike Black-Scholes, the Heston model assumes that volatility itself is a stochastic process, introducing randomness into its evolution

over time. This added flexibility enables the model to capture the observed smile effects.

Although other approaches exist, for example, the local volatility model developed by Dupire, which treats volatility as a deterministic function of time and the asset level. The core focus of this thesis lies in the implementation, calibration, and empirical testing of the Heston model. In particular, we aim to assess its ability to replicate the real market implied volatility surface by calibrating it to a dataset of European options.

The calibration is carried out using equity options on Bank of America Corporation (BAC), retrieved from Bloomberg Terminal. These data provide implied volatilities across a range of maturities and deltas.

The Heston model is calibrated by minimizing the discrepancy between observed market prices and those generated by the model. A nonlinear least squares optimization is applied to find out the five model parameters: long-term variance, mean-reversion speed, volatility of volatility, initial variance, and correlation between asset returns and variance. Calibration is performed across a grid of deltas and maturities, following standard market conventions.

To validate the calibration results, we compare the implied volatility surface generated by the Heston model with that obtained from market data. As a benchmark, we also compute the implied volatility surface resulting from the Black-Scholes model using the market implied volatilities directly. This comparison highlights the Heston model's ability to better fit the market implied volatility. The improvement is also supported by quantitative error metrics such as the mean absolute error and the root mean square error.

Through this work, we aim to offer both a rigorous theoretical foundation and a concrete empirical application of stochastic volatility modeling.

Chapter 1 provides an introduction to the world of derivatives, focusing on the

fundamental instruments such as forward contracts, futures, and swaps. While a general overview is presented, the core of the chapter is dedicated to option contracts.

Chapter 2 establishes the mathematical foundation for option pricing, beginning with the modeling of stock prices via stochastic processes. In particular, the geometric Brownian motion is introduced as the standard model for asset dynamics, reflecting the log-normality assumption of returns. From this basis, we derive the celebrated Black-Scholes model for European options, including the derivation of the Black-Scholes partial differential equation (PDE) and introduce the closed-form solutions for option prices. Moreover, we discuss the practical limitations that emerge when applying this model to real-world data.

Chapter 3 shifts the focus to the concept of implied volatility, a key market-based input for option pricing. By inverting the Black-Scholes formula, practitioners can infer the volatility that is "implied" by observed market prices. When plotted across different strikes and maturities, the implied volatility surface exhibits systematic patterns, commonly known as the volatility smile and volatility skew (or smirk). These patterns reveal the failure of the constant volatility assumption in the Black-Scholes framework and motivate the exploration of alternative models. In this context, we introduce the concept of local volatility and present the Dupire formula.

The core modeling framework explored in this thesis is presented in Chapter 4, which is devoted to the Heston model. This stochastic volatility model assumes that the variance of the underlying asset evolves according to a mean-reverting square-root process, introducing an additional source of randomness. The chapter includes the derivation of the Heston PDE for option prices, and show the well-known semi-analytical solution obtained via Fourier inversion techniques.

Chapter 5 constitutes the heart of this thesis, focusing on the calibration of

the Heston model to real market data. The calibration process refers to the procedure of adjusting model parameters so that the resulting option prices or implied volatilities closely match observed market values. The chapter begins with a general discussion on model calibration. We then delve into the specific calibration strategy implemented for this thesis, which includes the optimization of the Heston model parameters using MATLAB. A section of the chapter is dedicated to describing the empirical dataset used in the calibration. The results of the calibration are then thoroughly presented and interpreted. The fitted Heston surface is compared with the market surface and with the one produced by the Black-Scholes model. Finally, the chapter closes with a financial interpretation of the calibrated parameters and a brief explanation of the use of calibration in finance.

Chapter 1

Derivatives

Over the last decades, derivatives have become increasingly utilized in the financial industry. Some of these financial instruments, like futures and options, are widely traded on international exchanges, and other derivatives, like forwards and swaps, are traded over the counter. The derivatives market is huge, it is much bigger than the stock market when misured in terms of the value of underlying assets. Nowadays, it is essential that all professionals working in this field understand how derivatives work, how they are used, and how they are priced.

A derivative is a contract that involves two parties agreeing to a future transaction. Its value depends on the values of underlying variables. Often, the underlying variable is represented by the value of a tradable asset, such as in the case of stock option, that is a derivative whose value is dependent on the price of a stock. The fact that the underlying variables of derivatives can be of any imaginable type, such as interest rates, weather conditions, or even the outcome of a specific event, and that these variables can influence the terms and the payoff of the contract, makes derivatives really flexible and perfect for investment and risk management strategies.

Derivatives can be used for hedging, which involves reducing risk by providing a

way to insure against price movements. Investors also use derivatives to profit from price fluctuations of underlying assets, attempting to predict the direction in which the prices of assets will move. Moreover, traders use derivatives to profit from discrepancies in prices by engaging in arbitrage which is the practice of exploiting prices differences of the same or similar financial instruments on different markets or in different forms.

As we said, derivatives market is enormous, given that either exchange-traded derivatives and Over-the-counter(OTC) derivatives are enormous. However, even if the number of derivative transaction per year in OTC markets is smaller than in exchange-traded markets, the average size of the transactions is much greater. This means that the volume of the business in the OTC markets is much larger than in exchange-traded markets.

1.1 Forward contracts

Let's look at the various types of derivatives. A forward contract is an agreement to buy or sell an asset at a certain future time for a certain price. One of the parties to a forward contract assumes a long position and agrees to buy the underlying asset on a certain specified future date for a certain specified price. We can say that it is the opposite of a spot contract[1]. One of the main characteristics is that a forward contract is traded in the OTC market.

1.1.1 Forward Payoff

Let K be the delivery price and S_T the spot price of the asset at maturity of the contract. We can say that, in general, the payoff from a long position in a forward contract on one unit of an asset is $g(S_T) = S_T - K$. In a similar way, the

[1]It is an agreement to buy or sell an asset almost immediately

payoff from a short position in a forward contract to one unit of an asset is $K - S_T$. Obviously, the payoff can be positive or negative, because we must remember that the holder of the contract is obligated to buy an asset worth S_T for K, in case of a long position in the contract, or is obligated to sell an asset worth S_T for K, in case of a short position.

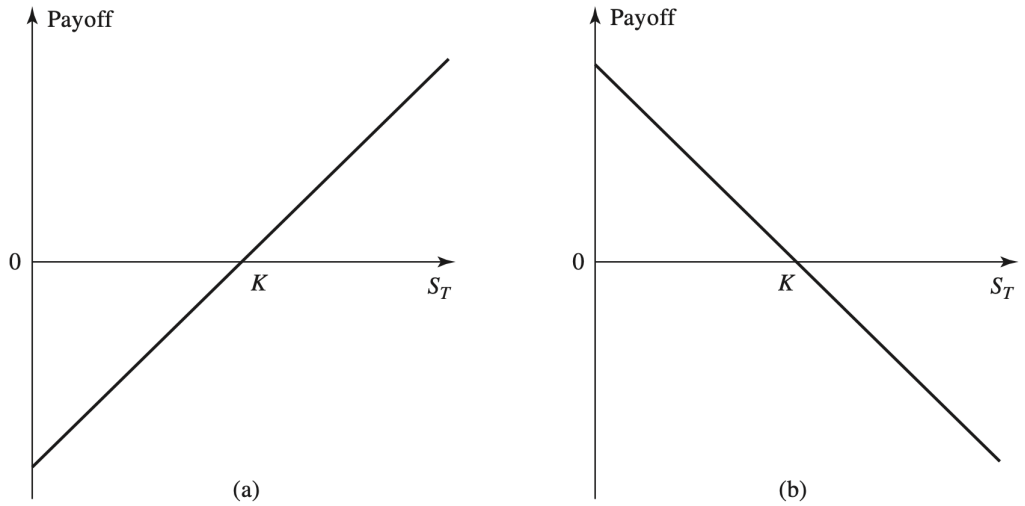

Fig. 1.1: Payoffs from forward contracts: (a) long position, (b) short position. Source: [10]

1.2 Futures Contracts

This type of derivative is similar to forward contract. The difference between the two contracts is that future contracts are usually traded on an exchange. To do this, the exchange specifies certain standardized characteristics of the contract, this is done to make the trading possible. The clearing house acts as an intermediary for all futures transactions. Infact, the main goal of the clearing house is to mitigate the risk of counterparty default by guaranteeing the performance of both parties in the contract. The advantage of this is that traders do not have to worry about

the creditworthiness of the people they are trading with. The clearing house meke this job by requiring each of the two traders to deposit funds with the clearing house to ensure that they will live up their obligations.

1.3 Option Contracts

Options ate contracts that are traded both on exchanges and in OTC market. The contract is done like the other two derivatives that we have seen but with the difference that the option gives the holder the right to buy or sell the underlying asset by a certain date, called *expiration date* or *maturity*, for a certain price, called *Exercise Price* or *strike price*. The holder has no duty to exercise the right. If the Holder has the right to buy, the option is called **call**, otherwise it is called **put**. Another important difference is between *American option* that can be exercised in any time up to the expiration date, while *European option* can be exercised only on the expiration date itself. In the exchange-traded equity option market, one contract is usually an agreement to buy or sell 100 shares. Thus, there are four types of participants:

- Buyers of calls,
- Sellers of calls,
- Buyers of puts,
- Sellers of puts.

Selling an option is known as *writing an option*.

1.3.1 Option Payoff

It is useful to characterize an option in terms of its payoff, that doesn't include the initial cost of the option. If K is the strike price and S_T is the price of the

underlying asset at the end of the contract, the payoff from a long position in a European call option is

$$g(S_T) = \max(S_T - K, 0). \tag{1.1}$$

It occurs because the option will be exercised only if $S_T > K$, otherwise the option will not be exercised, thus the payoff is 0. On the other hand, the payoff to the holder of a short position in the European call option is $-\max(S_T - K, 0) = \min(K - S_T, 0)$. The payoff to the holder of a long position in a European put option is $g(S_T) = \max(K - S_T, 0.)$ The payoff from a short position in a European put option is $g(S_T) = -\max(K - S_T, 0) = \min(S_T - K, 0)$.

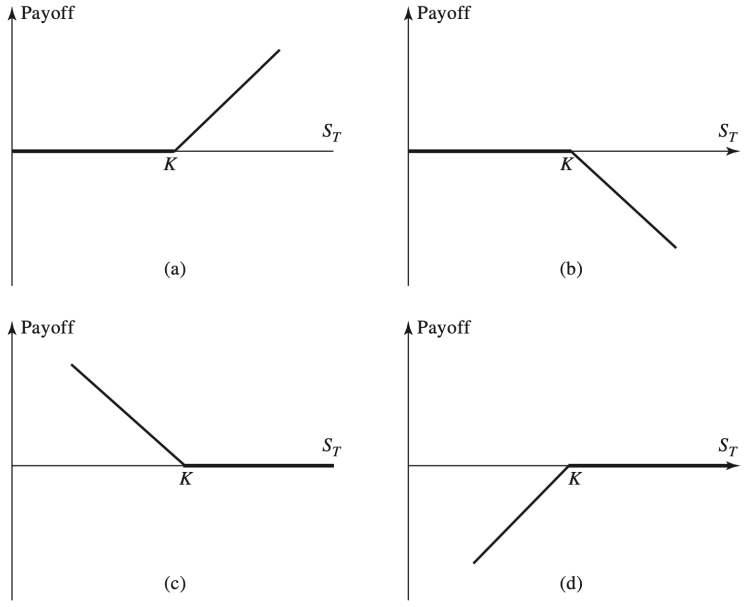

Fig. 1.2: Payoffs from option contracts: (a) long position on call, (b) short position on call, (c) long position on put, (d) short position on put. Source: [10]

1.4 Put-Call Parity

This relationship shows that the value of a European Call with a certain exercise price and exercise date can be deduced from the value of a European put with the same exercise price and exercise date, and vice versa.

Proposition 1.4.1 (Put-Call Parity)**.** *Let c be the price of a European Call option and p the price of a European Put option, then the following equation holds:*

$$c + Ke^{-rT} = p + S_0$$

Where:

- *K is the strike prize of the option*
- *e^{-rT} represents the exponential discount factor, where r is the risk-free rate and T is the time to expiration*
- *S_0 is the current price of the underlying asset*

This equation is known as put-call parity.

Note: If the put-call parity is not maintained, there would be arbitrage opportunities.

Chapter 2

Black-Scholes Model

The Black-Scholes model revolutionized financial economics by providing the first rigorous framework for option pricing. Developed in 1973, it introduced a systematic way to determinate the fair value of options under idealized market conditions. These "ideal conditions" in the market for the stock and for the options are:

1. The short-term interest rate is known and is constant through time.

2. The stock price follows a random walk in continuous time with a variance rate proportional to the square of the stock price. Thus the distribution of possible stock prices at the end of any finite interval is log-normal. The variance rate of the return on the stock is constant.

3. The stock pays no dividends or other distributions.

4. The option is a "European Option".

5. There are no transaction costs[1].

[1]Costs for buying or selling the stock or the option

6. It is possible to borrow any fraction of the price of a security to buy or to hold it, at the short-term interest rate.

7. There are no penalties to short selling[2].

Although its assumptions simplify reality, the model's analytical elegance and practical utility made it a cornerstone of modern finance.

2.1 Geometric Brownian Motion

Suppose that $S(t)$ denotes the stock price. A lot of random quantities are normally distributed but it is not so reasonable for $S(t)$, given that the stock price can not be negative, while with normal distribution, a normal random quantity X can satisfy $P[X < 0] > 0$, thus negative stock prices. Therefore, a more appropriate alternative is the log-normal distribution because preserve $S(t) > 0$. A random variable X is log-normal if it is strictly positive and its logarithm, $\log(X)$, is normally distributed. This property aligns well with the characteristics of stock prices. Define the accumulation factor from time t to $u > t$ as

$$A(t, u) = \frac{S(u)}{S(t)},$$

If $A(t, u) > 1$ means that the stock price increases. On the other hand, if $A(t, u) < 1$, it decreases. The proportional change in the stock price is

$$A(t, u) - 1 = \frac{S(u) - S(t)}{S(t)}.$$

We are interested in the values of $\mathbb{P}[A(t, u) < x]$ for any given x, and for any given t, u. To facilitate modeling, we introduce the following assumptions:

[2]A seller who does not own a security will simply accept the price of the security from a buyer, and will agree to settle with the buyer on some future date by paying him an amount equal to the price of the security on that date

- The distribution of $A(t, t+h)$ depends only on the length of the interval h and is independent of the starting time t.
- The distribution is even independent of the current stock price $S(t)$ and any historical prices.

Whether these assumptions imply that the probability distribution of the proportional change in the stock price over the next h units of time does not depend on its current value or on its historical values. Let us fix $t < u$. Consider n as a positive integer, and for each period $j = 0, 1, 2, ..., n$, define the points in time $t_j = t + (j/n)(u - t)$. As $n \to \infty$, the number of intermediate points increases and the time steps between them shrink to zero. In this limiting process, we transition from a discrete-time multiplicative model to a continuous-time formulation. This leads us naturally to model the stock price dynamics using a stochastic differential equation (SDE), which accurately captures the infinitesimal behavior of $S(t)$ over time. Hence, we have $t_0 = t$ and $t_n = u$. It follows that

$$\frac{S(u)}{S(t)} = \frac{S(t_n)}{S(t_0)} = \left(\frac{S(t_1)}{S(t_0)}\right)\left(\frac{S(t_2)}{S(t_1)}\right)\left(\frac{S(t_3)}{S(t_2)}\right)\cdots\left(\frac{S(t_n)}{S(t_{n-1})}\right).$$

Thus,

$$A(t, u) = A(t_0, t_1)A(t_1, t_2)\cdots A(t_{n-1}, t_n).$$

Now, if we define $L(t, u) = \log A(t, u)$, then we have

$$\begin{aligned} L(t, u) &= \log[A(t_0, t_1)A(t_1, t_2)\ldots A(t_{n-1}, t_n)] \\ &= \log[A(t_0, t_1)] + \log[A(t_1, t_2)] + \ldots + \log[A(t_{n-1}, t_n)] \\ &= L(t_0, t_1) + L(t_1, t_2) + \ldots + L(t_{n-1}, t_n). \end{aligned}$$

From our assumptions, the random variables $L(t_{j-1}, t_j)$, $1 \leq j \leq n$, are independent and identically distributed (i.i.d). Given their i.i.d nature, when n is

sufficiently large, we can invoke the central limit theorem to argue that $L(t, u)$ is approximately normally distributed.

