INSTITUTO
UNIVERSITÁRIO
DE LISBOA

Testing Alternative Methodologies for Calibrating the Heston Model with Jumps for S&P 500 Index Options

Francisco Santos Simões

Master in Financial Mathematics

Supervisor:
PhD, José Carlos Gonçalves Dias, Full Professor,
Iscte - IUL

October, 2024

BUSINESS
SCHOOL

Ciências ULisboa | Departamento de Matemática

Department of Finance

Department of Mathematics

Testing Alternative Methodologies for Calibrating the Heston Model with Jumps for S&P 500 Index Options

Francisco Santos Simões

Master in Financial Mathematics

Supervisor:
PhD, José Carlos Gonçalves Dias, Full Professor,
Iscte - IUL

October, 2024

Acknowledgment

This thesis would not have been possible without the support and guidance of many people, to whom I am deeply grateful. I would like to highlight the unwavering encouragement and motivation from my colleagues Ana Ferreira, Bernardo Catado, and Gonçalo Caldeirinha. I am also immensely thankful to my family, whose support and understanding have been a constant source of strength throughout this journey.

Lastly, I would like to express my sincere gratitude to Professor José Carlos Dias, who was always available to help and whose guidance and wisdom have been invaluable.

I dedicate this work to all those mentioned above.

Resumo

Esta tese foca-se na avaliação e comparação de dois modelos de determinação de preços de opções: o modelo de Heston (1993), que integra volatilidade estocástica, e o modelo de Bates (1996), que expande o modelo de Heston ao incluir saltos (jumps) nos preços dos ativos.

Embora o modelo de Black-Scholes (1973) e Merton (1973) tenha sido amplamente utilizado para a avaliação de opções, o seu pressuposto de volatilidade constante resulta em imprecisões quando comparado com os preços observados no mercado. Em contraste, os modelos de Heston e Bates procuram refletir de forma mais fiel as condições reais do mercado.

Descreveremos ambos os modelos com o devido destaque dos seus principais pressupostos e parâmetros. De seguida, o objetivo é calibrar e aplicar ambos os modelos utilizando dados reais de opções sobre o índice S&P500 e, após um análise empírica e devidos testes aos resultados obtidos, determinar se a introdução de jumps no modelo de Bates oferece uma vantagem significativa na determinação dos preços.

Palavras-chave: Modelo de Heston, Modelo de Bates, Volatilidade Estocástica, Jumps.

Abstract

This thesis focuses on the evaluation and comparison of two option pricing models: the Heston model (1993), which incorporates stochastic volatility, and the Bates model (1996), which extends the Heston model by adding jumps to asset prices.

Although the model of Black-Scholes (1973) and Merton (1973) has been widely used for option pricing, his assumption of constant volatility leads to inaccuracies when compared to market-observed prices. In contrast, the Heston and Bates models aim to more accurately reflect real market conditions.

We will describe both models, highlighting their key assumptions and parameters. Following this, the goal is to calibrate and apply both models using real data from S&P 500 index options and, after empirical analysis and relevant tests, determine whether the introduction of jumps in the Bates model provides a significant advantage in price determination.

Keywords: Heston Model, Bates Model, Stochastic Volatility, Jumps.

Contents

CHAPTER 1

Introduction

The world of financial derivatives has evolved significantly over the past few decades, with options playing a crucial role in investment strategies, risk management, and portfolio optimization. The Black-Scholes-Merton (BSM) model, introduced by Black and Scholes (1973) and Merton (1973), revolutionized options pricing with its closed-form solution for European options. However, it assumes constant volatility, which often leads to pricing inaccuracies in real-world scenarios. Empirical evidence from financial markets, particularly the volatility smile and the occurrence of price jumps, has prompted researchers to develop more sophisticated models.

One of the most prominent model to address these limitations is the Heston model (1993), which introduces stochastic volatility, allowing the volatility of the underlying asset to vary over time. This feature makes the model more capable of capturing the complex volatility patterns observed in financial markets. Furthermore, recognizing that asset prices can experience sudden jumps, the Bates (1996) model extends the Heston framework by incorporating both stochastic volatility and jumps, offering a more comprehensive approach to pricing options in volatile markets.

Given the S&P 500 index's pivotal role as a global financial benchmark, this thesis aims to explore the calibration and empirical performance of the Heston and Bates models when applied to S&P 500 index options. Through a comparative analysis, this research seeks to assess the pricing accuracy, ability to capture market dynamics, and practical applicability of both models. A key focus of the study is to evaluate whether the inclusion of jumps in the Bates model offers a significant advantage in option pricing, particularly in capturing sudden market shifts. The findings will provide insights into the limitations and strengths of these models, contributing to both academic literature and informed financial decision-making.

CHAPTER 2

Literature Review

This chapter provides a comprehensive review of some foundational concepts and definitions relevant to this study. We begin by introducing key stochastic processes and mathematical tools necessary for understanding option pricing models, including the Heston and Bates models. These initial definitions set the groundwork for the subsequent discussion on stochastic volatility and jump-diffusion models.

The chapter will then focus on the Heston model, a widely-used stochastic volatility model, followed by an exploration of its limitations, particularly in handling sudden price jumps. To address these limitations, we introduce the Bates model, which extends the Heston framework by incorporating jumps into the asset price dynamics. This extension aims for a more comprehensive representation of real-world market behaviors, particularly during periods of high volatility or significant price shifts.

2.1. Necessary Initial Definitions

Before delving into the specifics of the Heston and Bates models, it is essential to first review some concepts used in their formulation. This section covers key stochastic processes and transforms, such as martingale processes, Wiener processes, Poisson processes, characteristic functions, and the Fourier transform.

2.1.1. Martingale Processes

A *Martingale Process* is a stochastic process where the future expected value of the process is equal to its present value, conditional on all past information. Formally, a stochastic process $M_t, t \geq 0$, defined on a filtered probability space $(\Omega, \mathcal{F}_t, \mathbb{P})$ is called a martingale if it satisfies the following conditions:

(1) $\mathbb{E}[|M_t|] < \infty$ for all $t \geq 0$;
(2) $M_s = \mathbb{E}[M_t|\mathcal{F}_s]$ for all $s \leq t$, where $\mathcal{F}_s$ represents the information available up to time s;
(3) The process M_t is adapted to the filtration $\mathcal{F}_t$, meaning that the value of M_t at time t only depends on information up to that time.

