Robust Multi-Layer Calibration of the Heston Stochastic Volatility Model

The Balanced Premium Calibration Method

Mohammad Aldossari

Submitted for the degree of

Doctor of Philosophy



Heriot-Watt University

Department of Actuarial Mathematics and Statistics

School of Mathematical and Computer Sciences

September 9, 2026

The copyright in this thesis is owned by the author. Any quotation from the thesis or use of any of the information contained in it must acknowledge this thesis as the source of the quotation or information.

Abstract

This thesis presents the balanced premium calibration method (BPCM), a three-layer framework for robustly fitting the Heston stochastic volatility model to large option datasets. The BPCM method involves three layers. Layer 1 ensures market consistency by filtering and structurally repairing raw quotes via put–call parity and bid–ask bounds, separating stable observations from noise. Layer 2 performs daily least-squares calibration of the Heston parameters using closed-form characteristic function pricing and derived analytic gradients and Hessians, thereby achieving rapid convergence without finite-difference approximations. Layer 3 redistributes errors and allows for controlled adjustments to model inputs and outputs, absorbing residual pricing errors and restoring arbitrage-free consistency. Working with 1.5 million call and put quotes on a major equity from 2018 to 2024, BPCM ensured that model prices closely adhere to market bid–ask spreads (91.58% adherence) for stable regimes while maintaining realistic spot price behaviour. The calibrated model achieves high consistency with observed prices and reconstructs the underlying spot price trajectory with minimal deviation even during market crises. Moreover, residual error and correlation analyses reveal structural differences between stable and unstable data regimes. In the stable regime, implied option Greeks and calibration gradients behave smoothly and expectedly. These findings confirm that BPCM delivers reliable calibration under normal conditions (consistent with classical theory) and expose predictable model stress under turbulent market regimes.

In addition to BPCM, this thesis derives explicit closed-form expressions for the Heston model's Hessians. To the author's knowledge, these second-order derivatives lay the groundwork for future research in second-order optimisation techniques.

Keywords: Heston model • stochastic volatility • robust calibration • analytic derivatives • multi-layer refinement • residual diagnostics • arbitrage-free data engineering • quantitative finance

Acknowledgements

I would like to express my deepest gratitude to my supervisory team. I am particularly grateful to my first supervisor, Malham Simon, for his continuous encouragement, intelligent guidance, and generous patience during this research journey. I am deeply appreciative of my second supervisor, Anke Wiese, for her knowledge and support, which were useful in developing this work to its final version.

When I began this PhD, I came from a modest academic background in finance. Through their dedicated mentorship, my supervisors not only supported me technically but also helped raise my academic abilities to a level I could not have imagined at the outset. Their belief in my potential gave me the confidence to overcome challenges and grow as an independent researcher.

I would also like to thank the Heriot-Watt University for providing an outstanding research environment and student support services that truly empower postgraduate researchers. From academic resources to technical and personal assistance, the infrastructure and atmosphere have played an essential role in enabling my development. My appreciation also extends to the Mathematics Department for their support and for fostering an intellectually rich environment that complemented my research in financial modelling.

I am profoundly grateful to Umm Al-Qura University for sponsoring my doctoral studies. I would also like to acknowledge the support of the Saudi Arabian Cultural Bureau in the UK and the Ministry of Education, Kingdom of Saudi Arabia for facilitating this sponsorship. Their investment in education and research has made this journey possible.

Finally, I am deeply thankful to my wife, Rafah Aldosari, whose unwavering support, patience, and belief in me have sustained me throughout this journey. Her encouragement behind the scenes has been invaluable. I am also profoundly grateful to my children, Ibrahim and Bina, whose pure hearts and natural disposition have reminded me of the essence of balance, sincerity, and presence. Through them, I have rediscovered aspects of my own nature that have quietly supported and grounded me throughout this work.

The thesis text has been reviewed using Grammarly for grammar and clarity, and QuillBot for phrasing suggestions, without compromising the technical accuracy or academic tone of the original content.

Contents

List of Figures vii

List of Tables viii

List of Algorithms ix

1 Introduction 1
- 1.1 Option Pricing and the Heston Model . . . 1
- 1.2 Challenges in Model Calibration . . . 2
- 1.3 Data Engineering for Effective Model Calibration . . . 2
- 1.4 Limitations of Existing Calibration Approaches . . . 3
- 1.5 Contributions . . . 3

2 Option Price Model 6
- 2.1 Motivation and Model Context . . . 6
- 2.2 Black–Scholes Model . . . 7
- 2.3 Heston Model . . . 8
- 2.4 Reformulated Characteristic Function . . . 15

3 Calibration Methods 18
- 3.1 Introduction . . . 18
- 3.2 Heston Calibration Stability . . . 19
- 3.3 Residual Analysis in Financial Modelling . . . 19
- 3.4 Problem Statement . . . 20
- 3.5 Analytic Derivative Computation for Calibration . . . 22
- 3.6 Analytical Gradient Derivation . . . 25
- 3.7 Cross Derivatives Implementation and Interpretation . . . 30
- 3.8 Derivatives with respect to r, S_0, and τ . . . 34