Under our assumptions, the distribution of $L(s, s + t)$ depends only on t. As previously stated, we further assume it follows a normal distribution with mean $m(t)$ and variance $v(t)$. Additionally, the relationship

$$L(0, t + u) = L(0, t) + L(t, t + u), \tag{2.1}$$

rapresents a sum of independent normal variables. By computing the mean and variance in eq. (2.1), we derive

$$m(t + u) = m(t) + m(u),$$
$$v(t + u) = v(t) + v(u).$$

These functional equations imply that both $m(t)$ and $v(t)$ scale linearly with t. Hence, there exists a non-negative constant σ such that

$$m(t) = ct,$$
$$v(t) = \sigma^2 t.$$

This linearity has two key interpretations: First, the expected increment of the process grows proportionally with time, meaning that over longer intervals, the process tends to increase at a constant rate c. Second, the variability of the accumulated increments also increases linearly with time, indicating that extended observation periods lead to greater dispersion in the process's trajectory. Consequently, for any $t < u$, we have

$$\log \frac{S(u)}{S(t)} = \log A(t, u) = L(t, u)$$
$$\sim N(c(u - t), \sigma^2(u - t)).$$

Let's work through an example assuming $S(0) = 1$. We want to compute the expected value $\mathbb{E}[S(t)]$.

Recall that

$$S(t) = A(0,t) = \exp L(0,t),$$

and

$$L(0,t) \sim N(ct, \sigma^2 t),$$

meaning it can be expressed as $L(0,t) = ct + \sigma\sqrt{t}Z$, where Z is a standard normal random variable, i.e., $Z \sim N(0,1)$. (This holds because a random variable X that follows a normal distribution $N(\mu, \sigma^2 t)$ can be written as $X = \mu + \sigma\sqrt{t}Z$, where σ is the standard deviation.)

At this point, it is useful to clarify the distinction between Z, dZ, and dW_t, as these notations represent different but related concepts in the modeling framework. This will become especially important as we transition from discrete-time reasoning to continuous-time stochastic calculus.

- Z is a standard normal variable, typically used in expressions like $X = \mu + \sigma Z$, where $X \sim N(\mu, \sigma^2)$.
- W_t is a Brownian motion (or Wiener process), a continuous-time process with $W_0 = 0$ and increments such that $W_t - W_s \sim N(0, t-s)$ for $t > s$.
- dW_t is the infinitesimal increment of W_t over an interval dt, and is informally interpreted as $dW_t \sim N(0, dt)$. It satisfies $\mathbb{E}[dW_t] = 0$ and $\mathbb{E}[dW_t^2] = dt$.
- dZ is sometimes used interchangeably with dW_t in the finance literature, although dW_t is the more rigorous notation in continuous-time models.

For this reason

$$\begin{aligned} S(t) = A(0,t) &= \exp(ct + \sigma\sqrt{t}Z) \\ &= e^{ct} e^{\sigma\sqrt{t}Z}. \end{aligned}$$

Now, taking the expectation:

$$\mathbb{E}[S(t)] = \mathbb{E}[e^{ct}e^{\sigma\sqrt{t}Z}] = e^{ct}\mathbb{E}[e^{\sigma\sqrt{t}Z}].$$

We want to compute the expected value of $e^{\sigma\sqrt{t}Z}$, where $Z \sim \mathcal{N}(0,1)$. Using the definition of expectation for continuous random variables:

$$\begin{aligned}\mathbb{E}[e^{\sigma\sqrt{t}Z}] &= \int_{-\infty}^{\infty} e^{\sigma\sqrt{t}z} \cdot \frac{1}{\sqrt{2\pi}} e^{-z^2/2}\, dz \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{\sigma\sqrt{t}z - \frac{1}{2}z^2}\, dz.\end{aligned}$$

We complete the square in the exponent:

$$\sigma\sqrt{t}z - \frac{1}{2}z^2 = -\frac{1}{2}\left(z^2 - 2\sigma\sqrt{t}z\right) = -\frac{1}{2}(z - \sigma\sqrt{t})^2 + \frac{1}{2}\sigma^2 t.$$

Substituting back into the integral:

$$\begin{aligned}\mathbb{E}[e^{\sigma\sqrt{t}Z}] &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(z-\sigma\sqrt{t})^2 + \frac{1}{2}\sigma^2 t}\, dz \\ &= e^{\frac{1}{2}\sigma^2 t} \cdot \underbrace{\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(z-\sigma\sqrt{t})^2}\, dz}_{1}.\end{aligned}$$

The part that we highlighted is the total area under the probability density function of a normal distribution with mean $\sigma\sqrt{t}$ and variance 1, which is equal to 1. Hence:

$$\mathbb{E}[e^{\sigma\sqrt{t}Z}] = e^{\frac{1}{2}\sigma^2 t}.$$

$$= e^{ct}e^{\sigma^2 t/2} = e^{((c+(\sigma^2/2))t)}.$$

Writing $\mu = c + \sigma^2/2$, we have: $\mathbb{E}[S(t)] = e^{\mu t}$. Let μ denote the expected growth rate of the stock price, while the risk-free interest rate r represents the growth rate of a risk-free investment. For instance, an investment of one unit at the risk-free rate at time 0 will grow to e^{rt} by time t.

We say that the stock price process follows a Geometric Brownian Motion (or Log-normal process) with an expected rate of return (or infinitesimal drift) μ and volatility σ.

It is important to note that μ, the expected return of this process, is related to c via:

$$\mu = c + \frac{\sigma^2}{2} \quad \text{so,} \quad c = \mu - \frac{\sigma^2}{2}.$$

Definition 2.1. *Let $\sigma > 0$ be the volatility and μ a constant. A Geometric Brownian Motion (or Log-normal process) with expected return μ and volatility σ is a family of random variables $\{S(t)\}_{t\geq 0}$ satisfying the following properties:*

- *For each $t \geq 0$,$u \geq 0$, $\log \frac{S(t+u)}{S(t)} \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)u, \sigma^2 u\right)$.*
- *For all $t \geq 0$,$u \geq 0$ and any sequence $0 \leq s_1 < s_2 < \cdots s_n \leq t$, $\frac{S(t+u)}{S(t)}$ is independent of $(S(s_1), S(s_2), \ldots, S(s_n))$.*

Using reasoning similar to that employed in computing $\mathbb{E}[S(t)]$, when $S(0) = 1$, We conclude that $S(t)$,$t \geq 0$ follows a log-normal process with infinitesimal drift μ and volatility σ. Moreover, if the value of $S(t)$ is known at time $t = T\mu \geq 0$ We have: $\mathbb{E}[S(u)] = S(T)e^{\mu(\mu - T)}$. An alternative way to understand the behavior of the log-normal process is by examining stock price changes over small time intervals.

Suppose the stock price $S(t)$ follows a log-normal process with expected return rate μ and volatility σ. For $t \geq 0$ and $h \geq 0$,the definition of the log-normal process gives:

$$\log\left(\frac{S(t+h)}{S(t)}\right) \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)h, \sigma^2 h\right).$$

Letting Z denotes a standard normal variable, we can express this as,

$$\log\left(\frac{S(t+h)}{S(t)}\right) = \left(\mu - \frac{\sigma^2}{2}\right)h + \sigma\sqrt{h}Z,$$

which implies

$$\frac{S(t+h)}{S(t)} = \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)h + \sigma\sqrt{h}Z\right).$$

Let $\lambda(t,h)$ represent the proportional change in stock price from time t to $t+h$:

$$\lambda(t,h) = \frac{S(t+h) - S(t)}{S(t)} = \frac{S(t+h)}{S(t)} - 1.$$

Thus:

$$\lambda(t,h) = \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)h + \sigma\sqrt{h}Z\right) - 1.$$

Using the Maclaurin series expansion of the exponential function ($e^x = 1 + x + \frac{x^2}{2} + \frac{x^3}{3} + \cdots$), for small x we have:

$$e^x - 1 \simeq x + \frac{x^2}{2},$$

and for a small h, we can apply this to $x = (\mu - (\sigma^2/2))h + \sigma\sqrt{h}Z$. By ignoring higher-order terms in h, we obtain:

$$\begin{aligned}\lambda(t,h) &\approx \left(\left(\mu - \frac{\sigma^2}{2}\right)h + \sigma\sqrt{h}Z\right) + \frac{1}{2}\sigma^2 h Z^2 \\ &= \mu h + \sigma Z\sqrt{h} + \frac{\sigma^2 h}{2}(Z^2 - 1).\end{aligned}$$

In this expression:

- the first term μh is deterministic.
- the second term $\sigma Z\sqrt{h}$ is normally distributed with mean zero and variance $\sigma^2 h$
- the third term has mean zero (since $\mathbb{E}[Z^2] = 1$) and a variance much smaller than the second term's (as it scales h rather than $\sqrt{h}$).

For sufficiently small h, the third term becomes negligible, yielding:

$$\lambda(t,h) \simeq \mu h + \sigma\sqrt{h}Z,$$

which implies:

$$\lambda(t,h) \simeq N(\mu h, \sigma^2 h).$$

In differential notation, we often write:

- dt instead of h for an infinitesimal time increment
- $dS(t)$ for the stock price change in the time window $[t, t+dt]$
- dZ (or dW) for a normal random variable with mean 0 and variance dt, independent of prior events

This leads to the standard representation: $\frac{dS(t)}{S(t)} = \lambda(t,dt) = \mu dt + \sigma dZ$, or equivalently: $\frac{dS}{S} = \mu dt + \sigma dZ$. This expression describes the instantaneous rate of return of the stock. It consists of a deterministic part μdt (the drift) and a stochastic part σdZ (the volatility-driven random shock).

and:

$$dS = \mu S dt + \sigma S dZ. \tag{2.2}$$

2.2 Black-Scholes Model Derivation

To derive the model, we first explore Itô's Lemma then we obtain the Black-Scholes PDE, and with this we derive the Black-Scholes Formula for the option prices.

2.2.1 Itô's Lemma

Itô's Lemma is the stochastic counterpart of the chain rule in classical calculus. It allows us to compute the differential of a function that depends on a stochastic

process.

Theorem 2.2.1. *Suppose that X_t is a diffusion process that satisfies the following stochastic differential equation:*

$$dX_t = a(X_t, t)\,dt + b(X_t, t)\,dW_t, \tag{2.3}$$

where $a(X_t, t)$ is the drift term, $b(X_t, t)$ is the diffusion (volatility) term, and W_t is a standard Brownian motion.

Let $u(x, t)$ be a function such that $u \in C^{2,1}(\mathbb{R} \times [0, \infty))$, i.e., u is twice continuously differentiable in x and once in t. Define the process $Y_t = u(X_t, t)$. Then, Y_t evolves according to the following stochastic differential equation:

$$du(X_t, t) = \frac{\partial u}{\partial t}(X_t, t)\,dt + \frac{\partial u}{\partial x}(X_t, t)\,dX_t + \frac{1}{2}\frac{\partial^2 u}{\partial x^2}(X_t, t)\,b^2(X_t, t)\,dt. \tag{2.4}$$

The key difference from the classical chain rule lies in the presence of the second-order term. In stochastic calculus, we cannot ignore terms like dX_t^2. In fact, using the identity $dW_t^2 = dt$ and the Itô product rule, the term dX_t^2 gives rise to a non-zero contribution, specifically $b^2(X_t, t)\,dt$. This correction term is fundamental and is what allows Itô calculus to properly describe stochastic dynamics.

2.2.2 Application of Itô's Lemma to Geometric Brownian Motion

We now apply Itô's Lemma to a function $C = C(S, t)$, which represents the price of a derivative (e.g., an option) that depends on both the stock price S_t and time t. The function $C(S, t)$ is assumed to be continuously differentiable in t and twice differentiable in S.

Suppose that the stock price S_t follows a Geometric Brownian Motion described by the stochastic differential equation:

$$dS_t = \mu S_t\,dt + \sigma S_t\,dW_t, \tag{2.5}$$

where μ is the drift, σ is the volatility, and W_t is a standard Brownian motion.

We are now interested in finding the differential of the function $C(S_t, t)$ using the general form of Itô's Lemma:

$$du = \frac{\partial u}{\partial t}\, dt + \frac{\partial u}{\partial x}\, dX_t + \frac{1}{2}\frac{\partial^2 u}{\partial x^2} b^2(X_t, t)\, dt. \tag{2.6}$$

We identify the following correspondence:

- $X_t = S_t$,
- $u(S_t, t) = C(S_t, t)$,
- $a(S_t, t) = \mu S_t$,
- $b(S_t, t) = \sigma S_t$.

Substituting $X_t = S_t$ and $u = C$ into eq. (2.6), we obtain:

$$dC = \frac{\partial C}{\partial t}\, dt + \frac{\partial C}{\partial S}\, dS_t + \frac{1}{2}\frac{\partial^2 C}{\partial S^2}\sigma^2 S_t^2\, dt.$$

Now, replace dS_t from eq. (2.5):

$$\frac{\partial C}{\partial S}\, dS_t = \frac{\partial C}{\partial S}\left(\mu S_t\, dt + \sigma S_t\, dW_t\right) = \mu S_t \frac{\partial C}{\partial S} dt + \sigma S_t \frac{\partial C}{\partial S} dW_t$$

Substituting back into the expression for dC, we obtain:

$$dC = \frac{\partial C}{\partial t} dt + \mu S_t \frac{\partial C}{\partial S} dt + \sigma S_t \frac{\partial C}{\partial S} dW_t + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 C}{\partial S^2} dt$$

Finally, grouping all the terms involving dt, we arrive at the full expression:

$$dC = \left(\frac{\partial C}{\partial t} + \mu S_t \frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 C}{\partial S^2}\right) dt + \sigma S_t \frac{\partial C}{\partial S} dW_t \tag{2.7}$$

This is the stochastic differential equation that describes the infinitesimal evolution of the option price $C(S_t, t)$ as a function of the underlying asset price and time.

2.2.3 The Black-Scholes PDE

Now we are able to derive the Black-Scholes PDE, from our SDE (2.7). We now consider a self-financing portfolio strategy that replicates the value of the option. Let x_t be the number of units held in the risk-free asset and y_t the number of units held in the underlying stock. The value of this portfolio at time t is:

$$\pi_t = x_t B_t + y_t S_t, \tag{2.8}$$

where B_t is the value of the risk-free asset, which evolves deterministically as: $dB_t = rB_t dt$ We choose x_t and y_t such that the portfolio exactly replicates the option value at all times, i.e., $\pi_t = C(S_t, t)$ for all t. Applying Itô's Lemma to the portfolio, we differentiate π_t:

$$\begin{aligned} d\pi_t &= x_t dB_t + y_t dS_t \\ &= x_t r B_t dt + y_t \left(\mu S_t dt + \sigma S_t dW_t\right) \\ &= \left(r x_t B_t + y_t \mu S_t\right) dt + y_t \sigma S_t dW_t. \end{aligned} \tag{2.9}$$

Since the portfolio replicates the option, it must follow the same stochastic dynamics. Thus, comparing the drift and diffusion terms in eq. (2.9) and eq. (2.7), we match the coefficients:

$$y_t = \frac{\partial C}{\partial S}, \tag{2.10}$$

$$r x_t B_t = \frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 C}{\partial S^2}. \tag{2.11}$$

We now substitute eq. (2.10) and eq. (2.11) into the identity $C_t = \pi_t = x_t B_t + y_t S_t$ and derive the partial differential equation satisfied by C. Multiply eq. (2.11) by 1 (i.e., isolate $r x_t B_t$) and rewrite the identity:

$$rC = r x_t B_t + r y_t S_t = \left(\frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2}\right) + rS\frac{\partial C}{\partial S}$$

Rearranging terms gives the celebrated Black-Scholes Partial Differential Equation (PDE):

$$\frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + rS\frac{\partial C}{\partial S} - rC = 0 \tag{2.12}$$

This equation governs the evolution of the option price under the assumption of no arbitrage and in a complete, frictionless market with a constant risk-free rate.

2.2.4 Black-Scholes formula

The Black-Scholes PDE (2.12) describes the evolution of the option price $C(S,t)$. For this reason, we have to introduce the following boundary conditions for a European call option:

- Terminal condition: $C(S,T) = \max(S-K, 0)$
- Boundary conditions: $C(0,t) = 0$ and $C(S,t) \sim S$ as $S \to \infty$

To directly solve the PDE (2.12), the well-known closed-form solution is:

$$C(S,t) = S\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2) \tag{2.13}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution, and

$$d_1 = \frac{\log\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}},$$
$$d_2 = d_1 - \sigma\sqrt{T-t}.$$

To verify that eq. (2.13) indeed solves the PDE, we can substitute eq. (2.13) into eq. (2.12) and check analytically that the PDE is satisfied. The intuitive reasoning behind the solution, and thus its suitability for finance applications, arises from the concept of risk-neutral valuation. Specifically, the PDE obtained assumes a risk-free growth rate r, independent of investors' risk preferences. This approach is

the cornerstone of modern derivatives pricing, known as risk-neutral pricing. Thus, the Black-Scholes formula eq. (2.13) provides the analytic price for a European call option under the assumptions of the Black-Scholes model, aligning clearly with the practical financial reasoning used in markets.