In simpler terms, a martingale is a process where the conditional expected value of future observations, given the past and present, is always equal to the current value. This property makes martingales particularly useful in the modeling of financial markets, where no arbitrage opportunities are present, and price changes are considered to be "fair" over time.

Martingales play a central role in modern probability theory, particularly in the context of financial mathematics. They serve as the foundation for pricing derivative securities and constructing hedging strategies, as the absence of arbitrage in financial markets is closely related to the martingale property.

2.1.2. Wiener Processes

A *Wiener process* or *Brownian motion* is a real-valued stochastic process that plays a key role in modeling continuous-time random behavior in financial markets. Formally, a Wiener process $W_t, t \geq 0$, is defined on a filtered probability space $(\Omega, \mathcal{F}_t, \mathbb{P})$ and satisfies the following conditions:

(1) $W_0 = 0$, almost surely, meaning $\mathbb{P}(W_0 = 0) = 1$;
(2) W has independent increments, which implies that for any increasing sequence of times $t_0, t_1, \ldots, t_n$, the random variables $W_{t_0}, W_{t_1} - W_{t_0}, \ldots, W_{t_n} - W_{t_{n-1}}$ are independent;
(3) W has stationary increments, meaning the distribution of $W_{t+h} - W_t$ depends only on h, and not on t;
(4) $W_t \sim \mathcal{N}(0, t)$, meaning that the increments follow a normal (Gaussian) distribution with mean 0 and variance t.

It is possible to prove that this process is stochastically continuous and that it's sample path (trajectory) is continuous in t. This last property plays a critical role in the behavior of diffusion models. However, as demonstrated by Cont and Tankov (2004), it becomes less robust when jumps are introduced into the dynamics of asset prices. The Poisson process, a fundamental example of a pure jump process, will be introduced next.

2.1.3. Poisson Processes

A *Poisson process* is a fundamental building block in jump-diffusion models, particularly when modeling the random times at which jumps occur.

Let $(\tau_n)_{n \in \mathbb{N}}$ represent a sequence of random times, with each τ_n mapping from the probability space Ω to the positive real line, $\mathbb{R}^+$. The counting process associated with these random times is defined as:

$$N(t) = \sum_{n=0}^{\infty} 1_{\tau_n \leq t}, \tag{2.1}$$

where $1_{\tau_n \leq t}$ is an indicator function, equal to 1 if $\tau_n \leq t$ and 0 otherwise. This process $N(t)$ counts the number of random times smaller or equal to t.

A Poisson process with jump intensity $\lambda > 0$ is characterized by the probability distribution:

$$P[N(t) = k] = \frac{(\lambda t)^k e^{-\lambda t}}{k!}, \quad k \in \mathbb{N}. \tag{2.2}$$

This counting process is generally not a martingale. However, it can be transformed into a martingale by subtracting its expected value at time t, leading to the compensated Poisson process:

$$N_C(t) = N(t) - \mathbb{E}[N(t)] = N(t) - \lambda t. \tag{2.3}$$

To model the jump sizes, we can extend the basic Poisson process by introducing a compound Poisson process. Let $(Y_n)_{n\in\mathbb{N}}$ be a sequence of independent, identically distributed random variables representing the jump sizes, and $N(t)$ as the Poisson process that governs the jump times. The compound Poisson process is then defined as:

$$N_Y(t) = \sum_{n=0}^{N(t)} Y_n. \tag{2.4}$$

This extension allows for flexibility in financial modeling, enabling the inclusion of different jump size distributions, such as normal or exponential distributions. In practice, jump-diffusion models often use a combination of diffusion processes (captured by Brownian motion) and jumps (captured by the Poisson process) to better represent sudden market movements.

2.1.4. Characteristic Function

The characteristic function of a random variable X is defined as:

$$\phi_X(t) = \mathbb{E}[e^{itX}], \tag{2.5}$$

where:

- $\phi_X(t)$ is the characteristic function;
- t is a real variable;
- $\mathbb{E}$ denotes the expected value;
- i is the imaginary unit ($i^2 = -1$);
- X is the random variable.

The characteristic function has the following properties:

- $\phi_X(0) = 1$ and $|\phi_X(t)| \leq 1 \forall t \in \mathbb{R}$;
- The characteristic function always exists and is continuous;
- Random variables with the same characteristic function are identically distributed;
- It is possible to derive the moments of the random variable from $\phi_X(t)$.

When we say a characteristic function is in closed form, we mean that there is a direct analytical expression for $\phi_X(t)$. For example, for a normally distributed random variable $X \sim N(\mu, \sigma^2)$, the characteristic function has the following closed form:

$$\phi_X(t) = e^{i\mu t - \frac{1}{2}\sigma^2 t^2}. \tag{2.6}$$

This expression is direct and does not require integration or infinite sums to be calculated.

Having the characteristic function in closed form allows for efficient and precise option pricing and other financial derivatives evaluations.

- **Computational Efficiency**: Allows for quick calculations by avoiding the need for numerical integrations.
- **Precision**: Minimizes numerical errors that may occur in approximations or integrations.
- **Simplicity**: Facilitates theoretical analysis and algorithm implementation.

2.1.5. The Fourier Transform

The Fourier transform is widely used in mathematical finance, especially for option pricing models. There are several definitions but one of the most common forms encountered in financial literature is the following:

$$\hat{f}(u) = \int_{-\infty}^{\infty} e^{iux} f(x)\, dx, \tag{2.7}$$

where $i = \sqrt{-1}$ represents the imaginary unit. This form of the Fourier transform is frequently used by researchers such as Carr and Madan (1999), Duffie, Pan, and Singleton (2000), and others.