4 Data Collection and Preparation 43

4.1 Introduction and Data Overview 43
4.2 Supplementary Data Inputs 44
4.3 Data Challenges and Market Microstructure 45
4.4 Data Structure and Challenges 47

5 Balanced Premium Calibration Method 50
5.1 Introduction 50
5.2 The BPCM Framework 52
5.3 Implementation Details 59
5.4 Calibration Accuracy and Robustness 66
5.5 Computational Cost Analysis 73
5.6 Empirical Performance Overview 75
5.7 Interest Rate and Dividend Yield Adjustments 84
5.8 Gradient Dynamics of Heston Parameters 85
5.9 Option Greeks Under Stable and Unstable Regimes 89
5.10 Correlation Analysis Across Market Regimes 94

6 Conclusion and Future Research Directions 98
6.1 Summary and Contextualisation of Results 98
6.2 Novel Contributions 99
6.3 Limitations of BPCM 100
6.4 Future Research Directions 101

Bibliography 102

Appendices 113

A Derivation of the Heston Option Pricing PDE 113
A.1 Application of Itô's Lemma to the Heston Model 113
A.2 Log-Transformation of the Heston PDE 116

B Derivation of the Heston Model's Characteristic Function 119
B.1 Derivation of the Characteristic Function for the Heston Model 119
B.2 Solutions for the Riccati Equations $D(\tau)$ and $C(\tau)$ 121

C Second Derivatives of the Characteristic Function 130
C.1 Second Derivatives of φ with Respect to v_0 130
C.2 Second Derivatives of φ with Respect to ρ 131

C.3 Second Derivatives of φ with Respect to κ . . . 134
C.4 Second Derivatives of φ with Respect to σ . . . 141

D Derivatives with Respect to Time-to-Maturity 145

E Cross Derivatives of the Heston Characteristic Function 149
E.1 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial v_0 \, \partial \bar{v}}$ . . . 149
E.2 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial v_0 \, \partial \rho}$ . . . 150
E.3 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial v_0 \, \partial \kappa}$ . . . 152
E.4 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial v_0 \, \partial \sigma}$ . . . 153
E.5 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \bar{v} \, \partial \rho}$ . . . 154
E.6 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \bar{v} \, \partial \kappa}$ . . . 156
E.7 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \bar{v} \, \partial \sigma}$ . . . 157
E.8 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \rho \, \partial \kappa}$ . . . 159
E.9 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \rho \, \partial \sigma}$ . . . 161
E.10 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \kappa \, \partial \sigma}$ . . . 162

F Supplementary Data Tables 164
F.1 Correlation Analysis Tables . . . 164
F.2 Event-Based Context for Residual Spikes . . . 171

List of Figures

4.1 Density of implied volatility vs. log-moneyness. 44

5.1 Quote assignment relative to bid–ask spreads under different calibration methods. 69
5.2 Median absolute residuals in 5-day bins. 71
5.3 Fit diagnostics over 2018–2024. 72
5.4 Intra-day implied spot S_t over 6–11 May 2024. 76
5.5 Daily implied spot S_t, 2018–2024. 77
5.6 Implied spot vs. observed bounds (left) and spot differences (right). . . 78
5.7 Time series of Heston parameters for the stable dataset. 79
5.8 Time series of Heston parameters for the unstable dataset. 80
5.9 Time series of Heston parameters calibrated on stable data across 2018–2024. 81
5.10 Heston parameter estimates over a five-day stable sample (6–11 May 2024). 82
5.11 Time series of Heston parameters calibrated using unstable data. 83
5.12 Daily interest rate and dividend yield adjustments under stable and unstable regimes. 84
5.13 Gradients of option prices with respect to Heston parameters under stable market conditions (2018–2024). 86
5.14 Median Greeks under stable conditions: comparison of market-implied values, Black–Scholes fits, and Heston fits. 91
5.15 Heatmaps of differences in Greeks under alternative estimation methods. 93

List of Tables

4.1 Annual dividends per share and dividend yields for Apple Inc. over 2018–2024. 45

5.1 BPCM solutions to common Heston calibration limitations. 52
5.2 Calibration errors (MSE, variance of absolute errors, RMSE) for stable and unstable datasets across BPCM layers and the traditional calibration. 66
5.3 Spread-based diagnostics by method and regime: conditional mean position inside the spread, mean excess distance over all observations, and mean excess distance over violations only. 68
5.4 Runtime summary by window length for stable and unstable option sets (seconds). 74
5.5 Pairwise gradient correlations under stable and unstable regimes. . . . 88

F.1.1 Correlation results for stable and unstable data (showing $|\texttt{corr}| \geq 0.30$) 164
F.2.1 Months with large residual spikes (residual > 0.5) and associated events for Apple and the market 172

List of Algorithms

2.1 Heston Pricing with Characteristic Function 17
3.1 Hessian-Aware Trust-Region Calibration 33
5.1 Layer 1: Put–Call Parity Filtering 54
5.2 Layer 2: Daily Calibration of Heston Parameters (Single Day) 56
5.3 Layer 3: Joint Refinement Procedure 59
5.4 Levenberg–Marquardt solver with analytic Jacobian 62