2.3 Black-Scholes-Merton Model

In the original Black-Scholes framework, one of the simplifying assumptions is that the underlying asset pays no dividends. However, many real-world stocks distribute dividends, which directly affect the option's value. Thus, the Black-Scholes model can be extended to include a continuous dividend yield q, assuming that dividends are paid at a constant proportional rate over time. This extension, commonly referred to as the Black-Scholes-Merton model, was formally introduced by Merton (1973) [12]. It adjusts the risk-neutral valuation framework by incorporating the continuous dividend yield, leading to modified option pricing formulas that better reflect the reduced expected growth of the underlying asset.

2.3.1 Adjusted Geometric Brownian Motion

When a stock pays a continuous dividend yield q, the stock price dynamics under the risk-neutral measure are modified as follows:

$$dS_t = (r - q)S_t\,dt + \sigma S_t\,dW_t,$$

where q is the continuous dividend yield.

The key change here is that the expected growth rate of the stock under the risk-neutral measure becomes $r - q$, reflecting the fact that a portion of the stock's return is paid out as dividends and does not contribute to capital gains.

2.3.2 Adjusted Black-Scholes PDE

Following the same derivation steps as in the dividend-free case, we obtain the adjusted Black-Scholes Partial Differential Equation:

$$\frac{\partial C}{\partial t} + (r-q)S\frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} - rC = 0, \tag{2.14}$$

where $C(S,t)$ is the price of the option at time t when the underlying asset price is S.

2.3.3 Adjusted Black-Scholes Formula

For a European call option with continuous dividend yield q, the Black-Scholes formula becomes:

$$C(S,t) = Se^{-q(T-t)}\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2), \tag{2.15}$$

where:

$$d_1 = \frac{\log\frac{S}{K} + (r - q + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}},$$
$$d_2 = d_1 - \sigma\sqrt{T-t},$$

and $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution.

The key modification introduced by q is that the present value of the expected payoff is reduced to account for the continuous outflow of dividends. This impacts both the drift term in the underlying asset's stochastic process and the discounted value of the stock in the pricing formula. For put options, a similar adjustment applies, ensuring the model remains arbitrage-free under dividend-paying conditions.

2.4 Limitations

Despite its groundbreaking contribution to financial theory, the Black-Scholes model presents several notable limitations that restrict its accuracy in real-world applications. One of the most significant is the assumption of constant volatility: the model treats the underlying asset's volatility as fixed over time, whereas empirical evidence shows that market volatility fluctuates, often systematically, in ways that produce patterns like the *volatility smile* or *volatility skew*, that we will present in the next chapter. Additionally, the Black-Scholes framework assumes that asset prices follow a continuous log-normal process, implying normally distributed returns.

Another key limitation is the absence of price jumps or discontinuities: the model assumes that price paths evolve smoothly, ignoring the possibility of sudden jumps due to unexpected news, earnings announcements, or macroeconomic shocks. Furthermore, Black-Scholes presumes a world with no transaction costs, no bid-ask spreads, and no restrictions on borrowing or short selling, idealized conditions that rarely hold in practice. Finally, the model was originally developed for European options, which can only be exercised at maturity, whereas many traded options are American and can be exercised at any time.

These limitations have spurred the development of more advanced models that attempt to better capture the complexities of financial markets.

Chapter 3

Volatility Smile and Skew

3.1 Implied Volatility

The concept of implied volatility plays a central role in this thesis.

Implied volatility is defined as the value of the volatility parameter σ that, as we saw, when input into the Black-Scholes pricing formula, yields a model option price equal to the observed market price.

In mathematical terms, given the observed price C_{market} of a European call option, we solve the following equation for σ:

$$C_{\text{Black-Scholes}}(S, K, r, T, \sigma) = C_{\text{market}}, \tag{3.1}$$

where S is the current underlying price, K is the strike price, r is the risk-free rate, T is the time to maturity, and σ is the volatility.

This approach reflects the market's collective expectations about the future variability of the underlying asset, as implied by the option's price.

In the original Black-Scholes framework, the assumption is that volatility is constant over the life of the option, independent of strike or maturity. This leads to the theoretical expectation that implied volatility, when plotted across different

strike prices and maturities, should form a flat surface, meaning the same value of σ applies to all options on the same underlying, regardless of their moneyness or time to expiration.

However, when practitioners compute implied volatilities from real market data, they often observe systematic patterns and deviations from this flat surface. These empirical phenomena, known as the volatility smile and the volatility skew, reveal that the assumption of constant volatility is too simplistic to capture the complexities of financial markets.

In the following sections, we will explore these empirical patterns and highlight why they motivate the development of more advanced models.

3.2 Volatility Smile and Skew

One of the most well-known empirical patterns observed in options markets is the volatility smile. When implied volatilities are plotted against strike prices for options with the same maturity, practitioners often observe a curve that resembles a smile: implied volatilities are higher for deep in-the-money and deep out-of-the-money options, and lower for at-the-money options.

The reason behind the volatility smile lies in the market's recognition that extreme movements (large upward or downward price jumps) are more likely than predicted by the Black-Scholes model. Specifically, empirical distributions of asset returns tend to exhibit fat tails, meaning that large deviations from the mean occur more frequently than under the normal distribution assumed by Black-Scholes. Investors, aware of this, are willing to pay a premium for options that protect against extreme market moves, leading to higher implied volatilities for strikes far away from the current price. Furthermore, the implied volatility curve is not always symmetric: many markets display an asymmetric pattern known as the volatility

skew or volatility smirk. For example, in equity and index option markets, implied volatilities tend to increase as strike prices decrease, meaning that out-of-the-money puts are typically priced with higher implied volatilities than out-of-the-money calls.

This asymmetry is largely attributed to the market's anticipation of downside risk: large negative moves tend to be more abrupt and severe than large positive moves, so investors are willing to pay a premium for put options that hedge against such events.

Figure 3.1 provides a synthetic example of a volatility surface where both the smile and skew effects are visible, highlighting how implied volatilities vary with strike price and time to expiry.

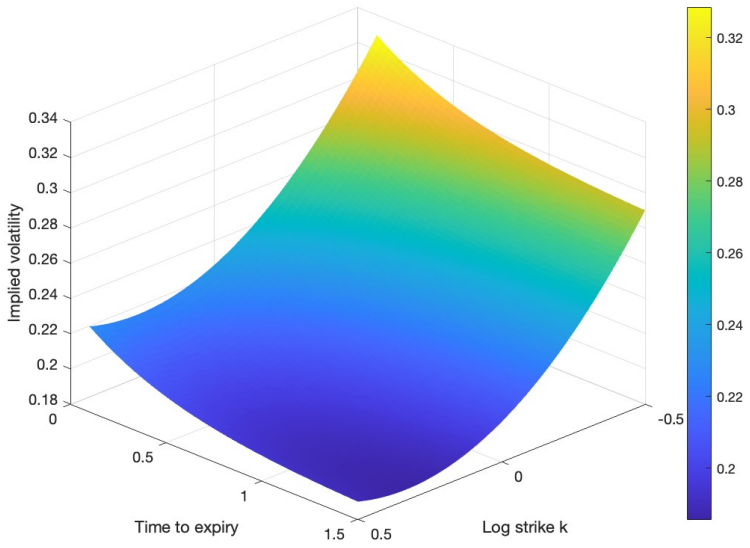

Fig. 3.1: Synthetic volatility surface with base level, skew, and smile components.
Source: own elaboration with Matlab

3.2.1 Dupire's Formula

The local volatility model, introduced by Bruno Dupire [7], extends the Black-Scholes framework by allowing volatility to vary deterministically with both time and the current level of the underlying. Under the risk-neutral measure, the asset price S_t follows the stochastic differential equation:

$$dS_t = (r - q)S_t\,dt + \sigma_{\text{loc}}(t, S_t)S_t\,dW_t, \tag{3.2}$$

where:

- r is the constant risk-free rate,
- q is the continuous dividend yield,
- W_t is a standard Brownian motion under the risk-neutral measure,
- $\sigma_{\text{loc}}(t, S_t)$ is the local volatility function.

To connect market-observed European call prices to the local volatility function, we derive the celebrated Dupire formula. The starting point is the Fokker-Planck equation, which governs the evolution of the density $p(T, S)$ of S_T:

$$\frac{\partial p}{\partial T} = -\frac{\partial}{\partial S}\Big[(r - q)Sp\Big] + \frac{1}{2}\frac{\partial^2}{\partial S^2}\Big[\sigma^2_{\text{loc}}(T, S)S^2 p\Big]. \tag{3.3}$$

The value of a European call option with strike K and maturity T is given by:

$$C(T, K) = e^{-rT}\int_K^\infty (S - K)p(T, S)\,dS. \tag{3.4}$$

We now compute the derivatives of $C(T, K)$:

$$\frac{\partial C}{\partial K} = -e^{-rT}\int_K^\infty p(T, S)\,dS, \tag{3.5}$$

$$\frac{\partial^2 C}{\partial K^2} = e^{-rT}p(T, K). \tag{3.6}$$

Differentiating eq. (3.4) with respect to T:

$$\frac{\partial C}{\partial T} = -re^{-rT}\int_K^\infty (S-K)p(T,S)\,dS + e^{-rT}\int_K^\infty (S-K)\frac{\partial p}{\partial T}(T,S)\,dS.$$

We now substitute the Fokker-Planck eq. (3.3) into the second integral:

$$\frac{\partial C}{\partial T} = -rC(T,K) + e^{-rT}\int_K^\infty (S-K)\left[-\frac{\partial}{\partial S}((r-q)Sp) + \frac{1}{2}\frac{\partial^2}{\partial S^2}(\sigma^2 S^2 p)\right] dS.$$

We apply integration by parts twice. For the drift term:

$$\int_K^\infty (S-K)\left(-\frac{\partial}{\partial S}[(r-q)Sp]\right) dS = (r-q)\int_K^\infty (S-K)p(T,S)\,dS.$$

For the diffusion term:

$$\int_K^\infty (S-K)\left(\frac{1}{2}\frac{\partial^2}{\partial S^2}[\sigma^2 S^2 p]\right) dS = \frac{1}{2}\sigma^2(T,K)K^2 p(T,K).$$

Putting all together:

$$\frac{\partial C}{\partial T} = -rC(T,K) + e^{-rT}\left[(r-q)\int_K^\infty (S-K)p(T,S)\,dS + \frac{1}{2}\sigma^2(T,K)K^2 p(T,K)\right]. \tag{3.7}$$

From the eq.3.6 we can write:

$$p(T,K) = e^{rT}\frac{\partial^2 C}{\partial K^2}. \tag{3.8}$$

Furthermore, taking the derivative of $C(T,K)$ with respect to K, eq.(3.5) Multiplying both sides by K:

$$K\frac{\partial C}{\partial K} = -Ke^{-rT}\int_K^\infty p(T,S)\,dS. \tag{3.9}$$

Now, summing the original expression $C(T,K)$ and $K\frac{\partial C}{\partial K}$:

$$C(T,K) + K\frac{\partial C}{\partial K} = e^{-rT}\left[\int_K^\infty (S-K)\,p(T,S)\,dS - K\int_K^\infty p(T,S)\,dS\right]. \tag{3.10}$$

$$= e^{-rT}\left[\int_K^\infty S\,p(T,S)\,dS - K\int_K^\infty p(T,S)\,dS\right] = e^{-rT}\int_K^\infty (S-K)\,p(T,S)\,dS. \tag{3.11}$$

Thus, we recover the identity:

$$\int_K^\infty (S-K)p(T,S)\,dS = e^{rT}\left[C(T,K) + K\frac{\partial C}{\partial K}\right], \tag{3.12}$$

so, taken this identity with eq.(3.8) an put inside eq.(3.7),

$$\frac{\partial C}{\partial T} = -rC(T,K) + (r-q)\left[C(T,K) + K\frac{\partial C}{\partial K}\right] + \frac{1}{2}\sigma^2(T,K)K^2\frac{\partial^2 C}{\partial K^2}.$$

Which simplifies to:

$$\frac{\partial C}{\partial T} + qK\frac{\partial C}{\partial K} + qC = \frac{1}{2}\sigma^2_{\text{loc}}(T,K)K^2\frac{\partial^2 C}{\partial K^2}. \tag{3.13}$$

Solving for the local volatility yields the Dupire formula:

$$\sigma^2_{\text{loc}}(T,K) = \frac{\frac{\partial C}{\partial T} + qK\frac{\partial C}{\partial K} + qC}{\frac{1}{2}K^2\frac{\partial^2 C}{\partial K^2}}. \tag{3.14}$$

Once the local volatility surface $\sigma_{\text{local}}(T,K)$ is computed from observed market option prices via Dupire's formula, it must be transformed into a function of time and the spot price, $\sigma_{\text{local}}(t,S_t)$, for use in the asset price dynamics. This transformation typically involves either interpolation techniques or a numerical inversion of the pricing PDE.

Once $\sigma_{\text{local}}(t,S)$ is available, the pricing of European options no longer relies on closed-form solutions, as in the Black-Scholes model. Instead, two primary numerical methods are commonly employed:

- The PDE method: One can solve the forward pricing partial differential equation using finite difference methods[1]. The PDE for the option price $C(t,S)$ becomes:

  $$\frac{\partial C}{\partial t} + (r-q)S\frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2_{\text{local}}(t,S)S^2\frac{\partial^2 C}{\partial S^2} - rC = 0, \tag{3.15}$$

  how we can see this is the eq.(2.14) but modified for the local volatility.

[1]Finite difference methods are numerical techniques for solving differential equations by approximating derivatives with differences over a discrete grid in time and space. This transforms the PDE into a system of algebraic equations that can be solved iteratively.

- Monte Carlo simulation: Alternatively, the dynamics of the underlying asset can be simulated using the SDE(3.2) and option prices are computed as discounted expectations of the payoff with a Monte Carlo Simulation[2].

These numerical approaches allow practitioners to evaluate option prices even in the absence of analytical pricing formulas.

Despite its ability to fit the implied volatility surface, the local volatility model presents notable limitations. Its deterministic nature prevents it from capturing the stochastic evolution of volatility observed in the markets, leading to unrealistic smile dynamics. Another important drawback is the instability in the calibration process: small perturbations in market data can lead to large variations in the estimated local volatility surface. So, these issues motivate the adoption of more sofisticated models.

[2]See appendix A.1

Chapter 4

Heston Model

The Heston model was developed to overcome one of the key limitations of the Black-Scholes framework: the assumption of constant volatility. While the Black-Scholes model assumes, as we said, that the volatility of the underlying asset remains fixed over the life of the option, empirical evidence shows that in real markets, volatility fluctuates over time and exhibits systematic patterns such as the volatility smile and volatility skew that we have shown in the last chapter.

The central idea behind the Heston model is to treat volatility itself as a stochastic process, allowing it to evolve randomly over time alongside the asset price. By doing so, the model captures richer market dynamics and provides a more realistic description of option prices.

4.1 Heston Model Derivation

4.1.1 Feller Square-Root Process

The FSR process (1954) is a stochastic process used to model quantities that must remain non-negative, such as interest rates or variances. It was introduced

in finance by Cox Ingersoll and Ross in 1985 to model the short interest rate. It is defined as solution of the following stochastic differential equation

$$dv_t = \kappa(\overline{v} - v_t)dt + \sigma\sqrt{v_t}dW_t \tag{4.1}$$

where:

- κ is the rate at which the process reverts to its long-term mean $\overline{v}$.
- $\overline{v}$ is the long-term average level of the process.
- σ is the volatility of the process (sometimes called "volatility of volatility" or vol of vol although it is actually the volatility of variance).
- W_t is a standard Brownian motion.

The FSR process guarantees the positiveness of v_t, provided the Feller condition[1] holds:

$$2\kappa\overline{v} \geq \sigma^2. \tag{4.2}$$

This condition ensures that the stochastic variance does not become negative, which is crucial when modeling financial variables like volatility.

4.1.2 Heston Model Dynamics

Built on the FSR process, the Heston model describes the joint evolution of the asset price S_t and its variance v_t through a system of two stochastic differential

[1]This condition arises because the FSR process is a square-root diffusion, and from stochastic process theory we have that if $2\kappa\overline{v} \geq \sigma^2$, the mean reversion force dominates random shocks, making the zero boundary inaccessible and ensuring variance remains strictly positive.

equations,

$$dS_t = rS_t dt + \sqrt{v_t} S_t dW_{1,t}, \tag{4.3}$$

$$dv_t = \kappa(\overline{v} - v_t)dt + \sigma\sqrt{v_t} dW_{2,t}, \tag{4.4}$$

$$E[dW_{1,t} dW_{2,t}] = \rho dt. \tag{4.5}$$

Here:

- r: risk-free interest rate.
- v_t: instantaneous variance at time t.
- ρ: correlation between the Brownian motions $W_{1,t}$ (driving the asset) and $W_{2,t}$ (driving the variance).