To recover the original function $f(x)$, we use the inverse Fourier transform:

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iux} \hat{f}(u)\, du. \tag{2.8}$$

Differentiation is straightforward with the Fourier transform, as it converts differentiation into multiplication. For example, the Fourier transform of the first derivative $f'(x)$ is given by:

$$\hat{f}'(u) = \int_{-\infty}^{\infty} e^{iux} f'(x)\, dx = -iu\hat{f}(u). \tag{2.9}$$

Repeated application of integration by parts shows that the Fourier transform of the n-th derivative of f is:

$$\hat{f}^{(n)}(u) = (-iu)^n \hat{f}(u). \tag{2.10}$$

This property makes Fourier transforms a powerful tool in the analysis and pricing of financial derivatives.

2.2. Heston Model

In the two decades following its introduction in 1993, the Heston model has become one of the most influential frameworks in the field of option pricing, particularly within the paradigm of stochastic volatility modeling. To understand its prominence, it is necessary to revisit a key event that reshaped financial markets: the stock market crash of October

1987. This event exposed limitations in the Black-Scholes-Merton model, particularly in its inability to handle the volatility skews and smiles that emerged post-crash. The Black-Scholes-Merton framework assumes normally distributed returns with constant volatility, a simplification that proved insufficient as empirical studies revealed more complex behavior in market data.

Market returns often exhibit skewness and kurtosis, characteristics such as 'fat tails', which deviate from normality, and volatility is known to fluctuate over time, typically showing an inverse relationship with stock prices. As a result, many researchers sought to address these shortcomings by incorporating time-varying volatility into their models. One widely adopted approach is to allow volatility to follow its own stochastic process, giving rise to stochastic volatility models, of which the Heston model is a key example. Earlier models by Hull and White (1987), Wiggins (1987), Chensey and Scott (1989), and Stein and Stein (1991) laid the foundation, but the Heston model has emerged as the most important and widely used among them.

What sets the Heston model apart is its ability to generate option prices that reflect skewness and kurtosis, while still allowing for an intuitive relationship between stock prices and volatility. By introducing correlation between the processes driving the stock price and its volatility, the Heston model is capable of reproducing the observed volatility smiles and skews in a manner that aligns with real market conditions. Importantly, the model offers a closed-form solution for option prices, which, although reliant on numerical integration, remains computationally efficient and practical.

Another breakthrough of the Heston model lies in its use of characteristic functions for option pricing. By focusing on the characteristic function of the terminal price distribution rather than the distribution itself, the model initiated a new and more efficient approach to pricing options, a methodology that has since become a cornerstone in financial mathematics. These factors collectively have cemented the Heston model as a benchmark in stochastic volatility modeling, against which other models are often compared.

2.2.1. Assumptions of the Heston Model

The Heston model relies on several fundamental assumptions to simplify the complex nature of option pricing. Below is a detailed look at these key assumptions:

(1) **Geometric Brownian Motion**: The model presumes that the underlying asset's price follows a geometric Brownian motion, meaning that the logarithm of the asset price evolves randomly over time.
(2) **Stochastic Volatility**: A distinctive feature of the Heston model is its treatment of volatility as a stochastic process, where volatility is not fixed but instead fluctuates according to its own random path over time.
(3) **Mean Reversion**: The volatility process in this model includes mean reversion, suggesting that while volatility can deviate from its average level, it will eventually revert to a long-term equilibrium value. This assumption moderates extreme changes in volatility.

(4) **Correlation between Returns and Volatility**: There is an assumed correlation between the asset's returns and volatility changes, where large price movements often coincide with increased volatility. This captures the tendency of volatility to rise during periods of significant market activity.

(5) **No Arbitrage**: The model operates under the no-arbitrage condition, implying that there are no opportunities for risk-free profit through trading strategies involving the asset or its options.

(6) **Constant Interest Rate**: A constant risk-free interest rate is assumed throughout the duration of the option, simplifying the discounting of future cash flows.

(7) **Log-Normal Distribution**: The asset price is assumed to follow a log-normal distribution, consistent with the geometric Brownian motion assumption, and is a common simplification in option pricing.

(8) **Continuous Trading**: The Heston model assumes that trading occurs continuously in the underlying asset, without accounting for any sudden market jumps or discontinuities.

(9) **Efficient Market Hypothesis**: The model presumes that the market is efficient, meaning that all available information is instantly reflected in the asset's price, excluding the possibility of market inefficiencies or behavioral biases.

These assumptions are essential for the mathematical simplicity of the Heston model, but they also come with limitations. While the model has become widely adopted, it is crucial to understand the impact of these assumptions on its applicability to real-world market conditions. As a result, practitioners and researchers often modify or extend the Heston model to address specific market behaviors or data.

2.2.2. Mathematical Formulation of the Heston Model

Heston (1993) devised the stochastic volatility model for option pricing. The model assumes that the underlying stock price, S_t, follows a Black-Scholes-Merton-type stochastic process, but with a stochastic variance, v_t, which is modeled by a Cox, Ingersoll, and Ross (1985) process. Under the physical measure, the stochastic volatility model defines the underlying asset process through the following bivariate system of stochastic differential equations (SDEs):

$$dS_t = \mu S_t \, dt + \sqrt{v_t} S_t \, dW_t^S, \tag{2.11}$$

$$dv_t = \kappa(\theta - v_t) \, dt + \sigma \sqrt{v_t} \, dW_t^v, \tag{2.12}$$

where:

- S_t is the price of the underlying asset at time t;
- v_t is the instantaneous variance at time t;
- μ is the drift of the process for the stock;
- $\kappa > 0$ is the mean reversion speed for the variance;

- $\theta > 0$ is the mean reversion level for the variance;
- $\sigma > 0$ is the volatility of variance (volatility of volatility);
- $v_0 >$ the initial (time zero) level of the variance;
- W_t^S and W_t^v are two correlated Wiener processes with correlation $\rho \in [-1, 1]$, i.e.:

$$dW_t^S dW_t^v = \rho\, dt. \tag{2.13}$$

As shown above, this model treats the asset's volatility as a stochastic process. This approach enables the model to more accurately reflect the behavior of stock returns, which often deviate from a normal distribution. It can account for features such as fat tails and asymmetry in the distribution of returns.