Chapter 1

Introduction

1.1 Option Pricing and the Heston Model

Financial options (contracts deriving value from an underlying asset's price) play a central role in modern markets. They grant the right (but not obligation) to buy or sell assets under specified terms and dates, enabling hedging and speculation (Hull, 2018). This thesis focuses on *European options*, which can only be exercised at maturity. A fundamental challenge in option pricing is choosing a model that accurately captures the asset's price dynamics, especially its volatility behaviour, and then *calibrating* the model's parameters to market data. The widely-used Black–Scholes model (Black and Scholes, 1973) assumes constant volatility, which is often unrealistic. By contrast, the *Heston stochastic volatility model* (Heston, 1993) incorporates a second stochastic process for volatility, providing a more flexible and realistic framework for option pricing. The Heston model is a cornerstone of quantitative finance, but its high complexity makes calibration significantly more difficult. Option prices under the Heston model can be highly sensitive to small parameter changes, and the calibration objective is a multi-dimensional, non-convex error surface with potential local minima. In practice, straightforward least-squares calibration (fitting model prices to market prices) may become unstable or inaccurate due to the complexity, noisiness and incompleteness of real market data (Albrecher et al., 2006). In response to these challenges, this thesis focuses on the development of a robust calibration framework that we call the *Balance Premium Calibration Method (BPCM)*. This method addresses several structural limitations in standard calibration methods (outlined in Section 1.4), and extracts additional information from calibration residuals.

1.2 Challenges in Model Calibration

A key prerequisite for efficient calibration is the ability to generate model-based option prices quickly and accurately for any given set of model parameters. There are several numerical approaches to pricing European options under the Heston model. Monte Carlo simulation is conceptually simple but computationally intensive, requiring many simulated asset paths to achieve low error. Partial differential equation (PDE) methods (solving the two-dimensional Fokker–Planck or pricing PDE) can be accurate but are also time-consuming and must be carefully discretised for stability. By contrast, the Heston model admits a semi-analytical solution for European option prices via its characteristic function (Heston, 1993). This allows prices to be obtained by numerical integration, which can be accelerated using the fast Fourier transform (Carr and Madan, 1999).

In this thesis, the characteristic function approach is used as the pricing engine for calibration. This approach is both deterministic and efficient (avoiding the sampling noise of Monte Carlo and the heavy computation of fine-grid PDE solvers (Glasserman, 2004; Carr and Madan, 1999)) which is crucial when thousands of pricing iterations are required inside an optimisation routine. By leveraging the closed-form characteristic function, option prices and sensitivities can be rapidly and reliably computed for each candidate parameter set during calibration, thereby enabling a thorough search for the best-fit parameters. After identifying key calibration challenges, this thesis investigates how the Heston stochastic volatility model can be robustly calibrated to real market data while accounting for data imperfections and extracting diagnostic insights from the calibration residuals.

1.3 Data Engineering for Effective Model Calibration

Successful model calibration requires high-quality market data. This thesis utilised a dataset of European call and put option prices written on a major publicly traded stock (AAPL) to examine the model's performance relative to a single asset's market dynamics and news events.

The dataset encompasses strike prices and maturities, reflecting the full shape of the implied volatility surface at multiple time snapshots. Rigorous data cleaning and preprocessing were performed to address data issues such as missing quotes, inconsistent date formats, data entry errors, and arbitrage violations. Arbitrage violations refer to pricing inconsistencies that enable a trader to construct a zero-cost portfolio

yielding non-negative payoffs in all future states and strictly positive payoff in at least one. According to (Glasserman, 2004), such arbitrage opportunities violate the conditions required for a valid risk-neutral pricing measure, undermining the integrity of model-based valuation frameworks.

1.4 Limitations of Existing Calibration Approaches

This thesis addresses three fundamental limitations of traditional Heston calibration methods. These are highlighted as follows.

Data Quality and Diagnostic Issues. Standard calibration approaches often treat recorded mid prices as ground truth, ignoring microstructure noise, arbitrage violations, and asynchronous quotes that can destabilise calibration (Clevenhaus et al., 2024; Wang and Bai, 2025). Such contaminated inputs not only reduce calibration robustness but also limit practitioners' ability to detect model failure and adapt to changing market conditions. Cleaning arbitrage data in option datasets can therefore substantially improve both calibration stability and diagnostic capability.

Numerical Instability. The five-dimensional Heston parameter space exhibits non-convex objective surfaces with multiple local minima (Cui et al., 2017), leading to unreliable parameter estimates. Many existing methods depend on initial parameter guesses and lack robust convergence guarantees.

1.5 Contributions

This thesis presents four novel contributions that improve Heston model calibration. In particular, we propose BPCM, a multi-layer calibration method fundamentally differing from traditional one-shot Heston calibration methods. Rather than fitting the model directly to raw option quotes, BPCM first enforces *structural consistency* in the data through a three-layer refinement process:

- *Layer 1: Data Filtering.* Applies structural repairs to raw option data by enforcing key economic constraints, including bid–ask spreads (ask-side quotes are prioritised), put–call parity, and no-arbitrage conditions. Additionally, it verifies that inferred spot prices lie within the observed daily high–low range of the underlying asset. These filters produce robust, economically consistent inputs for calibration.