As we can see, the asset price evolves similarly to Black-Scholes but now has a stochastic variance, while the variance itself follows a mean-reverting FSR process. The correlation ρ between asset and variance innovations allows the model to capture asymmetries and leverage effects observed in real markets.

4.1.3 Itô's Lemma

In the Black-Scholes framework, the option price $C(S_t, t)$ depends only on two variables: the asset price S_t and time t. Applying Itô's Lemma, we used the one-dimensional form:

$$du = \frac{\partial u}{\partial t} dt + \frac{\partial u}{\partial x} dX_t + \frac{1}{2} \frac{\partial^2 u}{\partial x^2} b^2(X_t, t)\, dt, \tag{4.6}$$

where $u(x,t) = C(S_t, t)$ and $X_t = S_t$. In the Heston model, the option price $C(S_t, v_t, t)$ depends on three variables: time t, asset price S_t, and variance v_t. We therefore apply Itô's Lemma for functions of two stochastic variables:

$$du = \frac{\partial u}{\partial t} dt + \frac{\partial u}{\partial x} dX_t + \frac{\partial u}{\partial y} dY_t + \frac{1}{2} \frac{\partial^2 u}{\partial x^2} (dX_t)^2 + \frac{1}{2} \frac{\partial^2 u}{\partial y^2} (dY_t)^2 + \frac{\partial^2 u}{\partial x \partial y} dX_t\, dY_t, \tag{4.7}$$

where $u(x, y, t) = C(S_t, v_t, t)$, $X_t = S_t$, $Y_t = v_t$. We now substitute the Heston dynamics:

$$dS_t = rS_t dt + \sqrt{v_t} S_t dW_{1,t}, \tag{4.8}$$

$$dv_t = \kappa(\overline{v} - v_t)dt + \sigma\sqrt{v_t} dW_{2,t}, \tag{4.9}$$

$$E[dW_{1,t} dW_{2,t}] = \rho dt, \tag{4.10}$$

We compute the quadratic terms:

$$(dS_t)^2 = v_t S_t^2 \, dt, \tag{4.11}$$

$$(dv_t)^2 = \sigma^2 v_t \, dt, \tag{4.12}$$

$$dS_t \, dv_t = \rho\sigma v_t S_t \, dt. \tag{4.13}$$

Therefore, applying Itô's Lemma gives:

$$dC = \frac{\partial C}{\partial t} dt + \frac{\partial C}{\partial S} dS_t + \frac{\partial C}{\partial v} dv_t + \frac{1}{2} \frac{\partial^2 C}{\partial S^2} v_t S_t^2 \, dt + \frac{1}{2} \frac{\partial^2 C}{\partial v^2} \sigma^2 v_t \, dt + \frac{\partial^2 C}{\partial S \partial v} \rho\sigma v_t S_t \, dt. \tag{4.14}$$

Finally, grouping all the terms involving dt, we arrive at the full stochastic differential equation describing the infinitesimal evolution of the option price under the Heston model:

$$dC = \left(\frac{\partial C}{\partial t} + rS_t \frac{\partial C}{\partial S} + \kappa(\overline{v} - v_t) \frac{\partial C}{\partial v} + \frac{1}{2} v_t S_t^2 \frac{\partial^2 C}{\partial S^2} + \frac{1}{2} \sigma^2 v_t \frac{\partial^2 C}{\partial v^2} + \rho\sigma v_t S_t \frac{\partial^2 C}{\partial S \partial v} \right) dt + \sqrt{v_t} S_t \frac{\partial C}{\partial S} dW_{1,t} + \sigma\sqrt{v_t} \frac{\partial C}{\partial v} dW_{2,t}. \tag{4.15}$$

4.1.4 From the SDE to the Heston PDE

Starting from the stochastic differential eq. (4.15), we construct a self-financing replicating portfolio:

$$\pi_t = x_t B_t + y_t S_t + z_t v_t, \tag{4.16}$$

where x_t is the position in the risk-free asset B_t, y_t is the position in the risky asset S_t, and z_t is the position in the variance process v_t. The dynamics of the portfolio is

$$d\pi_t = x_t r B_t dt + y_t(rS_t dt + \sqrt{v_t} S_t dW_{1,t}) + z_t(\kappa(\overline{v} - v_t)dt + \sigma\sqrt{v_t} dW_{2,t}). \quad (4.17)$$

Since the portfolio replicates the option, we require

$$\pi_t = C(S_t, v_t, t). \quad (4.18)$$

To eliminate the stochastic terms, we match

$$y_t = \frac{\partial C}{\partial S}, \quad z_t = \frac{\partial C}{\partial v}. \quad (4.19)$$

Focusing now on the deterministic (drift) part, we have

$$d\pi_t = \left(x_t r B_t + r S_t \frac{\partial C}{\partial S} + \kappa(\overline{v} - v_t)\frac{\partial C}{\partial v} \right) dt. \quad (4.20)$$

From the replicating condition, we solve for $x_t r B_t$

$$x_t r B_t = rC - r S_t \frac{\partial C}{\partial S} - \kappa(\overline{v} - v_t)\frac{\partial C}{\partial v}. \quad (4.21)$$

Therefore, the total drift becomes

$$d\pi_t = \left(rC - rS_t \frac{\partial C}{\partial S} - \kappa(\overline{v} - v_t)\frac{\partial C}{\partial v} + rS_t \frac{\partial C}{\partial S} + \kappa(\overline{v} - v_t)\frac{\partial C}{\partial v} \right. \quad (4.22)$$

$$\left. + \frac{\partial C}{\partial t} + \frac{1}{2} v_t S_t^2 \frac{\partial^2 C}{\partial S^2} + \frac{1}{2}\sigma^2 v_t \frac{\partial^2 C}{\partial v^2} + \rho\sigma v_t S_t \frac{\partial^2 C}{\partial S \partial v} \right) dt. \quad (4.23)$$

Semplifying:

$$d\pi_t = \left(rC + \frac{\partial C}{\partial t} + \frac{1}{2} v_t S_t^2 \frac{\partial^2 C}{\partial S^2} + \frac{1}{2}\sigma^2 v_t \frac{\partial^2 C}{\partial v^2} + \rho\sigma v_t S_t \frac{\partial^2 C}{\partial S \partial v} \right) dt. \quad (4.24)$$

So, now we can say that after simplifying the portfolio drift, we are left with a purely deterministic (riskless) component. As a result, the portfolio π_t behaves like

a risk-free asset, and under the risk-neutral measure, it must grow at the risk-free rate r. Therefore, we impose the no-arbitrage condition

$$d\pi_t = rC\,dt. \tag{4.25}$$

Equating the drift expression derived from the SDE (4.15) with the no-arbitrage condition eq. (4.25), and rearranging, we obtain the Heston PDE (4.26),

$$\frac{\partial C}{\partial t} + rS_t\frac{\partial C}{\partial S} + \kappa(\overline{v} - v_t)\frac{\partial C}{\partial v} + \frac{1}{2}v_t S_t^2\frac{\partial^2 C}{\partial S^2} + \frac{1}{2}\sigma^2 v_t\frac{\partial^2 C}{\partial v^2} + \rho\sigma v_t S_t\frac{\partial^2 C}{\partial S\partial v} - rC = 0. \tag{4.26}$$

4.1.5 Semi-Analytical Solution

Although the Heston PDE (4.26), formally governs the option price dynamics, solving this equation numerically can be computationally intensive. Fortunately, Heston (1993) [9] proposed a semi-analytical solution for European options that avoids solving the PDE directly. The idea is to express the European call option price under the risk-neutral measure as

$$C(S, v, t) = SP_1 - Ke^{-r(T-t)}P_2, \tag{4.27}$$

where P_1 and P_2 are risk-neutral probabilities defined as:

- P_1: the probability that the option finishes in-the-money under the Q_1 measure, linked to the asset's expected value discounted at the risk-free rate.
- P_2: the probability that the option finishes in-the-money under the Q_2 measure, related to the option's discounted strike.

These probabilities are computed using the characteristic function $\phi(u)$ of the log-asset price,

$$P_j = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \mathrm{Re}\left[\frac{e^{-iu\log K}\phi_j(u)}{iu}\right]du, \quad j = 1, 2, \tag{4.28}$$

where $\phi_j(u)$ is the characteristic function under each probability measure Q_j.

Specifically, the characteristic function under the Heston model takes the form

$$\phi(u;t,S,v) = e^{C(u,\tau)+D(u,\tau)v+iu\log S}, \tag{4.29}$$

where $\tau = T-t$ is the time to maturity, and the functions $C(u,\tau)$ and $D(u,\tau)$ are

$$C(u,\tau) = riu\tau + \frac{\kappa\overline{v}}{\sigma^2}\left[(b-d)\tau - 2\log\frac{1-ge^{-d\tau}}{1-g}\right], \tag{4.30}$$

$$D(u,\tau) = \frac{b-d}{\sigma^2}\left(\frac{1-e^{-d\tau}}{1-ge^{-d\tau}}\right), \tag{4.31}$$

with

$$d = \sqrt{(\rho\sigma iu - b)^2 - \sigma^2(2\alpha iu - u^2)}, \tag{4.32}$$

$$g = \frac{b-d}{b+d}, \tag{4.33}$$

$$b = \kappa + \lambda - \rho\sigma iu, \tag{4.34}$$

$$\alpha = -\frac{u^2}{2} - \frac{iu}{2}, \tag{4.35}$$

where λ is the market price of volatility risk, typically set to zero under risk-neutral valuation.

This semi-analytical solution provides a foundation for pricing options close to the observed market option prices.

Chapter 5

Calibration

As presented in the previous chapter, the Heston model extends the Black-Scholes framework by allowing volatility to evolve stochastically. This is achieved by introducing five parameters that are not directly observable from market data. Unlike the Black-Scholes model, where the volatility can be inferred from a single market price, the Heston model requires a more involved procedure to estimate all five parameters simultaneously. This leads to what is known as "calibration problem".

Calibration refers to the procedure of determining the model parameters that allow the model to reproduce observed market prices as closely as possible. Calibration becomes the key tool for inferring them from data.

In mathematical terms, the calibration problem can be formulated as an optimization task, where one aims to minimize the discrepancy between the model output and the market data.

Proposition 5.0.1. *Let V_{market} denote the observed market prices and $V_{model}(\boldsymbol{\theta})$ the corresponding model prices for the parameter set $\boldsymbol{\theta}$. The calibration problem is thus explain as*

$$\min_{\boldsymbol{\theta}\in\mathcal{P}} \; f\left(V_{model}(\boldsymbol{\theta}), V_{market}\right), \tag{5.1}$$

where f is a loss function measuring the error between model and market prices, and $\mathcal{P}$ is the domain of admissible parameter values (typically bounded by financial or mathematical constraints, as we can see later).

5.1 General Calibration Strategies

The solution of a calibration problem depends on several modelling and numerical choices. In this section, we describe the main components that define a calibration strategy, independently of the specific model under consideration.

The performance and stability of the calibration process are strongly influenced by some definitions, that are:

- The choice of the loss function;
- The numerical method adopted to solve the optimization problem;
- The use of constraints or penalty terms to ensure arbitrage-free and stable solutions.

The goal is to provide a general point of view, which will be applied in a specific way to the Heston model.

5.1.1 Objective Functions

As we said the definition of the objective function, that we called before "loss function", is a central component of any calibration because that quantifies the difference between model outputs and market observations as described in eq. (5.1). Beyond this, the specific form of the loss function has a significant impact on the behaviour of the calibration algorithm. In fact, it directly shapes the optimization landscape and thus influences three key aspects of the calibration process:

- Convergence: the smoothness and convexity of the objective function affect how easily the optimization algorithm can approach a minimum. For example, a poorly scaled[1] loss function may create flat regions or multiple local minima, making it harder to converge.

- Accuracy: the weighting and normalization applied inside the loss function determine which parts of the data are fitted more precisely. Using a relative or vega-weighted loss can improve accuracy in regions that are more relevant from a pricing or hedging perspective (e.g., at-the-money options).

- Robustness: a well-constructed loss function avoids overfitting to noisy data and produces stable parameter estimates. Incorporating regularization or normalizing by market prices helps reduce the sensitivity to small perturbations or data irregularities.

In general, the most common loss functions used in option pricing calibration are based on: Prices, volatilities and vega weighted loss.

Let $V_{\text{market},i}$ and $V_{\text{model},i}(\boldsymbol{\theta})$ be the market and model prices of option i, and let N[2] be the number of input data points used in calibration. A classical price-based objective function is the mean square error (MSE)

$$f_{\text{MSE}}(\boldsymbol{\theta}) = \frac{1}{N}\sum_{n=1}^{N}\left(V_{\text{model},n}(\boldsymbol{\theta}) - V_{\text{market},n}\right)^2. \tag{5.2}$$

Alternatively, one can use the relative mean square error (RMSE), normalized by market price,

$$f_{\text{rel}}(\boldsymbol{\theta}) = \frac{1}{N}\sum_{n=1}^{N}\left(\frac{V_{\text{model},n}(\boldsymbol{\theta}) - V_{\text{market},n}}{V_{\text{market},n}}\right)^2. \tag{5.3}$$

[1] A poorly scaled loss function may lead to numerical instability during optimization. This occurs when the errors associated with different options have vastly different magnitudes, typically because option prices vary widely across strikes and maturities.

[2] For example, if 10 maturities and 17 delta values are used to construct the implied volatility surface, then $N = 10 \times 17 = 170$.

Indeed, the relative error is often preferred over absolute differences because it naturally normalizes the pricing error with respect to the scale of each option. This allows for a more balanced treatment of deep in-the-money, at-the-money, and out-of-the-money options, avoiding the overweighting of options with high nominal prices. As discussed in [14] and [11], using relative errors leads to a more stable and robust calibration, especially when the price range is heterogeneous[3]. About Implied volatility-based loss functions, these are often preferred because volatilities are less sensitive to discounting, this is because option prices are affected by the discounting of future payoffs, and thus depend heavily on the interest rate and time to maturity. Implied volatilities, being derived from prices through inversion, tend to be less sensitive to such effects, making them more stable and easier to compare across different maturities. So, let $\sigma_{\text{model},n}(\boldsymbol{\theta})$ and $\sigma_{\text{market},n}$ be the model and the market implied volatility values of option n, thus

$$f_\sigma(\boldsymbol{\theta}) = \frac{1}{N}\sum_{n=1}^{N}\left(\sigma_{\text{model},n}(\boldsymbol{\theta}) - \sigma_{\text{market},n}\right)^2. \tag{5.4}$$

The choice between these alternatives depends on data availability and model structure. Price-based loss functions are more robust when working with deep ITM/OTM options [14, 11], whereas volatility-based loss functions better reflect market quoting conventions[4].

In some cases, the loss function is weighted by the sensitivity of each option to volatility, a quantity known as vega.

Proposition 5.1.1. *The vega $\mathcal{V}_n$ of an option measures how much the option price*

[3]A heterogeneous price range refers to the situation where option prices vary significantly across strikes and maturities. For example, deep in-the-money options may have prices close to the underlying asset, while deep out-of-the-money options may be worth only a few cents.

[4]Volatility is often the preferred quoting metric in derivative markets. Market participants typically refer to the implied volatility surface rather than to the raw option prices when quoting, trading, and calibrating.

changes in response to a small change in implied volatility, i.e.,

$$\mathcal{V}_n = \frac{\partial V_n}{\partial \sigma}. \tag{5.5}$$

It tends to be largest for options that are at-the-money (ATM) and with short to medium maturities, and smallest for deep in-the-money or out-of-the-money options. A common formulation is the vega-weighted loss

$$f_{\mathcal{V}}(\boldsymbol{\theta}) = \sum_{n=1}^{N} \left(\frac{V_{\text{model},n}(\boldsymbol{\theta}) - V_{\text{market},n}}{\mathcal{V}_n} \right)^2, \tag{5.6}$$

where $\mathcal{V}_n$ is the vega of option n.

This approach improves the calibration quality by emphasizing options with high sensitivity to volatility, such as ATM options, and down weighting those with low vega, which are often less liquid or more prone to pricing noise.

5.1.2 Optimization Algorithms

The calibration of financial models is typically formulated as a nonlinear constrained optimization problem. Solving such problems requires the use of optimization algorithms numerical procedures that explore the parameter space to minimize a given objective function.

A central concept in optimization is the distinction between local and global minima.