2.2.3. Characteristic Function of the Heston Model

One of the advantages of the Heston Model is that the characteristic function of the logarithm of the asset price can be expressed in closed form, which greatly facilitates the evaluation of options via the Fourier transform. The closed-form characteristic function for the Heston model can be obtained by the next theorem.

THEOREM 2.1. *Let $X_t = \log(S_t)$ represent the logarithm of the price variable. The characteristic function of X_t, as presented by BañoRollin, FerreiroCastilla, and Utzet (2009), is expressed as follows:*

$$\phi_T(u) = e^{A(u)+B(u)+C(u)}, \tag{2.14}$$

where

$$A(u) = iu(X_0 + rT), \tag{2.15}$$

$$B(u) = -\frac{(u^2 + iu)(1 - e^{\phi(u)T})v_0}{2\phi(u) - (\phi(u) - \tau(u))(1 - e^{-\phi(u)T})}, \tag{2.16}$$

$$C(u) = -\frac{k\theta}{\xi^2}\left[2\log\left(\frac{2\phi(u) - (\phi(u) - \tau(u))(1 - e^{-\phi(u)T})}{2\phi(u)}\right) + (\phi(u) - \tau(u))T\right], \tag{2.17}$$

with $\tau(u) = k - iu\rho\xi$ and $\phi(u) = \sqrt{\tau(u)^2 + \xi^2(u^2 + iu)}$.

2.2.4. Limitations and Critiques of the Heston Model

The Heston model, despite its significant contributions to option pricing, is not without its limitations. One primary critique lies in the complexity of calibrating the model's parameters, which can be computationally intensive and prone to errors, particularly when dealing with limited or noisy market data. Furthermore, the Heston model assumes a constant mean reversion speed for volatility, which may not accurately capture the variability in real-world market conditions.

Another major drawback is the model's failure to account for jumps or extreme events in asset prices, as it assumes continuous price movements and a log-normal distribution. This limits the model's ability to capture real market phenomena such as sudden shocks, which leads to inaccuracies in pricing options during times of high volatility or market stress. Additionally, while the model assumes a constant risk-free interest rate, this does not reflect the reality of fluctuating interest rates in financial markets.

In terms of practical application, the Heston model can be computationally demanding, requiring significant resources for simulations, especially when applied to complex financial instruments or for real-time decision-making. Moreover, the absence of closed-form solutions in many cases necessitates the use of numerical methods, further complicating its practical use.

Given these limitations, this thesis will now introduce a variation of the Heston model, known as the Bates model, which incorporates jumps in asset prices. The Bates model seeks to address some of the shortcomings of the Heston model by allowing for discontinuities in asset prices, which can better capture extreme market events and enhance the accuracy of option pricing under certain market conditions

2.3. Introduction of Jumps in the Heston Model

Having already discussed the limitations of the Heston model, particularly its inability to capture extreme events due to the assumption of continuous asset price movements, we now turn to an extension that addresses these shortcomings. While the Heston model incorporates stochastic volatility, it assumes that asset prices move continuously, which may not adequately reflect the reality of rare events such as market crashes or sharp price swings. These events result in fat tails and asymmetries in return distributions, which the Heston model alone struggles to model accurately. To overcome this limitation, we will now introduce the Bates model.

2.3.1. The Bates Model

The Bates model extends the Heston model by incorporating a jump process modeled as a compound Poisson process, that we defined in chapter 2, into the asset price dynamics. This extension is designed to address one of the key limitations of the Heston model: its inability to capture sudden and discrete price movements, such as market crashes or significant downturns, which are common in real financial markets. These jumps are often triggered by unexpected economic, political, or geopolitical events, leading to sharp deviations from normal price behavior that the continuous framework of the Heston model alone cannot adequately represent.

By combining stochastic volatility with a jump component, the Bates model offers greater flexibility in modeling both the continuous fluctuations in asset prices and the occurrence of extreme events. This enhanced framework allows the Bates model to more accurately capture the skewness and fat tails observed in the distribution of asset returns, features that are often reflected in the implied volatility surface. As a result, the Bates model is particularly effective in improving the pricing of options, especially in markets where large price swings and volatility spikes are more prevalent.

In the following sections, we will explore the mathematical formulation of the Bates model in detail, illustrating how the addition of jumps to the stochastic volatility process aims to enhance its ability to model real-world financial data and improves upon the Heston model's option pricing capabilities.

2.3.2. Mathematical Formulation of the Bates Model

The dynamics of the underlying asset price S_t in the Bates model are governed by the following stochastic differential equations (SDEs):

$$dS_t = \mu S_t\, dt + \sqrt{v_t} S_t\, dW_t^S + (Y-1)S_t - dN_t, \tag{2.18}$$

$$dv_t = \kappa(\theta - v_t)\, dt + \sigma\sqrt{v_t}\, dW_t^v, \tag{2.19}$$

where:

- $S(t)$ is the price of the underlying asset at time t;
- μ is the rate of return of the asset;
- $V(t)$ is the stochastic volatility process, which follows the second SDE;
- κ is the speed of mean reversion of the variance;
- θ is the long-term variance;
- ν is the volatility of volatility (vol of vol);
- $W_1(t)$ and $W_2(t)$ are two correlated Wiener processes with correlation ρ;
- $dN(t)$ is a Poisson process capturing the occurrence of jumps;
- Y is the jump size, which follows a log-normal distribution with parameters μ_J (mean jump size) and σ_J (jump volatility);
- λ is the jump intensity, representing the average frequency of jumps over time.

2.3.3. Characteristic Function of the Bates Model

The characteristic function of the Bates model combines the characteristic function of the Heston model, which describes the stochastic volatility, with a term that captures the jump distribution, modeled as a Poisson process with log-normal jumps.

The characteristic function of the Bates model is expressed as:

$$\phi_{Bates}(u,t) = \exp\left(\phi_{Heston}(u,t) + \lambda\left(\phi_{Jump}(u) - 1\right)T\right), \tag{2.20}$$

where:

- $\phi_{Heston}(u,t)$ is the characteristic function of the Heston model;
- $\phi_{Jump}(u)$ is the jump characteristic function, given by:

$$\phi_{Jump}(u) = \exp\left(iu\mu_J - \frac{1}{2}\sigma_J^2 u^2\right). \tag{2.21}$$

Here, μ_J represents the average jump size, while σ_J is the volatility associated with jumps. The term λ represents the intensity of the Poisson process, describing the expected frequency of jumps within the time interval T.