- *Layer 2: Parameter Calibration.* Fits Heston model parameters to the filtered *daily* dataset using closed-form characteristic-function pricing with analytic gradients, producing preliminary fits.
- *Layer 3: Joint Error Redistribution.* This layer treats stable and unstable data separately, reallocating noise under bounded constraints. In addition to the five Heston parameters, it also examines market inputs such as the risk-free rate and dividend yield.

This novel approach addresses well-known market data issues (microstructure noise, bid-ask spread biases, and asynchronous quotes) which are typically ignored by standard calibration methods.

Comprehensive Analytical Second-Order Sensitivity Framework. While closed-form expressions for the first-order Heston sensitivities (the gradient with respect to model parameters) have appeared in prior literature, this thesis is the first to derive and present the complete Hessian matrix of the model's pricing function, encompassing all second-order partial derivatives with respect to the five core parameters. In addition, we provide new closed-form expressions for both first- and second-order derivatives with respect to the risk-free rate and time to maturity, two key exogenous variables that are often held fixed in prior treatments. By delivering these analytic second-order sensitivities alongside established gradients, the framework supports rigorous local curvature analysis of the calibration objective, facilitates future development of second-order optimisation techniques.

Residual Analysis as a Market Diagnostic Tool. This thesis treats residuals (the differences between model predicted option prices and observed market prices) as valuable signals that can reveal market inefficiencies and structural patterns rather than as error terms to be minimised. This alternative perspective transforms calibration from a purely technical exercise into a practical analytical tool.

Economically Consistent Data Engineering Framework. This thesis presents a rigorous data engineering pipeline that constructs high-quality, structurally consistent market datasets through the enforcement of no-arbitrage conditions and structural relationships. This ensures calibration processes operate on reliable, arbitrage-free market data, making model outputs statistically and economically significant.

These contributions address the three core challenges that hinder broader Heston model adoption: numerical instability, computational inefficiency, and diagnostic

limitations. The proposed methodology demonstrates remarkable practical benefits, achieving 91.58% boundary adherence across over 1.4 million contracts while maintaining crisis period resilience and systematic adaptation to structural breaks. The remainder of this thesis is organised as follows. Chapter 2 introduces option pricing models, focusing on the Black–Scholes and Heston frameworks and their solution methods. Chapter 3 discusses calibration methods, reviews relevant literature, and presents the proposed analytic derivative framework. Chapter 4 describes the data sources and preprocessing pipeline developed for this research. Chapter 5 introduces the BPCM and presents empirical results from its application. Finally, Chapter 6 summarises the main contributions of this thesis, discusses their implications, and outlines directions for future research. Chapters 2 and 3 establish the theoretical foundation of the BPCM method, which will be implemented in the subsequent chapters of this thesis. Furthermore, the derivation of the second-order derivatives (specifically the Hessian) has been included to lay the groundwork for future developments aimed at enhancing the computational rigour of the Cui et al. (2017) method.

Chapter 2

Option Price Model

2.1 Motivation and Model Context

Options are fundamental instruments in modern financial markets, used extensively for hedging, speculation, and volatility trading. Accurately modelling their prices is essential for both theoretical finance and real-world risk management. The mathematical modelling of options, which provide the right but not the obligation to buy (call) or sell (put) an underlying asset at a predetermined strike price within a specified maturity period, has garnered significant attention. European options are exercisable only at maturity, whereas American options may be exercised at any time before expiry. Early contributions to option valuation include Boness (1964), while the modern option pricing theory was established with the classic works of Black and Scholes (1973) and Merton (1973). For a detailed overview, see also Hull and White (1987).

Over the past decades, numerous frameworks have emerged to determine the fair value of options. A seminal contribution is the Black–Scholes model (Black and Scholes, 1973), which states that the asset price follows a geometric Brownian motion with constant volatility. This model yields a closed-form solution for European options and catalysed extensive research in the field. Nonetheless, its simplifying assumptions soon proved insufficient (particularly the assumption of constant volatility) in capturing key features of real-world markets.

In particular, the assumption of constant volatility significantly limits the model's ability to reflect real market behaviour accurately. Financial markets exhibit well-documented phenomena such as stochastic volatility, volatility clustering, and asymmetric return distributions (features that are unaccounted for in the Black–Scholes framework). The Heston model (Heston, 1993) addresses these limitations by modelling volatility as a stochastic process governed by a mean-reverting square-root diffusion. This extension enhances theoretical realism and admits a semi-closed-form solution for European call options, rendering it particularly applicable in foreign ex-

change and equity derivatives markets.

Despite its advantages, the Heston model poses computational challenges and inefficiencies in pricing and calibration; for further discussion, see (Zhang, 2023; Cui et al., 2017). The formula of Carr and Madan (1999) marked a breakthrough by applying the fast Fourier transform (FFT) to the Heston characteristic function, enabling efficient option pricing.