Proposition 5.1.2. *A parameter set $\boldsymbol{\theta}^*$ is called a local minimum if there exists a neighborhood around $\boldsymbol{\theta}^*$ such that*

$$f(\boldsymbol{\theta}^*) \leq f(\boldsymbol{\theta}), \quad \textit{for all } \boldsymbol{\theta} \textit{ in a neighborhood of } \boldsymbol{\theta}^*. \tag{5.7}$$

Proposition 5.1.3. *It is a global minimum if this inequality holds for all $\boldsymbol{\theta}$ in the admissible set $\mathcal{P}$*

$$f(\boldsymbol{\theta}^*) \leq f(\boldsymbol{\theta}), \quad \forall\, \boldsymbol{\theta} \in \mathcal{P}. \tag{5.8}$$

Nonlinear models such as Heston often lead to loss functions with multiple local minima, making the choice of algorithm and initial guess particularly important, because if you do not guess very well the initial parameters, there is the possibility to stumble upon a local minimum.

Among the most commonly used methods in practice are gradient-based algorithms, which rely on information from the first or second derivatives of the loss function.

Gradient Descent This is the simplest iterative method based on first-order information. At each iteration, the parameters are updated along the direction of steepest descent,

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \nabla f(\boldsymbol{\theta}_k), \tag{5.9}$$

where $\eta > 0$ is the step size, also called learning rate. The gradient $\nabla f(\boldsymbol{\theta}_k)$ points in the direction of maximal increase of the loss function, so its opposite reduces it. The convergence depends critically on η: if too large, the method may overshoot; if too small, it may converge very slowly.

Newton's Method While gradient descent only uses first-order information, Newton's method includes curvature by incorporating the second derivative (the Hessian matrix $H(\boldsymbol{\theta})$). The update step becomes

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - H^{-1}(\boldsymbol{\theta}_k) \nabla L(\boldsymbol{\theta}_k). \tag{5.10}$$

This method converges much faster but requires the computation and inversion of the Hessian, which may be computationally intensive and unstable if H is non-invertible.

In summary, both methods aim to descend the loss surface:

- Gradient descent follows the slope, using only direction and a fixed step.

- Newton's method uses both slope and curvature, adapting the step to the geometry of the surface.

5.1.3 Constraints and Regularization

In practical calibration problems, simply minimizing a loss function may lead to solutions that are numerically unstable or economically unrealistic. To address this, one often incorporates constraints and regularization terms into the optimization framework. These mechanisms help ensure that the calibrated parameters remain useful for the purpose.

Constraints are imposed directly on the parameters, to enforce financial or mathematical validity. Typical constraints include:

- Non negativity: Ensuring that parameters such as variances or volatilities remain positive;
- Upper and lower bounds: bounding parameters to lie within a reasonable range derived from empirical studies or market intuition;
- No-arbitrage conditions.

These constraints are typically enforced through the optimizer.

Regularization augments the objective function with penalty terms that discourage certain undesirable behaviors. Formally, the modified loss function becomes

$$f_{\text{reg}}(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta}), \tag{5.11}$$

where $R(\boldsymbol{\theta})$ is a regularization term and $\lambda > 0$ is a penalty weight controlling the trade-off between data fit and regularity. The most common regularizations are:

- L_2 penalty or ridge or Tikhonov, that discourages large parameter values,

$$R(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_2^2 = \sum_{j=1}^{d} \boldsymbol{\theta}_j^2; \tag{5.12}$$

- L_1 penalty or LASSO that encourages sparsity in the parameter vector,

$$R(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_1 = \sum_{j=1}^{d} |\boldsymbol{\theta}_j|; \tag{5.13}$$

- Elastic net, a combination of both L_1 and L_2 penalties to balance both methods,

$$R(\boldsymbol{\theta}) = \alpha\|\boldsymbol{\theta}\|_1 + (1-\alpha)\|\boldsymbol{\theta}\|_2^2; \tag{5.14}$$

- Smoothness penalties that discourage rapid changes in volatility surfaces or other model-implied quantities.

In Appendix A.3 there is the difference between the two types of norms that we used. These regularization techniques are particularly useful when dealing with ill-posed calibration problems, or when market data is noisy. In [8] we saw that introducing regularization can significantly improve the robustness and interpretability of the calibrated model.

From a financial point of view, constraints reflect prior knowledge about the plausible range of parameters, while regularization acts as a safeguard against overfitting.

5.2 Data

For the calibration and empirical analysis of the Heston model, we use real market data obtained from Bloomberg for the stock "BANK OF AMERICA" or [BAC][5] as of 16 May 2025. We choose BAC as the underlying asset due to the high liquidity of its equity options, and a good range of available maturities and strikes, and comprehensive data coverage on Bloomberg. These properties make it an ideal candidate for our calibration. The data includes the implied volatilities,

[5]The market ticker of Bank of America Corporation.

option deltas, strikes, risk-free interest rates, and dividend yields for a range of maturities, in particular it goes from 23 May 2025, so 1 week after the current day, to 16 May 2035, so ten years after the current day.

Bloomberg provides implied volatilities across maturities and deltas, which represent the sensitivity of the option price with respect to the underlying asset. Rather than organizing the volatility surface by strike directly, Bloomberg often quotes implied volatilities as a function of delta and maturity, especially for OTC (over-the-counter) markets. This is because delta provides a more stable and model-independent measure of moneyness across time and market conditions. Unlike strike prices, which are absolute and shift in relevance as the spot price moves, the delta directly reflects the option's sensitivity to the underlying asset and maintains a consistent economic meaning across maturities and volatilities. Therefore, quoting in delta allows market participants to compare options with similar exposure across time. This stability is one of the main reasons why financial data providers, including Bloomberg, prefer delta-based grids for quoting volatility.

In particular, we obtain implied volatilities for both put and call options across a grid of deltas and maturities. For each entry, Bloomberg also reports the corresponding strike price, so that the implied volatility surface can be mapped directly to the strike and maturity space. Volatilities are quoted in annualized percentage terms and were converted to decimal form for calibration. All maturities were also transformed into fractional years to obtain time-to-maturity values τ consistent with continuous-time option pricing models.

Moreover, from Bloomberg we also collect the term structures of the risk-free rate and the dividend yield. These are required inputs for both the Black-Scholes formula and the Heston model pricing equations. For each maturity, we obtain:

- the zero-coupon risk-free rate (from government bond yields),

- the implied dividend yield (computed from the spot-forward relationship).

Figure 5.1 shows the implied volatility surface organized by delta and maturity, along with the associated strike prices. Figure 5.2 presents the term structures of the risk-free rate and the dividend yield. Note that in the Bloomberg volatility surface interface, values shown in orange represent market quotes that are directly observed or interpolated using liquid instruments, while values shown in gray are extrapolated estimates beyond the liquid region of the surface. These gray entries are typically the result of Bloomberg's internal interpolation and extrapolation algorithms when no direct quotes are available for that specific combination of delta and maturity.

To improve the quality and stability of the calibration, we apply a truncation of the dataset. In particular, we exclude option maturities that are either too short or too long, specifically, less than 2 months or more than 5 years, since it is well documented that the Heston model performs poorly in these regimes due to instability in short-term implied volatility fits and unrealistic extrapolations in long-term dynamics [6, 14], that we have just described. Therefore, the final maturity range used for calibration is restricted to $[\frac{2}{12}, 5]$ years. The following is the Matlab code related to the truncation,

```
% Truncation to keep maturities between 2 months and 5 years
valid_idx = (tau >= 2/12) & (tau <= 5);
vol0 = vol0(valis_idx)
K = K(valid_idx);
tau = tau(valid_idx);
r = r(valid_idx)
q = q(valid_idx)
```

Expiry	Exp Date	ImpFwd	5DP	10DP	15DP	25DP	35DP	40DP	45DP	47.5DP	50D	47.5DC	45DC	40DC	35DC	25DC	15DC	10DC	5DC
1W	23 May 2025	43.87	33.00	29.23	27.51	25.92	25.13	24.86	24.63	24.53	24.43	24.34	24.26	24.11	23.99	23.87	24.06	24.50	25.71
			40.736	41.687	42.203	42.848	43.312	43.515	43.708	43.802	43.895	43.988	44.08	44.267	44.46	44.883	45.436	45.845	46.545
1M	16 Jun 2025	43.75	34.90	29.69	27.47	25.36	24.16	23.69	23.27	23.08	22.89	22.71	22.54	22.23	21.94	21.50	21.36	21.60	22.66
			37.205	39.307	40.398	41.738	42.686	43.095	43.48	43.666	43.849	44.029	44.209	44.568	44.934	45.729	46.756	47.52	48.876
2M	16 Jul 2025	43.91	40.22	34.01	31.32	28.73	27.30	26.75	26.27	26.05	25.84	25.63	25.44	25.07	24.73	24.14	23.73	23.70	24.24
			33.962	37.102	38.77	40.846	42.324	42.965	43.571	43.864	44.154	44.441	44.727	45.301	45.886	47.159	48.779	49.946	51.931
3M	16 Aug 2025	44.07	39.39	34.05	31.47	28.75	27.18	26.58	26.06	25.82	25.59	25.37	25.16	24.76	24.38	23.70	23.15	23.05	23.54
			32.465	35.924	37.886	40.402	42.205	42.987	43.724	44.082	44.435	44.785	45.134	45.832	46.546	48.091	50.048	51.456	53.896
6M	16 Nov 2025	44.32	39.37	35.10	32.81	30.00	28.16	27.42	26.76	26.45	26.16	25.87	25.60	25.08	24.59	23.71	23.04	22.91	23.38
			29.103	33.222	35.773	39.272	41.863	42.994	44.062	44.58	45.09	45.596	46.098	47.105	48.131	50.356	53.211	55.313	59.039
9M	16 Feb 2026	44.56	40.14	35.01	32.47	29.53	27.70	26.98	26.35	26.06	25.78	25.52	25.26	24.77	24.32	23.50	22.87	22.76	23.23
			26.675	31.597	34.607	38.729	41.804	43.158	44.444	45.071	45.691	46.307	46.922	48.16	49.431	52.218	55.847	58.552	63.392
1Y	16 May 2026	44.74	39.94	34.99	32.58	29.87	28.18	27.49	26.85	26.54	26.24	25.94	25.64	25.06	24.50	23.43	22.59	22.45	23.04
			25.124	30.382	33.663	38.247	41.764	43.338	44.844	45.579	46.307	47.03	47.75	49.194	50.665	53.855	58.00	61.168	67.106
18M	16 Nov 2026	44.96	38.47	34.06	31.83	29.18	27.45	26.75	26.11	25.82	25.53	25.26	25.00	24.50	24.03	23.23	22.71	22.74	23.48
			23.131	28.725	32.375	37.652	41.789	43.658	45.454	46.337	47.214	48.09	48.968	50.748	52.596	56.731	62.363	66.804	75.237
2Y	16 May 2027	45.45	39.06	34.33	32.12	29.52	27.81	27.11	26.49	26.20	25.92	25.66	25.40	24.93	24.50	23.78	23.36	23.45	24.22
			21.344	27.451	31.473	37.421	42.201	44.389	46.512	47.561	48.608	49.659	50.718	52.881	55.151	60.332	67.594	73.444	84.652
3Y	16 May 2028	46.21	34.28	31.95	30.47	28.45	27.03	26.45	25.95	25.72	25.51	25.31	25.13	24.81	24.54	24.24	24.44	24.98	26.32
			20.753	26.497	30.736	37.42	43.05	45.698	48.313	49.625	50.948	52.289	53.657	56.507	59.583	66.984	78.373	88.362	108.549
4Y	16 May 2029	46.97	34.14	31.48	29.88	27.85	26.54	26.04	25.62	25.44	25.27	25.12	24.98	24.76	24.60	24.56	25.15	25.99	27.79
			19.291	25.557	30.229	37.677	44.078	47.148	50.225	51.786	53.374	54.999	56.669	60.20	64.091	73.81	89.775	104.688	136.761
5Y	16 May 2030	47.83	34.16	31.33	29.62	27.47	26.11	25.60	25.18	25.01	24.85	24.71	24.59	24.40	24.30	24.45	25.45	26.64	28.84
			18.231	24.91	29.981	38.17	45.292	48.74	52.22	53.997	55.813	57.68	59.612	63.737	68.356	80.309	101.418	122.519	170.052
7Y	16 May 2032	49.84	33.41	30.66	28.92	26.70	25.32	24.83	24.44	24.28	24.15	24.04	23.95	23.85	23.87	24.38	25.99	27.47	29.64
			17.21	24.488	30.218	39.717	48.182	52.355	56.63	58.84	61.12	63.488	65.967	71.364	77.60	94.824	128.764	164.727	246.183
10Y	16 May 2035	53.37	34.04	30.24	28.06	25.59	24.28	23.87	23.58	23.48	23.41	23.36	23.35	23.42	23.66	24.96	28.64	32.02	37.11
			16.218	24.755	31.542	42.902	53.314	58.609	64.172	67.11	70.189	73.442	76.909	84.708	94.207	124.091	205.625	326.073	732.388

Fig. 5.1: Implied volatilities and strikes by delta and maturity for Bank of America (Bloomberg, 16 May 2025).

Source: Bloomberg Terminal

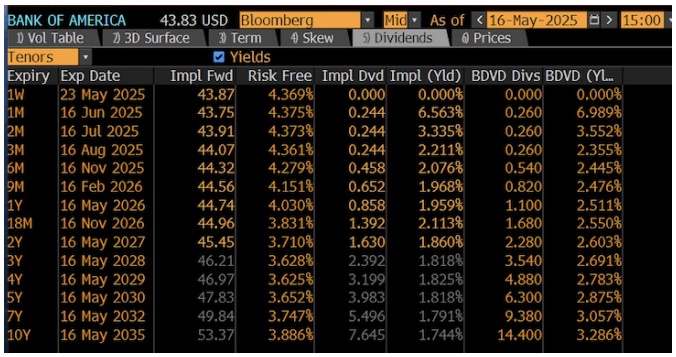

Expiry	Exp Date	Impl Fwd	Risk Free	Impl Dvd	Impl (Yld)	BDVD Divs	BDVD (Yl…
1W	23 May 2025	43.87	4.369%	0.000	0.000%	0.000	0.000%
1M	16 Jun 2025	43.75	4.375%	0.244	6.563%	0.260	6.989%
2M	16 Jul 2025	43.91	4.373%	0.244	3.335%	0.260	3.552%
3M	16 Aug 2025	44.07	4.361%	0.244	2.211%	0.260	2.355%
6M	16 Nov 2025	44.32	4.279%	0.458	2.076%	0.540	2.445%
9M	16 Feb 2026	44.56	4.151%	0.652	1.968%	0.820	2.476%
1Y	16 May 2026	44.74	4.030%	0.858	1.959%	1.100	2.511%
18M	16 Nov 2026	44.96	3.831%	1.392	2.113%	1.680	2.550%
2Y	16 May 2027	45.45	3.710%	1.630	1.860%	2.280	2.603%
3Y	16 May 2028	46.21	3.628%	2.392	1.818%	3.540	2.691%
4Y	16 May 2029	46.97	3.625%	3.199	1.825%	4.880	2.783%
5Y	16 May 2030	47.83	3.652%	3.983	1.818%	6.300	2.875%
7Y	16 May 2032	49.84	3.747%	5.496	1.791%	9.380	3.057%
10Y	16 May 2035	53.37	3.886%	7.645	1.744%	14.400	3.286%

Fig. 5.2: Risk-free rate and dividend yield term structures for Bank of America (Bloomberg, 16 May 2025).

Source: Bloomberg Terminal

The final goal of our calibration is to create a model that can fit the volatility surface displays in Figure 5.3, that is the raw implied volatility surface retrieved from Bloomberg, before any model fitting.

To plot our volatility surface, for convention, we use only negative values, but

they mean from [-0.5,0] for delta put and [-1,-0.5] for delta call.