CHAPTER 3

Methodology

3.1. Data Preparation

In order to compare the performance of the Bates and Heston models, S&P500 index call option prices were utilized. The market data for the S&P500 index options, traded on exchanges such as the Chicago Board Options Exchange and the New York Stock Exchange Arca Options, was sourced from the Refinitiv database. The analysis was conducted using data from specific days from February 2021 to February 2022. We calibrated the models with data from the days 2021-02-03, 2021-05-03, 2021-08-03, 2021-11-03 and 2022-02-03 in order to later forecast the days 2021-02-04, 2021-05-04, 2021-08-04, 2021-11-04 and 2022-02-04, respectively.

Moneyness is defined as the percentage difference between the current underlying price and the strike price, calculated as:

$$\text{Moneyness}(\%) = \frac{S}{K}.$$

For this analysis, moneyness is categorized as follows: OTM if $S/K < 0.95$, ATM if $0.95 \leq S/K < 1.05$, and ITM if $S/K \geq 1.05$. Options that were significantly deep in-the-money or deep out-of-the-money are less liquid and their market prices may deviate substantially from their true values. Specifically, an option is classified as very deep in-the-money when its moneyness exceeds 12%, and very deep out-of-the-money when its moneyness falls below -12%. We excluded these type of options. Additionally, options with fewer than 7 days or more than 180 days until expiration were excluded due to their high sensitivity to liquidity-related biases.

Table 1 provides a summary of the in-sample properties of the S&P500 index options used in this study, including the average call prices, standard deviation of the prices, the total number of observations for each financial metric, and the time to maturity.

TABLE 1.
Properties of the S&P500 index call options in-sample. The annotated numbers represent the average closing prices of the call options, the standard deviation of the prices {provided in braces}, and the total number of observations (given in parentheses) for each category of moneyness and time to maturity. S refers to the spot level of the S&P500 index, and K denotes the strike price. ITM, ATM, and OTM correspond to in-the-money, at-the-money, and out-of-the-money options, respectively.

Moneyness (S/K)	$\tau < 1$ month	$1 \leq \tau < 3$ months	$\tau \geq 3$ months
OTM < 0.95	0.4174 {0.6029} (197)	8.0655 {8.4699} (404)	36.5964 {22.4696} (488)
ATM $[0.95, 1.05[$	69.0310 {68.0045} (435)	114.3557 {76.3940} (769)	180.9689 {79.7588} (883)
ITM ≥ 1.05	339.4929 {74.1654} (255)	373.4587 {70.4385} (452)	427.4420 {66.4807} (523)

3.2. Procedures

We used the implementation of the Heston and Bates models for pricing S&P 500 index options using the Attari (2004) approach. This method provides an alternative to traditional Fourier-based pricing models by reducing the calculation to a single numerical integration, which results in a significant improvement in computational efficiency while maintaining accuracy. Attari's method builds upon earlier work by Lewis (2001), offering a more straightforward approach for computing option prices by leveraging the characteristic function of the underlying asset's log return.

3.2.1. Attari (2004) Representation

The terminal stock price under the risk-neutral measure will be expressed as:

$$S_T = S_t e^{r\tau + x(t,T)},$$

where $x(t,T)$ denotes the stochastic component of the stock price process. The call price is given by:

$$C(K) = e^{-r\tau}\mathbb{E}^Q[S_T|S_T > K] - Ke^{-r\tau}\mathbb{E}^Q[1_{S_T>K}]. \tag{3.1}$$

For the sake of simplicity, we assume a zero dividend yield. The formula (3.1) can be rewritten in terms of expectations of the stochastic variable x, which leads to:

$$C(K) = e^{-r\tau} E^Q[S_T | S_T > K] - K e^{-r\tau} E^Q[1_{S_T > K}]$$

$$= S_t E^Q[e^x | x > \zeta] - K e^{-r\tau} E^Q[1_{x>\zeta}] \tag{3.2}$$

$$= S_t \Phi_1 - K e^{-r\tau} \Phi_2,$$

where $\zeta = \ln\left(\frac{Ke^{-r\tau}}{S_t}\right)$. The two expectations, $\mathbb{E}_Q[e^x | x > \zeta]$ and $\mathbb{E}_Q[1_{x>\zeta}]$, are taken under the risk-neutral density for x, $q(x)$. These expectations lead to two probability terms:

$$\Phi_1 = \mathbb{E}^Q[e^x | x > \zeta] = \int_\zeta^\infty e^x q(x)\, dx = \int_\zeta^\infty p(x)\, dx. \tag{3.3}$$

$$\Phi_2 = \mathbb{E}^Q[1_{x>\zeta}] = \int_\zeta^\infty q(x)\, dx.$$

Since $e^x q(x) > 0$ and $0 \leq \Phi_1 \leq 1$, then $p(x) = e^x q(x)$ can be the density function. The characteristic functions for $q(x)$ and $p(x)$ are denoted as $\varphi_2(u)$ and $\varphi_1(u)$, respectively. Using the relationship between these characteristic functions, we have:

$$\varphi_1(u) = \int_{-\infty}^\infty e^{iux} p(x)\, dx = \int_{-\infty}^\infty e^{iux} e^x q(x)\, dx = \int_{-\infty}^\infty e^{i(u-i)x} q(x)\, dx = \varphi_2(u-i). \tag{3.4}$$

This allows the computation of the call price to be transformed into integrals involving the characteristic function. Specifically, having in consideration that we can express $\Pr(\ln S_T > k)$, where $k = Ln(K)$, using the density:

$$\Pr(\ln S_T > k) = \int_k^\infty f(x)dx = \frac{1}{2\pi} \int_k^\infty \left(\int_{-\infty}^\infty e^{-iux} \varphi(u) du \right) dx$$
$$= \frac{1}{2\pi} \int_{-\infty}^\infty \varphi(u) \left(\int_k^\infty e^{-iux} dx \right) du.$$