However, this approach also introduced numerical complications, including sensitivity to the damping parameter, aliasing artefacts, and instability during gradient-based calibration, as noted by (Lee, 2005; Kahl and Jäckel, 2006; Rouah, 2013).

To mitigate these issues, Cui et al. (2017) proposed a stabilised formulation of the Heston characteristic function. Their method resolves branch-switching complications and improves numerical continuity, especially for long maturities or illiquid contracts.

Alternative numerical techniques (such as finite differences and Monte Carlo simulations (e.g. Andersen, 2008; Lord et al., 2010)) exist for solving the Heston model. These are generally better suited to exotic or American options, where path dependence or early exercise features necessitate full state-space discretisation. Though flexible, such methods are computationally intensive, particularly for iterative calibration.

This thesis focuses on European vanilla options, where characteristic-function-based approaches offer superior scalability and efficiency. Prioritising stability and speed, this research adopts the Cui et al. (2017) stabilised formulation to address the numerical instabilities inherent in earlier implementations of the Heston model. The objective is to present accurate, robust models suitable for rapid automated calibration pipelines. The remainder of this chapter systematically introduces and evaluates the Black–Scholes model, the Heston model, the Carr–Madan approach, and the stabilised formulation of Cui et al. (2017), which forms the core of the pricing mechanism used throughout this thesis.

2.2 Black–Scholes Model

The Black–Scholes model was the first model developed for calculating the price of European options. The model is based on the following assumptions (Black and Scholes, 1973): no dividends are paid out during the life of the option, there are no transaction costs in buying the option, and the risk-free rate and volatility of the underlying asset are known and constant. Moreover, the returns of the underlying asset are normally distributed, and the option is exercisable only at expiration. Let

$u(s,t)$ denote the option price for the underlying asset price s at time t. The option price satisfies the PDE:

$$\frac{\partial u}{\partial t} + \tfrac{1}{2}\sigma^2 s^2 \frac{\partial^2 u}{\partial s^2} + (r-q)s\frac{\partial u}{\partial s} - ru = 0,$$

where r is the risk-free rate (e.g. from government bonds) and σ is the constant volatility. This equation is derived under the risk-neutral measure $\mathbb{Q}$, where the asset price S_t follows the SDE:

$$\mathrm{d}S_t = (r-q)S_t\mathrm{d}t + \sigma S_t \mathrm{d}W_t^{\mathbb{Q}}.$$

Under the real-world (physical) measure $\mathbb{P}$, the asset earns an expected return μ that includes a risk premium. However, for arbitrage-free pricing, we transition to the risk-neutral measure $\mathbb{Q}$, under which the asset's drift becomes the risk-free rate r. This change of measure is justified by Girsanov's theorem, which mathematically replaces μ with r in the pricing dynamics:

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \, \mathrm{d}W_t^{\mathbb{P}}.$$

Considerations of constant volatility have raised doubts about the Black–Scholes model's realism; accordingly, model extensions by Hull and White (1987), Wiggins (1987), and Scott (1987) incorporated stochastic volatility. However, these more realistic models lacked simple numerical methods and closed-form solutions. In response to these limitations, Heston (1993) proposed a stochastic variance model for European options that governs the asset price S_t and variance V_t, and offers a closed-form solution.

2.3 Heston Model

Heston's (1993) model for pricing European options combines two SDEs involving two stochastic quantities: the asset price S_t and the variance V_t, whose square root $\sqrt{V_t}$ represents the instantaneous volatility. The variance is a mean-reverting square-root process. It was introduced by Cox et al. (1985) and is now often called the Cox–Ingersoll–Ross process. In addition, the Heston model takes into account the correlation between the Brownian shocks driving the asset price and the variance, commonly referred to as the leverage effect, where an increase in volatility often contributes to a decrease in the asset price (Shen, 2020).

The underlying PDEs for the option price are obtained using the Itô's Lemma, see (Appendix A.1). The characteristic function was then utilised to simplify the option pricing calculation, notably via the Feynman–Kac theorem. The Heston model is formulated as follows:

$$\begin{aligned} \mathrm{d}S_t &= (r-q)\,S_t\,\mathrm{d}t + \sqrt{V_t}\,S_t\,\mathrm{d}W_t^s, \\ \mathrm{d}V_t &= \kappa(\bar{v}-V_t)\mathrm{d}t + \sigma\sqrt{V_t}\mathrm{d}W_t^v, \end{aligned} \tag{2.1}$$

where W_t^s and W_t^v are $\mathbb{Q}$-Brownian motions with $\mathbb{E}[\mathrm{d}W_t^s \mathrm{d}W_t^v] = \rho \mathrm{d}t$. The parameters are the risk-free rate r, mean reversion rate κ, long-term variance $\bar{v}$, volatility-of-volatility σ, and correlation ρ. Under $\mathbb{Q}$ the money–market account grows at rate r and the stock's drift is $r-q$ (so the discounted cum-dividend stock is a martingale).