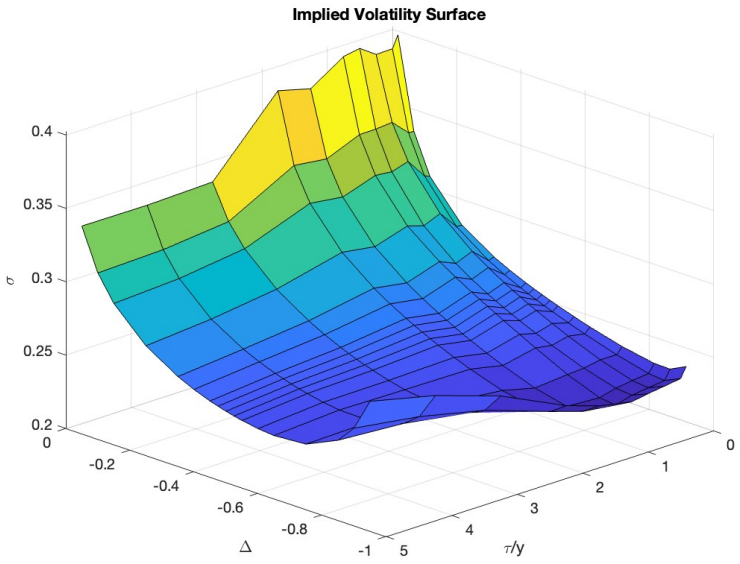

Fig. 5.3: Initial implied volatility surface from Bloomberg data.
The surface is expressed as a function of delta and maturity.
Source: own elaboration with Matlab

5.3 Methodology

To implement the calibration strategy in practice, we adopt a structured workflow, summarized in Figure 5.4. This scheme traces the full calibration path highlighting each conceptual step and its corresponding Matlab function. Following the workflow, we describe each component of the methodology adopted for calibrating the Heston model.

5.3.1 Market Prices Using the Black-Scholes Formula

Given Strikes K, maturities $\tau(\tau = T - t)$, implied volatilities σ, the spot price S_0, the risk-free rate $r(\tau)$, and the continuous dividend yield $q(\tau)$, we compute the

market-consistent option prices using the standard Black-Scholes formula for calls and puts that we have derived in eq. (2.3.3)

$$V = \overline{v}\left[S_0 e^{-q\tau}\Phi(d_1) - Ke^{-r\tau}\Phi(d_2)\right], \tag{5.15}$$

with

$$d_1 = \frac{\log\left(\frac{S_0}{K}\right) + (r - q + \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}, \tag{5.16}$$

$$d_2 = d_1 - \sigma\sqrt{\tau}, \tag{5.17}$$

where $\overline{v} = +1$ for calls and $\overline{v} = -1$ for puts. The main differences between this formula is and eq. (2.15) is that this formula is valid for both puts and calls that depends on the sign of $\overline{v}$ and that we use S_0 to describe the value of the underlying asset.

This step transforms the observed implied volatilities into market prices that will serve as the reference for model calibration. The relevant code in Matlab is

```
d1 = (log(S0./K)+(r_mat-q_mat+0.5*vol.^2).*tau_mat) ...
                                    ./(vol.*sqrt(tau_mat));
d2 = d1 - vol.*sqrt(tau_mat);
V = theta_mat.*(S0*exp(-q_mat.*tau_mat).*normcdf(theta_mat.*d1) ...
                -K.*exp(-r_mat.*tau_mat).*normcdf(theta_mat.*d2));
```

These computed prices V_{market} constitute the target surface to which the Heston model will be calibrated in subsequent steps.

5.3.2 Option Prices with the Heston Model

As described in the previous chapter 4, given the Heston model parameters $p = (\kappa, \overline{v}, \sigma, \rho, v_0)$, that before we called $\boldsymbol{\theta}$, we compute the theoretical option prices using a Fourier-based method inspired by the original approach of Heston

(1993) [9].

The function `europeanPricing.m` implements this method, incorporating even the term structures of the risk-free rate $r(\tau)$ and dividend yield $q(\tau)$ as retrieved from Bloomberg.

Below we describe the key steps in the pricing routine, with inline explanations of the Matlab code:

```
% Input: model parameters and market data
kappa = p(1);        % Rate of mean reversion
theta = p(2);        % long-term variance level
sigma = p(3);        % volatility of volatility
rho   = p(4);        % correlation between weiner processes
v0    = p(5);        % initial variance
```

The input also includes the option maturity τ, strike price K, spot price S_0, and the yield curves r and q for each maturity.

```
% Fourier grid parameters
u = 80;                    % upper bound of integration domain
ngrid = 2^12;              % number of evaluation points
dxi = pi / u;              % spacing in Fourier space
xi = dxi * (0:ngrid-1);    % frequency grid
```

- `u` determines how far we evaluate the option payoff across possible asset price changes. A wider range (larger `u`) improves accuracy but increases computation time.

- `ngrid` is the number of evaluation points on the frequency grid. Using a power of two allows efficient computation via the Fast Fourier Transform (FFT)[6].

[6]The Fast Fourier Transform (FFT) is an algorithm that allows efficient evaluation of Fourier

- `dxi` sets the spacing between points in the frequency grid and ensures numerical stability.
- `xi` is the actual grid of frequencies where we evaluate the characteristic function.

To ensure numerical stability when evaluating the Fourier transform, a damping factor α is introduced. This technique smooths the payoff function and guarantees convergence.

```
% Damping factor
if theta > 0
    alpha = -3;  % call option
else
    alpha = 4;   % put option
end
xi_shifted = xi + 1i * alpha;
```

We then compute the characteristic function of the log-price under the Heston model, which captures the full distributional behaviour of the asset price under stochastic volatility.

```
% Auxiliary variables for the characteristic function
c = kappa - rho * sigma * 1i * xi_shifted;
d = sqrt(c.^2 + sigma^2 * (1i * xi_shifted + xi_shifted.^2));
g = (c - d) ./ (c + d);
```

The characteristic function is then assembled. Note that we adjust it to match the given term structures of interest rate and dividend yield.

```
% Final characteristic function
```

integrals by exploiting symmetries in the data structure. It is commonly used in option pricing with characteristic functions.

```
phi = exp(1i * xi_shifted * log(S0) ...
      + (kappa * theta / sigma^2) * ((c - d) * tau ...
      - 2 * log((1 - g .* exp(-d * tau)) ./ (1 - g))) ...
      + (v0 / sigma^2) * ((c - d) .* (1 - exp(-d * tau))  ...
      ./ (1 - g .* exp(-d * tau))));
```

The final option price is obtained through numerical integration in the Fourier domain.

The function returns the theoretical Heston prices V_{model} for each (K, τ) pair.

5.3.3 Implementation of the Objective Function

In Section 5.1.1, we discussed various possible definitions of the loss function for model calibration, such as price-based, implied volatility-based, and vega-weighted objective functions. Among these, we opted for the relative price-based loss function defined in eq. (5.3), as it offers a good balance between interpretability and numerical stability.

The objective function is implemented in the Matlab script `objfun.m`, which is called during the optimization procedure. Below, we describe its structure.

```
function err = objfun(p, tau, K, theta, S0, r, q, Vmarket)
```

The function takes as input the same that we took in the last two section, given that the loss function is built on the errors of the prices with the two different methods described.

```
% Compute model prices using Heston pricing
Vmodel = europeanPricing(p, tau, K, theta, S0, q, r);
```

We call the pricing function `europeanPricing.m`, previously discussed, which returns the Heston prices V_{model}.

```
% Compute squared relative error
```

```
squared_diff = ((Vmodel - Vmarket).^2) ./ Vmarket.^2;
```

This implements the relative squared error for each option, following eq. (5.3). Normalizing by `Vmarket.^2` prevents large-priced options from dominating the objective function.

To avoid overfitting and enforce parameter regularity, we augment the loss with a regularization term based on the Elastic Net, which is the regulation term that we described erealier. Thus, adding the term in eq. 5.14 at the loss function, we have:

$$f_{\text{reg}}(p) = f_{\text{rel}}(p) + \lambda \left(\alpha \|p\|_1 + (1 - \alpha) \|p\|_2^2 \right). \tag{5.18}$$

This combines the sparsity-promoting effect of the L_1 norm with the stability of L_2, helping control both the parameter magnitude and interpretability. In Matlab:

```
% Elastic Net regularization
lambda = 0.01;    % Regularization weight
alpha  = 0.5;     % Balance between L1 and L2
reg_term = lambda * (alpha * sum(abs(p)) + (1 - alpha) * sum(p.^2));

% Final objective: relative error + Elastic Net penalty
err = mean(squared_diff, 'all') + reg_term;
```

The result is a smooth, convexified objective that favors both data fit and model parsimony. This encourages the optimizer to explore only meaningful regions of the parameter space while maintaining stability and avoiding large or erratic parameter estimates.

5.3.4 No-Arbitrage Constraints

As anticipated in the general discussion on constraints section 5.1.3 are used to ensure economically and mathematically meaningful solutions. While the regularization term was added directly to the objective function (see Section 5.3.3), here

we describe the implementation of external constraints enforced separately via the function `nonlcon.m`. These are used to ensure that model prices are arbitrage-free, following the conditions described in [1] and [4].

The constraints are implemented in the following Matlab script:

```
function [c,ceq] = nonlcon(p,tau_mat,K,theta_mat,S0,r_mat,q_mat,N)
```

Below we explain the financial meaning of each condition included in the constraint vector `c`, which must satisfy $c \leq 0$[7]:

Monotonicity in maturity: $\frac{\partial V}{\partial \tau} \geq 0$

```
dV1 = diff(V_model,1,1);
ctau = -dV1(:);
```

This constraint enforces that the price of an option should not decrease as maturity increases. Financially, a longer maturity means more time value, therefore greater uncertainty, so the option should be worth at least as much as the same one with shorter maturity.

Monotonicity in strike: $\frac{\partial V_{\text{call}}}{\partial K} \leq 0$, $\frac{\partial V_{\text{put}}}{\partial K} \geq 0$

```
dV2_call = diff(V_model_call,1,2);
dV2_put  = diff(V_model_put,1,2);
cK_call  = dV2_call(end);
cK_put   = -dV2_put(1);
```

As the strike increases, the price of a call should decrease, and the price of a put should increase.

Slope bounds with respect to the strike: $-e^{-r\tau} \leq \frac{\partial V}{\partial K} \leq 0$

[7]The condition $c(p) \leq 0$ follows a standard convention in constrained optimization: each constraint must be expressed in a way that it is satisfied when the corresponding element is less than or equal to zero. This is the required format for solvers like `fmincon` in Matlab. Therefore, when we write a condition such as $\frac{\partial V}{\partial \tau} \geq 0$, it is reformulated as $-\frac{\partial V}{\partial \tau} \leq 0$ to fit this convention.

```
dVdK_call = dV2_call ./ diff(K_call,1,2);
dVdK_put  = dV2_put  ./ diff(K_put,1,2);
discount_factor = exp(-r_mat(:,1).*tau_mat(:,1));
cdf_call = -discount_factor(:) - dVdK_call(1);
cdf_put  = -discount_factor(:) + dVdK_put(end);
```

These slope bounds ensure that the rate of change of the option price with respect to strike is within theoretical limits. The discount factor appears as a theoretical upper bound on the slope, derived from the Breeden-Litzenberger relation[8].

Convexity with respect to strike: $\frac{\partial^2 V}{\partial K^2} \geq 0$

```
d2VdK2_call = diff(dVdK_call,1,2);
d2VdK2_put  = diff(dVdK_put,1,2);
cK2_call = -d2VdK2_call(:);
cK2_put  = -d2VdK2_put(:);
```

Convexity of the option price in strike is a key no-arbitrage requirement. It reflects the idea that the price of an option must not have "kinks" or "non-smooth" behavior as the strike varies, and ensures the existence of a non-negative probability distribution.

Thus, this is the final constraint vector:

```
c = [ctau; cK_call; cK_put; cdf_call; cdf_put; cK2_call; cK2_put];
ceq = [];
```

These constraints are passed to the optimizer `fmincon`, that we will describe in the following section, as a separate nonlinear constraint function, independent from the objective function and enforce fundamental structural conditions that option prices must satisfy in order to avoid arbitrage opportunities.

[8]See Appendix A.2 for the explanation of the Breeden-Litzenberger relation and its role in this section.

5.3.5 Initial Parameter Guess and Black-Scholes Volatility

To initialize the calibration routine, we must specify both a starting point for the Heston parameters and a reference volatility for the Black-Scholes model. These choices significantly affect the convergence behavior of the optimization algorithm, especially for a non-linear model like this which may exhibit multiple local minima.

The initial guess for the Heston parameters was selected based on economic intuition and empirical studies from the literature (e.g., [8, 14]). Specifically, we used:

- $\kappa_0 = 2.0$ (moderate mean reversion),
- $\overline{v}_0 = 0.05$ (long-run volatility level of 22%),
- $\sigma_0 = 0.5$ (moderate volatility of volatility),
- $\rho_0 = -0.5$ (typical leverage effect),
- $v_0 = 0.04$ (initial volatility of 20%).

These values reflect a market with negative correlation between asset returns and volatility (consistent with equity markets) and moderate variance fluctuations.

Regarding the Black-Scholes model, we assumed a constant volatility equal to the implied volatility of the 1-year at-the-money (ATM) call option. This choice is motivated by the fact that the ATM region is where the Black-Scholes assumptions are most accurate and where implied volatility tends to be most stable. The estimated ATM volatility at this point was approximately $\sigma_{\text{BS}} = 26.5\%$, which was used uniformly across all strikes and maturities in the Black-Scholes pricing and implied volatility surfaces.

5.3.6 Optimization

To calibrate the Heston model, we solve a nonlinear constrained optimization problem. For this purpose, we use Matlab built-in solver `fmincon`, which is specifically designed to handle such problems.

Among the algorithms supported by `fmincon`, we select the interior-point method. This method is well-suited for problems with nonlinear constraints and provides good convergence behavior even in the presence of multiple local minima or ill-conditioned objective functions.

The interior-point method operates by transforming the original constrained problem into a sequence of approximate unconstrained problems. This is achieved by incorporating the inequality constraints directly into the objective function using a so-called barrier term, which penalizes infeasible solutions by becoming very large near the boundary of the feasible region. The modified objective function, known as the barrier function, takes the form

$$f_{\text{barrier}}(p, \mu) = f(p) - \mu \sum_j \log(-c_j(p)), \tag{5.19}$$

where $c_j(p) \leq 0$ are the inequality constraints, and $\mu > 0$ is a parameter that controls the strength of the penalty.

The optimization proceeds iteratively. At each step, the method:

1. solves a subproblem by minimizing the barrier-augmented objective using Newton-type steps,
2. updates the parameters p within the interior of the feasible region,
3. and gradually reduces the barrier parameter μ so that the approximation improves over time.

As μ becomes smaller, the solution of the barrier problem converges to the solution of the original constrained problem.

The result of this is a set of parameters p that should be the most efficient parameters that our calibration could reach.

5.3.7 Recovering Implied Volatilities

Once the model prices have been computed using the calibrated parameters, we recover the corresponding implied volatilities to enable a direct comparison with market data. This step is crucial because implied volatilities, rather than raw option prices, are typically used by market practitioners to evaluate model accuracy and behavior.

To this end, we use the Matlab built-in function `blsimpv`, which computes the Black-Scholes implied volatility given an option price and the other contract terms. Specifically, for each model-generated price V_{model}, the function solves the following inverse problem

$$C_{\mathrm{BS}}(S_0, K, r, q, \sigma, \tau) = V_{model}, \tag{5.20}$$

where the goal is to find the value of σ (volatility) that matches the model price under the Black-Scholes formula.

Since this is a nonlinear equation in σ, `blsimpv` uses numerical root-finding methods internally. The result is a matrix of implied volatilities corresponding to each delta and maturity, which can then be plotted and compared directly against the market implied volatility surface.

5.4 Results

5.4.1 Calibrated Parameters and Performance Metrics

After performing the calibration using the procedure described previously, we obtain the following optimal parameters for the Heston model with market data

that we described in sec. 5.2.

Parameter	Description	Calibrated value
κ	Rate of mean reversion	7.528391
$\overline{v}$	Long-run variance	0.069999
σ	Volatility of variance	1.165200
ρ	Correlation	-0.505768
v_0	Initial variance	0.096673

Table 5.1: Calibrated parameters for the Heston model. Source: own elaboration.

The calibration process converged to a final value of the objective function

Loss function value: 0.038346 and required approximately CPU time: 93 seconds

Each parameter of the Heston model influences the shape of the implied volatility surface in a distinct way. Understanding these effects is crucial to interpret the calibration results.

- **Mean-reversion rate** κ: Determines how quickly the variance reverts to its long-run level $\overline{v}$. A higher κ results in faster mean reversion, which generally reduces instant volatility persistence. On the implied volatility surface, increasing κ tends to flatten the term structure of volatility, especially for short-term maturities.

- **Long-run variance** $\overline{v}$: Represents the level around which the variance oscillates in the long run. It sets the baseline level of volatility. A higher $\overline{v}$ raises the entire volatility surface uniformly, particularly affecting long-dated options.

- **Volatility of volatility** σ: Governs how volatile the variance itself is. It controls the curvature and smile effects on the implied volatility surface.

Larger values of σ enhance the smile and skew, especially for options far from the money, and cause more pronounced convexity.

- **Correlation ρ between asset and variance**: This parameter is key to reproducing the skew observed in equity markets. A negative ρ (typically observed empirically) implies that volatility tends to rise when the asset price drops. This causes implied volatility to increase for lower strikes (put options) and decrease for higher strikes, generating a leftward skew on the surface.