The first expectation Φ_1 can then be written as:

$$\Phi_1 = \frac{1}{2\pi} \int_{-\infty}^\infty \varphi_1(v) \left(\int_\zeta^\infty e^{-ivx}\, dx \right) dv = \frac{1}{2\pi} \int_{-\infty}^\infty \varphi_2(v-i) \left(\int_\zeta^\infty e^{-ivx}\, dx \right) dv. \tag{3.5}$$

Substitute the change of variable $u = v - i$ in function (3.5):

$$\Phi_1 = \frac{1}{2\pi} \int_{-\infty}^\infty \varphi_2(u) \left(\int_\zeta^\infty e^{-i(u+i)x}\, dx \right) du. \tag{3.6}$$

By performing a change of variables and simplifying the inner integral, this becomes:

$$\Phi_1 = \frac{e^\zeta}{2\pi} \int_{-\infty}^\infty \frac{\varphi_2(u) e^{-iu\zeta}}{i(u+i)}\, du + 1. \tag{3.7}$$

For the second expectation Φ_2, we use the Gil-Pelaez (1951) formula

$$\Pr(\ln S_T > k) = \frac{1}{2} + \frac{1}{2\pi}\int_{-\infty}^{\infty} \varphi(u)\frac{e^{-iuk}}{iu}du,$$

which leads to:

$$\Phi_2 = \frac{1}{2} + \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{\varphi_2(u)e^{-iu\zeta}}{iu}\,du. \tag{3.8}$$

Substituting these into the call price (3.2) equation results in:

$$C(K) = S_t\left[1 + \frac{e^{\zeta}}{2\pi}\int_{-\infty}^{\infty} \frac{\varphi_2(u)e^{-iu\zeta}}{i(u+i)}\,du\right] - Ke^{-r\tau}\left[\frac{1}{2} + \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{\varphi_2(u)e^{-iu\zeta}}{iu}\,du\right]. \tag{3.9}$$

If we substitute $\zeta = \ln\left(\frac{Ke^{-r\tau}}{S_t}\right)$ and use the fact that we only need to consider the real part of the integrals, we obtain:

$$C(K) = S_t - \frac{Ke^{-r\tau}}{2} - \frac{Ke^{-r\tau}}{\pi}\int_0^{\infty} \Re\left[\varphi_2(u)e^{-iu\zeta}\left(\frac{1}{iu} - \frac{1}{i(u+i)}\right)\right]du. \tag{3.10}$$

In the bracketed term of the last integral in Equation (3.37), multiply the second fraction by $u - i$ in the numerator and denominator. The integrand becomes

$$\Re\left[\varphi_2(u)e^{-iu\zeta}\left(\frac{1-i/u}{u^2+1}\right)\right]. \tag{3.11}$$

Now expand $\varphi_2(u) = R_2(u) + iI_2(u)$, where $R_2(u)$ and $I_2(u)$ are the real and imaginary parts of $\varphi_2(u)$, respectively, and expand $e^{-iu\zeta} = \cos(u\zeta) - i\sin(u\zeta)$. Substitute into the integrand (3.38) and regroup the real terms. The integrand becomes

$$A(u) = \frac{R_2(u) + I_2(u)/u}{1+u^2}\cos(u\zeta) + \frac{I_2(u) - R_2(u)/u}{1+u^2}\sin(u\zeta). \tag{3.12}$$

Attari's (2004) formula for the call price is, therefore:

$$C(K) = S_t - \frac{1}{2}Ke^{-r\tau} - Ke^{-r\tau}\frac{1}{\pi}\int_0^{\infty} A(u)\,du. \tag{3.13}$$

This formulation is particularly efficient because it reduces the problem of pricing a European call option to a single numerical integration, which can be handled effectively using numerical methods like Gauss-Laguerre quadrature.

3.2.2. Numerical Integration and Implementation

The implementation was performed using *Python*, with adaptations based on code provided by Rouah (2013) and by Kienitz and Wetterau (2012). For both the Heston and Bates models, the integral formulations derived above were computed using *Gauss-Laguerre quadrature*, a numerical method that is particularly well-suited for handling integrals over infinite domains. The integrand decays rapidly as u increases, allowing the integral to be truncated efficiently without losing accuracy.

3.2.3. Calibration

Both models were calibrated using S&P 500 Index call option data, with the calibration process aimed at minimizing the Root Mean Squared Error (RMSE). It is calculated as follows:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(P_{\text{observed},i} - P_{\text{predicted},i}\right)^2}, \tag{3.1}$$

where:

- $P_{\text{observed},i}$ is the observed option price;
- $P_{\text{predicted},i}$ is the predicted option price from the model;
- n is the total number of observations.

3.3. Forecasting Accuracy

To assess whether there were significant differences in the forecasting accuracy between the models, Heston and Bates, we conducted statistical tests. Specifically, we compared the mean squared errors of the forecasts using two well-known methods: The Diebold-Mariano (DM) test and the Wilcoxon Signed-Rank (WS) test (Diebold and Mariano, 1995). These tests allowed us to evaluate if the observed differences in forecasting performance of both models were statistically significant.

3.3.1. Diabold Mariano Test

The Diebold-Mariano (DM) test is used to evaluate whether there are significant differences in the predictive accuracy of two competing forecasting models.

Following Park, Kim, and Lee (2014), we define the series of forecasting errors as follows:

$$e_{t,i} = \hat{y}_{t,i} - y_t \quad \text{for} \quad i = 1, 2,$$

where y_t represents the actual option prices of the time series and $\hat{y}_{t,i}$ corresponds to the forecasted values. To compare the forecasting performance of the two models, we calculate the difference in the errors using a loss function $h(.)$. We will be using the squared-error loss function which gives us:

$$d_t = h(e_{t,1}) - h(e_{t,2}) = e_{t,1}^2 - e_{t,2}^2.$$

The null hypothesis of the DM test is:

$$H_0 : E(d_t) = 0,$$

implying that there is no significant difference between the forecasting errors of the two models. The test statistic is computed as:

$$DM = \frac{d}{\sqrt{2\pi\hat{f}_d(0)/T}},$$

where $d = \sum_{t=1}^{T}(h(e_{t,1}) - h(e_{t,2}))/T$ and $\hat{f}_d(0)$ is the spectral density of d_t.