Lemma 1 (Itô's Lemma)**.** *Let X_t be a stochastic process defined by the SDE:*

$$\mathrm{d}X_t = \mu(X_t,t)\mathrm{d}t + \sigma(X_t,t)\,\mathrm{d}W_t,$$

where $\mu(X_t,t)$ is the drift term and $\sigma(X_t,t)$ is the diffusion term. The differential of the twice-differentiable function $f(X_t,t)$ is:

$$\mathrm{d}f(X_t,t) = \frac{\partial f}{\partial t}\,\mathrm{d}t + \frac{\partial f}{\partial X}\,\mathrm{d}X_t + \tfrac{1}{2}\,\frac{\partial^2 f}{\partial X^2}\,(\mathrm{d}X_t)^2.$$

Theorem 1 (Feynman–Kac)**.** *Let X_t be a stochastic process governed by the SDE*

$$\mathrm{d}X_t = \mu(X_t,t)\mathrm{d}t + \sigma(X_t,t)\,\mathrm{d}W_t,$$

and let $u(x,t)$ be a function satisfying the following PDE:

$$\frac{\partial u}{\partial t} + \mu(x,t)\frac{\partial u}{\partial x} + \tfrac{1}{2}\sigma^2(x,t)\frac{\partial^2 u}{\partial x^2} = ru.$$

The solution $u(x,t)$ represents the conditional expectation

$$u(x,t) = \mathbb{E}\left[e^{-r(T-t)}g(X_T) \mid X_t = x\right],$$

where r is the risk-free interest rate under the risk-neutral measure and $g(X_T)$ is a given payoff function at the final time T, defining the value of $u(x,T)$ as $u(x,T) = g(X_T)$.

When applied to u in terms of S_t and V_t, Itô's Lemma has the following form:

$$\mathrm{d}u = \frac{\partial u}{\partial t}\mathrm{d}t + \frac{\partial u}{\partial S}\,\mathrm{d}S_t + \frac{\partial u}{\partial V}\,\mathrm{d}V_t + \tfrac{1}{2}\frac{\partial^2 u}{\partial S^2}\,(\mathrm{d}S_t)^2 + \tfrac{1}{2}\frac{\partial^2 u}{\partial V^2}\,(\mathrm{d}V_t)^2 + \frac{\partial^2 u}{\partial S \partial V}\,\mathrm{d}S_t\,\mathrm{d}V_t. \quad (2.2)$$

Incorporating the $\mathrm{d}S_t$ and $\mathrm{d}V_t$ SDEs from the Heston model into Equation (2.2) produces the following PDE:

$$\frac{\partial u}{\partial t} + \tfrac{1}{2}\,V\,S^2\,\frac{\partial^2 u}{\partial S^2} + \rho\,\sigma\,V\,S\frac{\partial^2 u}{\partial S \partial V} + \tfrac{1}{2}\,\sigma^2\,V\,\frac{\partial^2 u}{\partial V^2} + (r-q)\,S\frac{\partial u}{\partial S} + \kappa\,(\bar{v}-V)\,\frac{\partial u}{\partial V} - r\,u = 0.$$

Next, define the log spot $s_t = \ln S_t$. Applying Itô's lemma to (2.1) (see Appendix A.2 for steps) gives

$$\begin{aligned} \mathrm{d}s_t &= \Big(r - q - \tfrac{1}{2}v_t\Big)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^s, \\ \mathrm{d}v_t &= \kappa(\bar{v} - v_t)\,\mathrm{d}t + \sigma\sqrt{v_t}\,\mathrm{d}W_t^v, \qquad \mathbb{E}[\mathrm{d}W_t^s \mathrm{d}W_t^v] = \rho\,\mathrm{d}t. \end{aligned} \quad (2.3)$$

By the standard replication argument (or Feynman–Kac), the option price $C = C(s, v, t)$ solves

$$\frac{\partial C}{\partial t} + \Big(r - q - \tfrac{1}{2}v\Big)\frac{\partial C}{\partial s} + \kappa(\bar{v} - v)\frac{\partial C}{\partial v} + \tfrac{1}{2}v\,\frac{\partial^2 C}{\partial s^2} + \tfrac{1}{2}\sigma^2 v\,\frac{\partial^2 C}{\partial v^2} + \rho\sigma v\,\frac{\partial^2 C}{\partial s\,\partial v} = r\,C, \quad (2.4)$$

with terminal condition $C(s, v, T) = (e^s - e^k)^+$.

Semi-analytical call formula. Let $k = \ln K$ and $\tau = T - t$. The Heston call price can be written in the usual form

$$C(s_t, K, T) \;=\; e^{\,s_t - q\tau}\,P_1(s_t, v_t, \tau; k) \;-\; e^{\,k - r\tau}\,P_2(s_t, v_t, \tau; k). \quad (2.5)$$