- **Initial variance** v_0: Reflects the instantaneous variance at the time of calibration. It influences the short end of the implied volatility term structure. If v_0 is far from $\overline{v}$, the volatility curve will initially drift toward the mean, producing term structure effects at short maturities.

In our case, the calibrated parameters provide meaningful insights into the dynamics captured by the Heston model. The estimated value of the mean-reversion rate $\kappa \approx 7.53$ indicates that the variance is expected to revert quickly to its long-run average, implying a market where volatility shocks are short-lived and the volatility term structure tends to be relatively flat for shorter maturities. Moreover, the proximity between the long-run variance $\overline{v} \approx 0.07$ and the initial variance $v_0 \approx 0.097$ suggests that the current volatility regime is not far from its expected long-term level, pointing to a market in equilibrium or steady state.

The volatility of volatility parameter, $\sigma \approx 1.16$, is relatively high, reflecting the presence of strong smile and skew features in the market-implied volatility surface. This value allows the model to replicate sharp curvature, especially for options that are deep in-the-money or out-of-the-money. Finally, the negative correlation $\rho \approx -0.51$ between the asset price and the variance confirms the typical leverage effect observed in equity markets, where negative returns are associated

with increased volatility. This contributes to the asymmetric shape of the implied volatility surface, especially for lower strikes, which display higher implied volatilities due to increased downside risk.

Taken together, these parameters highlight the ability of the Heston model to reproduce both the level and the structure of observed market volatilities with greater flexibility.

5.4.2 Heston vs Black-Scholes

As we said, the calibration process is a fundamental component in any model-based pricing framework. Its importance goes far beyond the mere interpretation of the model parameters.

What truly validates a model is its ability to replicate observed market phenomena. In the context of option pricing, this translates into the model's capability to reproduce both the implied volatility surface and the option price surface with high accuracy.

This is where the Heston model demonstrates an advantage over the Black-Scholes framework.

In the following analysis, we compare the two models in terms of pricing accuracy and their ability to replicate the implied volatility surface. Indeed, as we said the calibration started from our market implied volatility surface Fig. 5.3.

As observed from the figures below, both the Black-Scholes and Heston models are able to reproduce the overall shape of the market option price surface with notable precision. Despite their fundamental differences, the resulting pricing surfaces appear visually close to that of the market.

In particular, the shape, slope, and general curvature of the surfaces align well in all three cases, suggesting that both models are capable of capturing the main features of observed option prices. The figures presented below show, respectively,

the price surface generated from the market data, the one obtained using the Black-Scholes model, and the one generated by the Heston model.

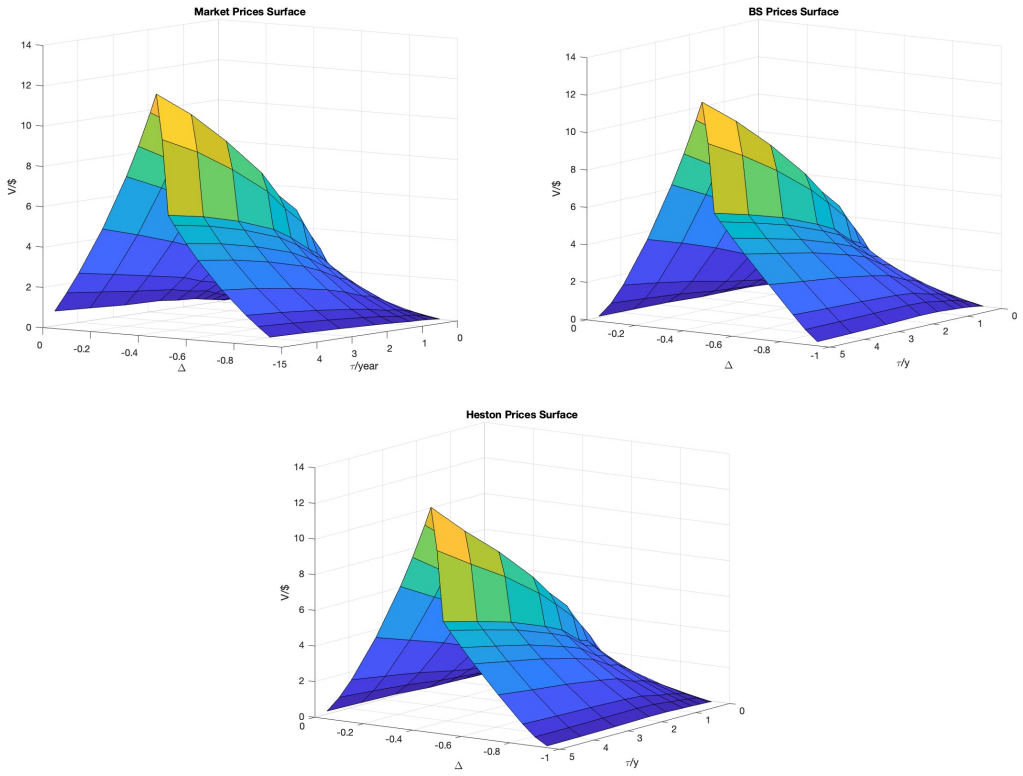

Fig. 5.5: Top-Left: Market prices surface; Top-Right: Black-Scholes prices surface; Bottom: Heston prices surface. Source: own elaboration with Matlab.

To assess the fidelity of the models in replicating real market behavior, we now compare the theoretical surfaces directly against the empirical surface. The following Fig. 5.6 overlay the option prices computed from both the Black-Scholes and Heston models with those observed in the market.

From these plots, we notice that the Heston model tends to match the market data more closely across the entire range of deltas and maturities. In particular, it better captures curvature and variation in the medium-term and long-term regions of the surface. The Black-Scholes model, while generally accurate, appears slightly

less flexible in reproducing such nuances.

An exception to this trend occurs in the region of the surface characterized by long maturities and a delta close to -0.5, which corresponds to the transition zone between put and call options. In this specific area, the Black-Scholes model exhibits an advantage, managing to align slightly better with the market prices than the Heston model.

This phenomenon can be interpreted as follows. When delta is close to -0.5, we are considering put options that are at-the-money forward. These options have a strike price close to the expected future value of the underlying asset. In this region, the market prices are particularly stable and sensitive to volatility variations, which makes the evaluation highly accurate and reliable.

The Black-Scholes model, although based on the strong assumption of constant volatility, tends to perform very well in this central region of the distribution. Indeed, the lognormal assumption underlying Black-Scholes produces its best approximations around the current spot price[9], where the probability density is highest. This makes the model naturally suited to price ATM options with good precision, especially when the implied volatility surface is relatively smooth, as it often is for these maturities and deltas.

On the other hand, the Heston model is designed to capture complex features of the volatility surface, such as skew and smile, through a stochastic variance process. While this flexibility is essential to fit the entire surface, it introduces more complexity into the pricing formula, involving numerical inversion of characteristic functions or Fourier transforms. These numerical procedures can introduce minor

[9]In the Black-Scholes model, the logarithm of the underlying price follows a normal distribution, meaning the price itself follows a lognormal distribution. This implies that the highest probability density is located near the expected future price of the asset, which, under normal market conditions, is close to the current spot price. As a result, the model's approximations are particularly accurate around the at-the-money (ATM) region.

approximation errors, especially in regions of the surface where the curvature is low and precision is critical, such as near $\Delta = -0.5$ and long maturities.

As a result, while Heston performs better globally across most of the surface, the Black-Scholes model shows a slight edge in this narrow band where its structural assumptions are naturally more aligned with market dynamics and where even small deviations in model output are more easily noticed.

These visual impressions are confirmed by the numerical error metrics. By summing the absolute differences between model and market prices across all points on the surface, we obtain the following results:

```
--- Error Prices Analysis ---
Black-Scholes vs Market:
  Absolute Error: 41.557900
  Mean Error:     0.244458

Heston vs Market:
  Absolute Error: 25.033199
  Mean Error:     0.147254
```

These values quantitatively confirm that the Heston model achieves a better overall fit than the Black-Scholes model in terms of pricing accuracy. However, the most significant improvements are observed not in the price surfaces, but in the ability of the Heston model to replicate the market implied volatility surface.

As observed from the figures below, the differences between the two models are now more pronounced. The Black-Scholes model assumes a constant volatility, which results in a completely flat surface, a clear contrast with the observable market patterns.

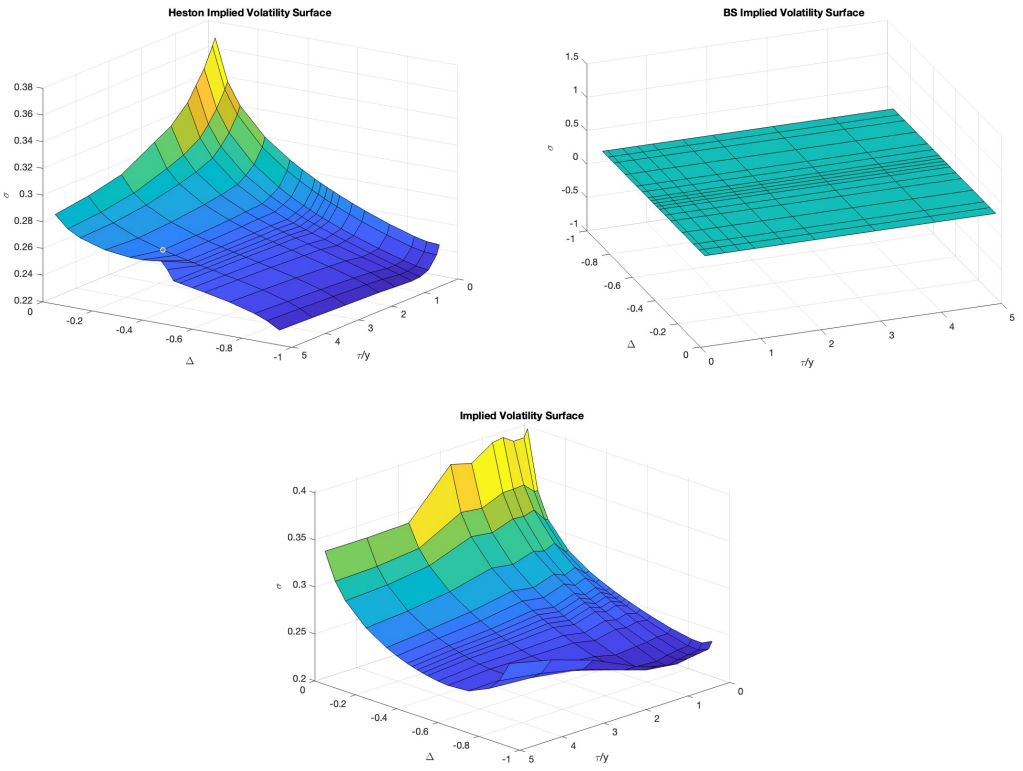

Fig. 5.7: Top-Left: Heston implied volatility surface; Top-Right: Black-Scholes implied volatility surface; Bottom: Market implied volatility surface. Source: own elaboration with Matlab.

To better appreciate the models's performances, we now compare their implied volatility surfaces directly with the one derived from market data. Figure 5.8 overlays the theoretical implied volatilities from both Black-Scholes and Heston with the market-implied volatility surface.

Visually, the Heston-implied volatility surface aligns much more closely with the market one. Indeed, if we see the shape of the surfaces are really similar. The main goal, that was to replicate the smile is done with the exception that Heston can't replicate perfectly the smile over deep in-the-money call, in fact in this area, there is a difference between the surfaces. In particular the main difference is around the long-maturity options, as we can see when $\tau/y \approx 5$.

This discrepancy can be attributed to the fact that, although the Heston model incorporates stochastic volatility, it still lacks mechanisms such as jumps or higher-order corrections that could better reflect market dynamics in extreme moneyness regimes and even the fact that for long-maturites the market data are not so precise, as we said in sec.5.2.

These visual findings are confirmed by numerical error metrics. As with prices, we compute both the absolute and mean error between model and market volatilities over the entire surface:

```
--- Error Volatility Analysis ---
Black-Scholes vs Market:
  Absolute Error: 4.408500
  Mean Error:     0.025932

Heston vs Market:
  Absolute Error: 2.268112
  Mean Error:     0.013342
```

The Heston model reduces the mean error by nearly half compared to Black-Scholes and this quantitative evidence reinforces the visual impression, Indeed, while Black-Scholes provides a useful benchmark and performs not bed, it cannot reproduce the full structure of the implied volatility surface as we said at the beginning of the thesis. The Heston model, offers a significantly more accurate framework for capturing market volatility dynamics.

5.5 Financial Relevance of Calibration

The ultimate purpose of model calibration in financial contexts is not merely academic or theoretical, it is intrinsically tied to concrete tasks such as pricing,

hedging, and risk management of derivative instruments. Among all the quantities derived from option models, the implied volatility plays a pivotal role, as it encapsulates the market's expectation of future asset fluctuations in a model-consistent framework.

Accurately fitting the implied volatility surface enables traders and risk managers to assess option values more consistently across different strikes and maturities, which is essential in constructing arbitrage-free pricing systems. Moreover, the implied volatility surface serves as a core input in the computation of Greeks, which are used to build dynamic hedging strategies. Even small deviations in volatility estimates can lead to significant errors in delta or vega hedging, ultimately exposing portfolios to unintended risk.

In addition, a well-calibrated model supports the evaluation of exotic options, which do not have closed-form solutions and must be priced using Monte Carlo or tree-based methods built on top of a volatility model and if the base model fails to replicate observed market behavior the pricing of these instruments may be materially biased.

From a regulatory and risk management point of view, institutions must also compute Value-at-Risk (VaR) and stress scenarios[10] that depend on forward-looking volatility estimates. A model that fits the implied volatility surface well provides a more accurate and consistent risk profile across different market conditions.

As discussed in [8], the calibration of volatility models is therefore not simply a technical exercise, but a critical component in achieving reliable pricing and hedging in practice. Furthermore, as emphasized in [14], improved fit to implied

[10]Stress scenarios refer to hypothetical situations used to evaluate the behavior of a financial portfolio under extreme market conditions, such as sudden interest rate spikes, volatility surges, or market crashes.

volatility surfaces translates into better replication of market dynamics, reducing the need for ad hoc volatility adjustments and enhancing model credibility in front-office applications.

In summary, high-precision calibration enables more faithful replication of market conditions, minimizes hedging errors, and strengthens the reliability of risk metrics and all of these tasks are essential for effective decision-making in financial markets.

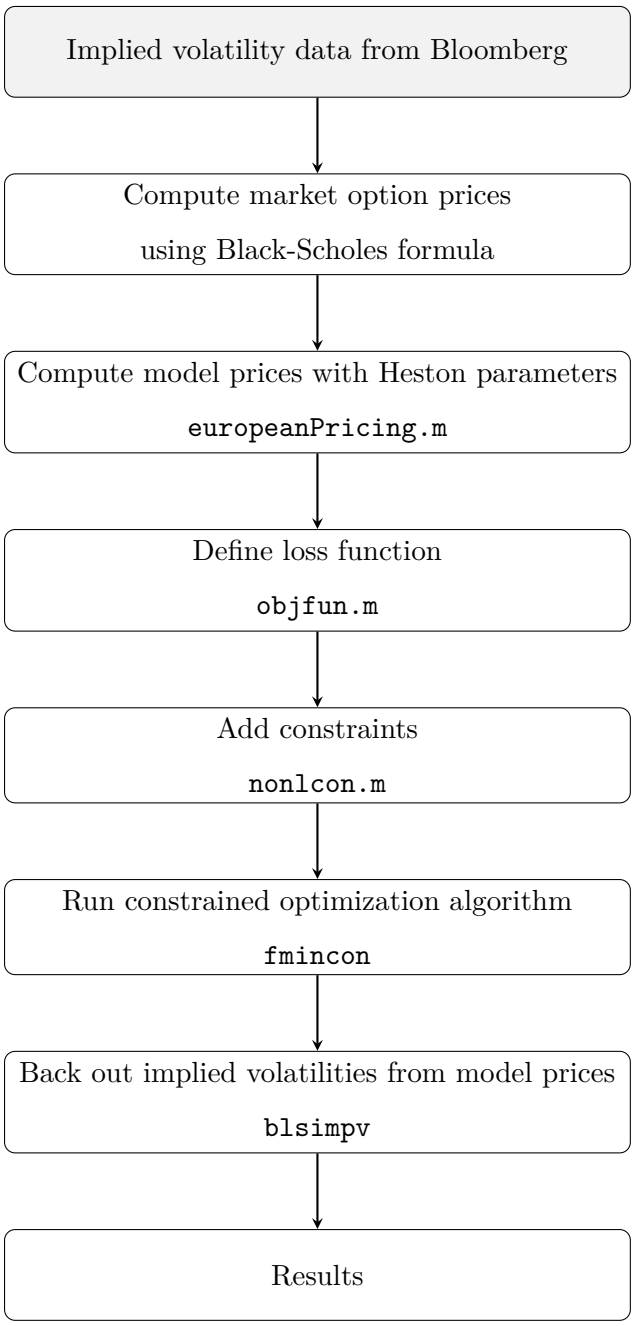

Fig. 5.4: Calibration flowchart and Matlab functions. Source: Own elaboration with Matlab.