3.3.2. Wilcoxon Signed Tank Test

The Wilcoxon Signed Rank Test is a non-parametric statistical test used to compare paired observations. The null hypothesis of the test is:

$$H_0 : \text{Median}(d_t) = 0,$$

where d_t represents the difference in forecasting errors at time t. The Wilcoxon test is based on the ranks of the differences, testing whether the median of these differences is significantly different from zero. The test statistic is given by:

$$WS = \sum_{t=1}^{T} I^+(d_t)\text{Rank}(|d_t|),$$

where:

$$I^+(d_t) = \begin{cases} 1 & \text{if}\, d_t > 0 \\ 0 & \text{otherwise} \end{cases}.$$

CHAPTER 4

Results

This chapter outlines the findings from the comparative analysis between the Heston and Bates models. It covers a detailed discussion of the performance evaluation metrics, and the statistical tests employed to assess the effectiveness of both models.

4.1. Forecasting Performance

Table 2 presents the RMSE values obtained on the forecasts done for each model, categorized by sample groups. The columns are organized by 'days to expiration,' and moneyness is represented as S/K. The best performance in each category is highlighted in **bold**.

TABLE 2. RMSE values obtained forecasting with each model, categorized by moneyness and time-to-maturity.

Moneyness (S/K)	**Model**	$\tau < 1$ month	$1 \leq \tau < 3$ months	$\tau \geq 3$ months
< 0.95	Heston	**0.5796**	**1.5500**	4.0715
	Bates	0.8539	1.6139	**4.0011**
$[0.95, 1.05[$	Heston	5.9338	**6.3584**	6.9336
	Bates	**5.9047**	6.4258	**6.9252**
≥ 1.05	Heston	2.7480	5.0031	6.2874
	Bates	**2.4958**	**4.9149**	**6.2419**

For **out-of-the-money (OTM)** options ($S/K < 0.95$), the Heston model demonstrates superior accuracy for shorter maturities, achieving a lower RMSE of 0.5796 compared to the Bates model's 0.8539. For maturities between 1 and 3 months, the Heston model continues to outperform with an RMSE of 1.5500, while the Bates model's RMSE is 1.6139. For longer maturities (above 3 months), the Bates model shows a slight edge, with an RMSE of 4.0011 versus 4.0715 for the Heston model. These results suggest that the Heston model is more accurate in pricing OTM options for shorter maturities, while the Bates model becomes more effective for extended periods.

For **at-the-money (ATM)** options ($0.95 \leq S/K < 1.05$), both models show a comparable performance for maturities shorter than one month, with the Bates model delivering a slightly lower RMSE of 5.9047 for compared to 5.9338 for the Heston model. In the 1 to 3-month maturity range, the Heston model marginally outperforms with an RMSE of 6.3584, while the Bates model follows closely with an RMSE of 6.4258. For longer maturities (over 3 months), both models perform similarly, with RMSE values of 6.9336 and 6.9252 for the Heston and Bates models, respectively. These results indicate that for ATM options, both models perform similarly across all maturities, with a slight advantage for the Heston model in the 1 to 3-month maturity range.

For **in-the-money (ITM)** options ($S/K \geq 1.05$), the Bates model consistently outperforms the Heston model across all maturity categories. For maturities under one month, the Bates model records a lower RMSE of 2.4958, compared to 2.7480 for the Heston model. In the 1 to 3-month range, the Bates model again demonstrates superior accuracy with an RMSE of 4.9149, compared to the Heston model's RMSE of 5.0031. For longer maturities, the Bates model maintains its lead, showing an RMSE of 6.2419 versus 6.2874 for the Heston model. This suggests that the Bates model is a stronger choice for pricing ITM options.

In summary, these results suggest that the **Heston model** generally provides better accuracy for **OTM** options at shorter maturities, while the **Bates model** is more effective for longer maturities in these options. For **ATM** options, both models perform similarly with minor variations, and the **Bates model** shows a consistent advantage for **ITM** options, regardless of maturity.

4.1.1. Results of the Statistical Tests

Table 3 presents the comparison statistics for pricing errors between the Heston and Bates models. It provides the p-values at the 95% significance level, indicating whether there is a statistically significant difference between the two models. The DM-test and WS-test refer to the p-values obtained from the Diebold-Mariano and Wilcoxon Signed-Rank tests, respectively. Significant results at $\alpha = 0.05$ are highlighted in **bold**.

TABLE 3. Comparison statistics for pricing errors. The table provides p-values at the 95% significance level to evaluate the difference between the Heston and Bates models.

Moneyness (S/K)	**Test**	$\tau < 1$ month	$1 \leq \tau < 3$ months	$\tau \geq 3$ months
< 0.95	DM	**0.0000**	**0.0494**	**0.0030**
	WS	**0.0000**	**0.0347**	**0.0000**
$[0.95, 1.05[$	DM	0.2272	**0.0000**	0.2449
	WS	**0.0093**	**0.0002**	**0.0000**
≥ 1.05	DM	**0.0000**	**0.0000**	**0.0000**
	WS	**0.0000**	**0.0000**	**0.0000**

The results of both the Diebold-Mariano (DM) test and the Wilcoxon signed-rank (WS) test provide statistical evidence of significant differences in forecasting performance between the Heston and Bates models across various categories of moneyness and time-to-maturity.

For **OTM options** ($S/K < 0.95$), the DM test shows highly significant results for maturities below 1 month and above 3 months, with p-values of 0.0000 and 0.0030, respectively. For the maturity range of 1 to 3 months, the p-value is 0.0494, which is just within the significance threshold, suggesting a significant difference less marked. The WS test also reveals significance across all maturities, with p-values of 0.0000 for maturities shorter than 1 month and above 3 months, and a p-value of 0.0347 for the 1 to 3 months range. These results indicate that for OTM options, both tests suggest statistically significant differences between the Heston and Bates models for all maturities, although the level of significance is less pronounced in the intermediate maturity range.