Here P_1 and P_2 are *risk-neutral exercise probabilities* in the Heston sense. With backward time $\tau = T - t$,

$$P_j(s, v, \tau; k) \;=\; \mathbb{E}^{\mathbb{Q}}\big[\mathbf{1}_{\{S_T > K\}} \;\big|\; S_t = e^s,\, V_t = v\big], \quad j = 1, 2,$$

and the initial condition at $\tau = 0$ (expiry) is

$$P_j(s, v, 0; k) = \mathbf{1}_{\{s \geq k\}}. \quad (2.6)$$

Substituting (2.5) into (2.4) in the standard way Heston (1993) shows that P_1 and P_2 each satisfy a linear parabolic PDE of the same form, differing only by a constant u_j in the s-drift:

$$\frac{\partial P_j}{\partial \tau} = \tfrac{1}{2} v \frac{\partial^2 P_j}{\partial s^2} + \rho\sigma v \frac{\partial^2 P_j}{\partial s \partial v} + \tfrac{1}{2}\sigma^2 v \frac{\partial^2 P_j}{\partial v^2} + (a - b_j v)\frac{\partial P_j}{\partial v} + ((r-q) + u_j v)\frac{\partial P_j}{\partial s}, \quad (2.7)$$

with initial condition $P_j(s, v, 0; k) = \mathbf{1}_{\{s \geq k\}}$, where

$$u_1 = \tfrac{1}{2}, \qquad u_2 = -\tfrac{1}{2}, \qquad a = \kappa\bar{v}, \qquad b_1 = \kappa + \lambda - \rho\sigma, \qquad b_2 = \kappa + \lambda. \quad (2.8)$$

(Here λ is Heston's volatility risk premium; set $\lambda = 0$ for pure risk-neutral pricing.)

Characteristic-function solution (ansatz). Following Heston, one solves (2.7) via the exponential-affine ansatz for the *characteristic functions* of the log-price under the two specifications $j = 1, 2$:

$$\varphi_j(\phi; s, v, \tau) \;=\; \exp\big(C_j(\tau, \phi) \;+\; D_j(\tau, \phi)\, v \;+\; i\phi\, s\big), \quad (2.9)$$

A full derivation of this exponential–affine form is provided in Appendix B.1. where $C_j(\tau, \phi)$ and $D_j(\tau, \phi)$ satisfy the associated Riccati equations.

$$\frac{\partial D_j}{\partial \tau} = \tfrac{1}{2}\phi^2 - \rho\sigma i\phi\, D_j - \tfrac{1}{2}\sigma^2 D_j^2 + u_j i\phi + b_j D_j, \quad (2.10)$$

$$\frac{\partial C_j}{\partial \tau} = -i\phi(r-q) - a\, D_j, \quad (2.11)$$

with $C_j(0, \phi) = 0$, $D_j(0, \phi) = 0$. The closed-form solutions are

$$D_j(\tau, \phi) = \frac{b_j - \rho\sigma i\phi + d_j}{\sigma^2}\,\frac{1 - e^{-d_j\tau}}{1 - h_j e^{-d_j\tau}}, \quad (2.12)$$

$$C_j(\tau, \phi) = -i\phi(r-q)\tau + \frac{a}{\sigma^2}\left[(b_j - \rho\sigma i\phi + d_j)\tau - 2\ln\frac{1 - h_j e^{-d_j\tau}}{1 - h_j}\right], \quad (2.13)$$

where

$$h_j = \frac{b_j - \rho\sigma i\phi + d_j}{b_j - \rho\sigma i\phi - d_j}, \qquad d_j = \sqrt{(b_j - \rho\sigma i\phi)^2 - \sigma^2(2u_j i\phi - \phi^2)}.$$

Full derivations of D and C are provided in Appendix B.2. Finally, P_1 and P_2 are recovered by Fourier inversion:

$$P_1(s, v, \tau; k) = \tfrac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re\left[\frac{e^{-i\phi k}\, \varphi_1(\phi; s, v, \tau)}{i\phi}\right] \mathrm{d}\phi, \tag{2.14}$$

$$P_2(s, v, \tau; k) = \tfrac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re\left[\frac{e^{-i\phi k}\, \varphi_2(\phi; s, v, \tau)}{i\phi}\right] \mathrm{d}\phi. \tag{2.15}$$

Equations (2.5)–(2.8) together with the characteristic-function solution (2.9) reproduce Heston's semi-analytical pricing result, explicitly including the dividend yield q via the drift $(r - q)$ and the discount factors $e^{-q\tau}$ and $e^{-r\tau}$.

Carr–Madan Formula for Option Pricing

The previous section derived the characteristic function and applied the Fourier transform to determine the option price. The FFT can be employed to further increase calculation efficiency; however, its direct application poses two key challenges. First, the function is undefined at $\phi = 0$. Second, the option call price function $C(k)$ is not square-integrable because the option $C \to S_0$ has $k \to -\infty$. To address these issues, Carr and Madan (1999) scaled $C(k)$ by multiplying it with $\exp(\alpha k)$, which helps to make the function square-integrable. Here, $\alpha > 0$ is chosen so that the damped call price $\exp(\alpha k)C(k)$ decays at the tails, guaranteeing the existence of its Fourier transform. This enhancement only requires one characteristic function, which directly gives the option price for a log-strike price k as a transform of the log-price s of the characteristic function (Equation (2.9)).