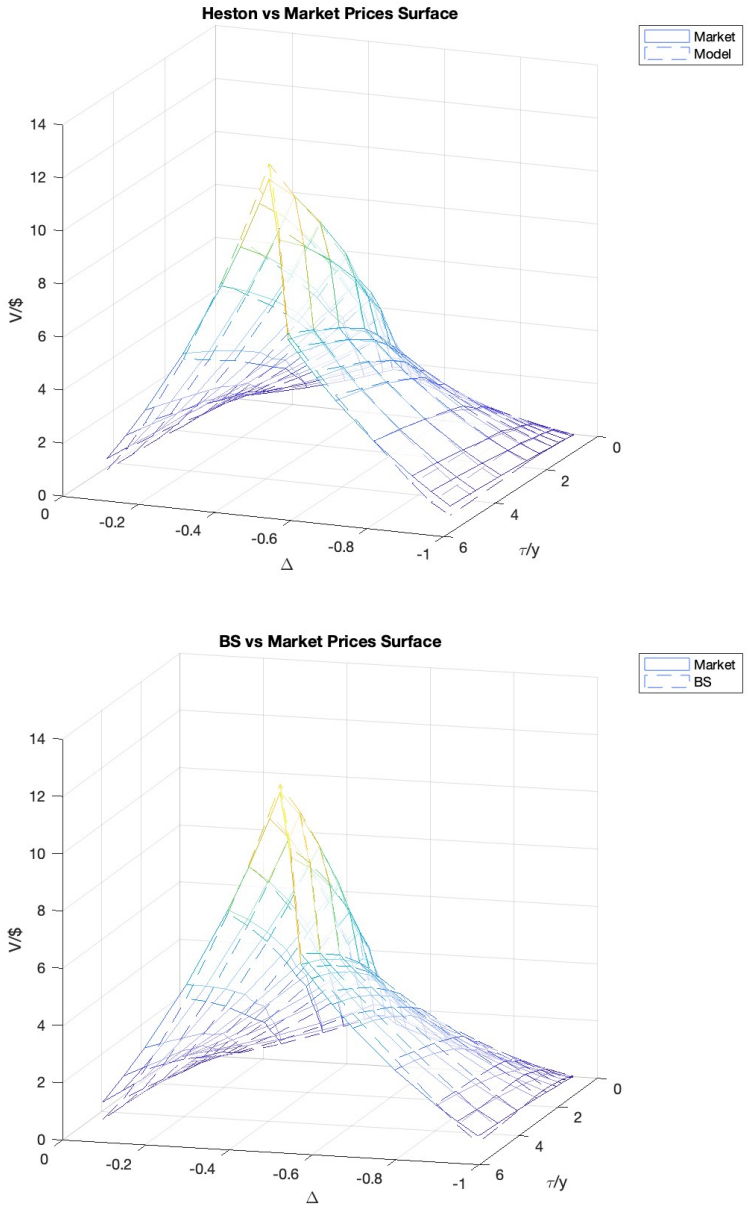

Fig. 5.6: Top: Heston vs Market ; Bottom: Black-Scholes vs Market. Source: own elaboration with Matlab.

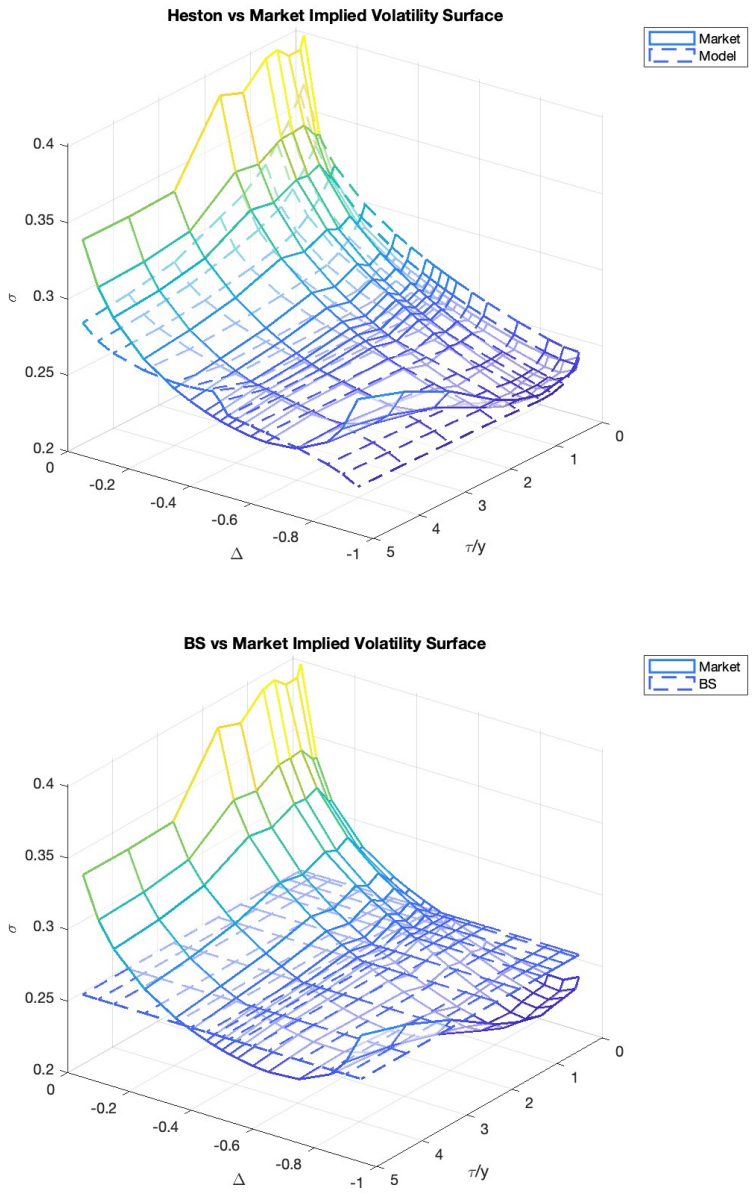

Fig. 5.8: Top:Heston vs Market implied volatility surface; Bottom: Black-Scholes vs Market implied volatility surface. Source: own elaboration with Matlab.

Conclusion

In this thesis, we have investigated option pricing models with a particular focus on the calibration process of the Heston model. Our journey began by introducing the notion of pricing models, with special attention to one of the most foundational and widely used frameworks: the Black-Scholes model. While models such as Merton,Äôs, with adjustments for dividend payments, extend Black-Scholes in certain directions, they still retain one of its major limitations: the assumption of constant volatility.

To address this, we explored the empirical inconsistencies between market data and the theoretical implications of constant volatility. Specifically, we discussed the phenomena of the volatility smile and skew, which manifest in the market's implied volatility surface and are incompatible with the flat volatility predicted by Black-Scholes. As a response, we briefly introduced the local volatility approach via Dupire's formula, acknowledging its ability to model the smile statically, yet also highlighting its limitations due to determinism and instability under certain market conditions.

Consequently, we turned our attention to a more flexible, though more complex, model: the Heston model. This stochastic volatility model allows the volatility itself to evolve as a random process over time, better capturing the dynamic behavior observed in financial markets. Introducing such a model, however, significantly departs from the Black-Scholes framework and introduces considerable

mathematical and computational complexity.

After presenting the theoretical structure of the Heston model, including its dynamics, assumptions, and semi-analytical solution for option pricing, we shifted our focus to its calibration. We explained the general problem of calibration. This step represents a crucial bridge between theory and practice.

The calibration task proved to be the most technically challenging part of this thesis. As discussed, effective calibration requires several components to be addressed simultaneously: the selection of a suitable objective function, the choice of an appropriate optimization algorithm, and the imposition of realistic constraints to ensure meaningful and interpretable parameter values. Even decisions such as the initial guess of parameter values have a significant impact on the convergence and stability of the algorithm.

To validate our methodology, we extracted an implied volatility surface from the Bloomberg Terminal for Bank of America (BAC) equity options. This real-world dataset served as the foundation for our calibration test. We then implemented our calibration routine in MATLAB, aiming to align the Heston-implied volatilities with the market-observed ones. To assess the effectiveness of our calibration, we used the Black-Scholes model as a benchmark. The comparison clearly demonstrated that, although more computationally demanding, the calibrated Heston model succeeded in closely replicating the shape and curvature of the market's implied volatility surface.

Nevertheless, our results also revealed some limitations of the Heston model. While the model successfully captures the general smile and skew structure, there are still discrepancies in certain regions of the volatility surface, for instance for deep in-the-money call options or long maturities, where the model fails to fully replicate market behavior. This may be due to the equity underlying exhibiting jump behaviors, which the Heston model does not account for.

These observations suggest promising directions for future research. One natural extension would be to explore jump-diffusion models or hybrid frameworks that combine stochastic volatility with discontinuous processes, such as the Bates model. These approaches may offer further improvements in aligning theoretical models with complex market realities.

In conclusion, this thesis has highlighted the importance and challenges of calibrating advanced stochastic models in quantitative finance. While the Black-Scholes model provides a useful starting point, it is insufficient for capturing the nuanced features of market volatility. The Heston model, although more intricate, offers a viable and more accurate alternative. Our calibration procedure demonstrated that, when implemented with care, it can significantly reduce the gap between theoretical and market prices, providing a valuable tool for both academic study and practical application in derivatives pricing.

Appendix

Proposition A.1 (Law of Large Numbers)**.** *Let $X_1, X_2, \ldots$ be a swquence of i.i.d. random variables, each with finite meanμ. Thus*

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{X_1 + \ldots + X_n}{n} \to \mu\right) = 1\,,$$

or equivalently, almost surely,

$$\lim_{n\to\infty} \frac{X_1 + \ldots + X_n}{n} = \mu\,.$$

Proposition A.2 (Central Limit Theorem)**.** *Let $X_1, X_2, \ldots$ a swquence of i.i.d. random variables with mean $\mu < \infty$ and variance $\sigma^2 < \infty$. Define: $S_n = \sum i = 1^n X_i$ e $Z_n = \frac{S_n - n\mu}{\sigma\sqrt{n}}$ as the normalized sum, then, for every $t \in \mathbb{R}$we have:*

$$F_{Z_n}(t) := \mathbb{P}(Z_n \leq t) \xrightarrow{q.c.} \mathbb{P}(Z \leq t) \quad per\ n \to \infty\,,$$

dove $Z \sim \mathcal{N}(0,1)$. Equivalently, Z_n converges in distribution to $Z \sim \mathcal{N}(0,1)$ for $n \to \infty$.

A.1 Monte Carlo Simulation

The Monte Carlo method is a class of computational algorithms that rely on repeated random sampling to estimate numerical results. In finance, it is especially valuable for evaluating expectations under stochastic processes, such as option prices, where analytical solutions may be unavailable.

Formally, suppose one is interested in estimating the expected value of a random variable X under a certain probability distribution:

$$\mathbb{E}[X] = \int X(\omega)\, d\mathbb{P}(\omega).$$

Monte Carlo estimation proceeds by generating N independent random samples $\omega_1, \ldots, \omega_N$ from the distribution $\mathbb{P}$, and then computing the sample average:

$$\mathbb{E}[X] \approx \frac{1}{N} \sum_{i=1}^{N} X(\omega_i).$$

According to the law of large numbers, this estimate converges almost surely to the true expected value as $N \to \infty$.

In practice, the method involves three key steps:

1. Simulation of paths: Generate realizations of the underlying stochastic variables .

2. Evaluation of the payoff: Compute the value of the contingent claim (e.g., option payoff) on each simulated path.

3. Discounting and averaging: Average the discounted payoffs across all paths to estimate the present value.

This method is widely used due to its generality, particularly in high-dimensional problems or in models where the underlying dynamics (e.g., local or stochastic volatility) prevent analytical treatment.

Source: Note del corso "Probabilità e applicazioni alla finanza"

A.2 The Breeden-Litzenberger Relation

The Breeden-Litzenberger relation [3] shows how European call option prices embed information about the risk-neutral distribution of the underlying asset at

maturity. Specifically, as discussed in [10] under no-arbitrage and smoothness assumptions, the second derivative of the call price with respect to the strike gives the risk-neutral density,

$$\frac{\partial^2 C(K,\tau)}{\partial K^2} = e^{-r\tau} f^{\mathbb{Q}}(K), \tag{21}$$

where $f^{\mathbb{Q}}(K)$ is the risk-neutral probability density at strike K.

This result has two important implications:

- Since $f^{\mathbb{Q}}(K) \geq 0$, call prices must be convex with respect to the strike,

$$\frac{\partial^2 C}{\partial K^2} \geq 0. \tag{22}$$

- The call price must satisfy a slope condition with respect to the strike price: the derivative of the call price with respect to the strike must lie between 0 and minus the discount factor $e^{-r\tau}$,

$$-e^{-r\tau} \leq \frac{\partial C}{\partial K} \leq 0, \tag{23}$$

  which reflects the fact that call prices decrease with strike, but not too rapidly.

These conditions are useful when constructing no-arbitrage constraints for calibration, as they ensure that the prices are consistent with an underlying probability distribution.

A.3 L_1 and L_2 Norms

In this section, we briefly recall the two most commonly used vector norms in optimization:

Proposition A.1. *The L_1 norm or Manhattan norm of a vector $x = (x_1, x_2, \ldots, x_n)$ is defined as:*

$$\|x\|_1 = \sum_{i=1}^{n} |x_i|.$$

Proposition A.2. *The L_2 norm or Euclidean norm is defined as:*

$$\|x\|_2 = \left(\sum_{i=1}^{n} x_i^2 \right)^{1/2}.$$

It corresponds to the standard Euclidean distance from the origin.

Bibliography

[1] Yacine Aıt-Sahalia and Jefferson Duarte, *Nonparametric option pricing under shape restrictions*, Journal of Econometrics **116** (2003), no. 1-2, 9–47.

[2] Fischer Black and Myron Scholes, *The pricing of options and corporate liabilities*, Journal of political economy **81** (1973), no. 3, 637–654.

[3] Douglas T Breeden and Robert H Litzenberger, *Prices of state-contingent claims implicit in option prices*, Journal of business (1978), 621–651.

[4] Peter Carr and Dilip B Madan, *A note on sufficient conditions for no arbitrage*, Finance Research Letters **2** (2005), no. 3, 125–130.

[5] Yiran Cui, Sebastian del Baño Rollin, and Guido Germano, *Full and fast calibration of the heston stochastic volatility model*, European Journal of Operational Research **263** (2017), no. 2, 625–638.

[6] José Da Fonseca and Martino Grasselli, *Riding on the smiles*, Quantitative Finance **11** (2011), no. 11, 1609–1632.

[7] Bruno Dupire, *Pricing with a smile*, Risk Magazine **7** (1994), no. 1, 18–20.

[8] Jim Gatheral, *The volatility surface: a practitioner's guide*, John Wiley & Sons, 2011.

[9] Steven L Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, The review of financial studies **6** (1993), no. 2, 327–343.

[10] J. C. Hull and E. Barone, *Opzioni, futures e altri derivati*, Pearson, 2009.

[11] Shuaiqiang Liu, Anastasia Borovykh, Lech A Grzelak, and Cornelis W Oosterlee, *A neural network-based framework for financial model calibration*, Journal of Mathematics in Industry **9** (2019), no. 1, 9.

[12] Robert C. Merton, *Theory of rational option pricing*, Bell Journal of Economics and Management Science **4** (1973), no. 1, 141–183.

[13] Fabrice Douglas Rouah, *Derivaton of local volatility*, 1998.

[14] ———, *Simplified derivation of the heston model*, The Heston Model and Its Extensions in Matlab and C# (2013), 4–7.

Ringraziamenti

Un ringraziamento speciale ai due relatori di questa tesi.

Il professor. Hlafo Alfie Mimun per avermi accompagnato con pazienza e disponibilità lungo l'intero percorso. I suo consigli e la sua passione mi hanno ispirato a fare sempre del mio meglio.

Al professor. Guido Germano per il supporto metodologico e per avermi fornito strumenti fondamentali, che hanno arricchito in maniera significativa il contenuto della tesi. La sua esperienza e competenza sono stati per me un punto di riferimento fondamentale.