For **ATM options** ($0.95 \leq S/K < 1.05$), the DM test shows a significant difference only for maturities between 1 and 3 months, with a p-value of 0.0000. For shorter (below 1 month) and longer (above 3 months) maturities, the p-values are 0.2272 and 0.2449, respectively, which are above the 0.05 threshold, indicating no statistically significant difference between the models in these categories. Conversely, the WS test shows a significant difference for all maturities, with p-values of 0.0093, 0.0002, and 0.0000 for the three maturity intervals, respectively. This implies that for ATM options, the DM test identifies significant differences only in the intermediate maturity range, while the WS test suggests significant differences across all maturities.

For **ITM options** ($S/K \geq 1.05$), both the DM and WS tests indicate significant differences across all maturity categories, with p-values of 0.0000 in each case. This suggests a highly significant difference between the Heston and Bates models for ITM options, regardless of time-to-maturity.

In summary, the results highlight considerable differences in the forecasting performance of the Heston and Bates models, particularly for OTM and ITM options.

CHAPTER 5

Conclusions

In conclusion, the empirical comparison between the Heston and Bates models in the context of S&P 500 index options has provided important insights into the pricing dynamics of these sophisticated financial instruments. This thesis has thoroughly explored both models, highlighting their relative strengths and weaknesses through a detailed analysis of real market data.

Our findings indicate that the **Heston model** demonstrates better performance in capturing the pricing dynamics of out-of-the-money (OTM) options with small maturities. The model's ability to account for time-varying volatility allows for a more accurate representation of option prices, particularly during periods of increased volatility. However, despite its effectiveness, the complexity and computational demands of the Heston model remain a key consideration, especially when used in real-time trading or risk management applications.

The introduction of jumps in the **Bates model** was intended to enhance pricing accuracy by incorporating sudden market shifts. The results suggest that the Bates model provides an improvement for in-the-money (ITM) options, particularly for short-term maturities. However, for at-the-money (ATM) options, the Bates model does not appear to offer significant advantages over the Heston model, suggesting that the added complexity of jump modeling may only yield meaningful benefits in markets experiencing extreme volatility or frequent abrupt price movements.

Additionally, this research has underscored the importance of accurate calibration and parameter estimation for both models. Poorly calibrated parameters can lead to significant pricing errors, diminishing the theoretical advantages of these advanced models. Moreover, the computational efficiency of the numerical methods used to solve the models' equations can substantially impact their practical applicability, especially in high-frequency trading environments.

In summary, the choice between the Heston and Bates models largely depends on the specific requirements of market participants. While the Bates model may be advantageous in markets with frequent large jumps, the Heston model remains a robust and computationally efficient choice for most practical applications. Investors, risk managers, and traders should carefully evaluate the trade-offs between model complexity and performance, selecting the approach that best aligns with the market environment and the characteristics of the options being priced.

References

Attari, M. (2004), "Option Pricing Using Fourier Transforms: A Numerically Efficient Simplification", *Working Paper*, Charles River Associates, Boston, MA, *https://ssrn.com/abstract=520042.*

BanõRollin, S., FerreiroCastilla, A., & Utzet, F. (2009), "A New Look at the Heston Characteristic Function", *Centre de Recerca Matemàtica*.

Bates, D. S. (1996), "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options", *The Review of Financial Studies*, 9(1), 69-107.

Black, F., & Scholes, M. (1973), "The Pricing of Options and Corporate Liabilities", *The Journal of Political Economy*, 81(3), 637-654.

Carr, P., & Madan, D. (1999), "Option Valuation Using the Fast Fourier Transform", *Journal of Computational Finance*, 2(4), 61–73.

Chesney, M., & Scott, L. (1989), "Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model", *Journal of Financial and Quantitative Analysis*, 23(3), 267-284.

Cont, R., & Tankov, P. (2004), "Financial Modelling with Jump Processes", *Chapman and Hall*, Boca Raton, FL.

Cox, J.C., Ingersoll, J.E., & Ross, S.A. (1985), "A Theory of the Term Structure of Interest Rates", *Econometrica*, 53(2), 385–408.

Diebold, F. X., & Mariano, R. S. (1995), "Comparative Predictive Accuracy", *Journal of Business and Economic Statistics*, 20(1), 134-144.

Duffie, D., Pan, J., & Singleton, K. (2000), "Transform Analysis and Asset Pricing for Affine Jump Diffusions", *Journal of Computational Finance*, 68(6), 1343–1376.

Gil-Pelaez, J. (1951), "Note on the Inversion Theorem", *Biometrika*, 38(3-4), 481–482.

Heston, S. (1993), "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", *The Review of Financial Studies*, 6(2), 327-343.

Hull, J., & White, A. (1987), "The Pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance*, 42(2), 281-300.

Kienitz, J., & Wetterau, D. (2012), *Financial Modelling: Theory, Implementation and Practice with MATLAB Source*, Wiley Finance Series, John Wiley & Sons Ltd., Chichester, United Kingdom.

Lewis, A.L. (2001), "A Simple Option Formula for General Jump-Diffusion and Other Exponential Levy Processes", *https://ssrn.com/abstract=282110.*

Merton, R. (1973), "The Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science*, 4(1), 141-183.

Park, H., Kim, N., & Lee, J. (2014), "Parametric Models and Non-Parametric Machine Learning Models for Predicting Option Prices: Empirical Comparison Study over KOSPI 200 Index Options", *Expert Systems with Applications*, 41(11), 5227-5237.

Rouah, F. D. (2013), "The Heston Model and Its Extensions in Matlab and C#", *Wiley Finance Series*, John Wiley & Sons, Inc., Hoboken, New Jersey.

Stein, E., & Stein, J. (1991), "Stock Price Distributions with Stochastic Volatility: An Analytic Approach", *The Review of Financial Studies*, 4(4), 727-752.

Wiggins, J. (1987), “Option Values under Stochastic Volatility: Theory and Empirical Estimates”, *Journal of Financial Economics*, 19(2), 351-372.