Given the risk-neutral density of the terminal log-price $q(s)$ at maturity T, the characteristic function is

$$\varphi(\phi; s, v, t) = \int_{-\infty}^{\infty} e^{i\phi s}\, q(s)\, \mathrm{d}s$$

Let an initial call price value $C(k)$ for the asset log price s at maturity T be:

$$C(k) = e^{-r\tau}\, \mathbb{E}\left[\left(e^{s_T} - e^{k}\right)^{+}\right], \tag{2.16}$$

$$C(k) = e^{-r\tau} \int_k^\infty \left(e^{s} - e^{k}\right) q(s)\, \mathrm{d}s. \tag{2.17}$$

As $k \to -\infty$, $C(k)$ tends to $S_0\, e^{-q\tau}$, which refers to the call pricing function; thus, it is not square-integrable. To make the function square-integrable, the modified call

price $c(k)$ is defined as

$$c(k) = e^{\alpha k} C(k), \tag{2.18}$$

where $\alpha > 0$ is chosen such that $c(k)$ is square-integrable and adjusts the exponential growth to make $c(k)$ well-behaved for the Fourier transform. Furthermore, the Fourier transform of the modified call price $c(k)$ is defined as follows:

$$\psi(\phi) = \int_{-\infty}^{\infty} e^{i\phi k} c(k)\, \mathrm{d}k. \tag{2.19}$$

Substituting the definition of $c(k)$ from Equation (2.18) produces

$$\psi(\phi) = \int_{-\infty}^{\infty} e^{i\phi k} e^{\alpha k} C(k)\, \mathrm{d}k. \tag{2.20}$$

Next, Equation (2.17) is substituted into Equation (2.20):

$$\psi(\phi) = \int_{-\infty}^{\infty} e^{i\phi k} e^{\alpha k} \left(\int_{k}^{\infty} e^{-r\tau} \left(e^{s} - e^{k} \right) q(s)\, \mathrm{d}s \right) \mathrm{d}k.$$

Assuming the integrals converge appropriately, the order of integration is interchanged:

$$\psi(\phi) = e^{-r\tau} \int_{-\infty}^{\infty} q(s) \left(\int_{-\infty}^{s} e^{i\phi k} e^{\alpha k} \left(e^{s} - e^{k} \right) \mathrm{d}k \right) \mathrm{d}s.$$

To evaluate the inner integral, the integrand is separated into two terms:

$$\begin{aligned} \psi(\phi) &= e^{-r\tau} \int_{-\infty}^{\infty} q(s) \left[\frac{e^{(\alpha+1+i\phi)s}}{\alpha + i\phi} - \frac{e^{(\alpha+1+i\phi)s}}{\alpha + 1 + i\phi} \right] \mathrm{d}s \\ &= e^{-r\tau} \int_{-\infty}^{\infty} q(s) \left[\frac{e^{i(-i(\alpha+1)+\phi)s}}{\alpha^2 + \alpha - \phi^2 + i\phi(2\alpha+1)} \right] \mathrm{d}s \\ &= e^{-r\tau} \frac{\varphi(\phi - i(\alpha+1); s_0, v_0, t)}{\alpha^2 + \alpha - \phi^2 + i\phi(2\alpha+1)}. \end{aligned} \tag{2.21}$$

The call price is obtained by applying the inverse Fourier transform to Equation (2.20):

$$C(k) = e^{-\alpha k} \frac{1}{2\pi} \int_{-\infty}^{\infty} \Re \left[e^{-ik\phi} \psi(\phi) \right] \mathrm{d}\phi,$$

due to the symmetry of the integrand, the domain of integration reduces to

$$C(k) = \frac{e^{-\alpha k}}{\pi} \int_{0}^{\infty} \Re \left[e^{-i\phi k} \psi(\phi) \right] \mathrm{d}\phi. \tag{2.22}$$

Then, the price of a put option is given by:

$$P(k) = C(k) + e^{k-r\tau} - S_0\, e^{-q\tau}.$$

Given the complexity of evaluating the integral in the inverse Fourier transform, the trapezoid rule is applied to estimate Equation (2.22):

$$C(k) \approx \frac{e^{-\alpha k}}{\pi} \sum_{j=1}^{N} e^{-i\varphi_j k} \psi(\varphi_j) \Delta\phi. \tag{2.23}$$

Carr and Madan (1999) developed an efficient algorithm that uses the FFT to compute the sum in Equation (2.23). The discrete Fourier transform of ω of y is given by:

$$\omega(k) = \sum_{j=1}^{N} e^{-i\frac{2\pi(k-1)(j-1)}{N}} y(j), \quad \text{for } k = 1, \ldots, N. \tag{2.24}$$

Following Carr and Madan (1999), Equation (2.23) is matched with Equation (2.24) to discretise the FFT:

$$C(k_n) = \Delta\phi \frac{e^{-\alpha k_n}}{\pi} \sum_{j=1}^{N} \Re\left[e^{-i\Delta k \Delta\phi (j-1)(n-1)} e^{iu_j \left[\frac{N\Delta k}{2} - \ln(S_t)\right]} \psi(\varphi_j) \right] w_j,$$

where ϕ is the characteristic function variable and α is the damping factor. The weights $w_j = (\Delta\phi/3)(3 + (-1)^j - \delta_{j-1})$ are the Simpson's rule weights for quadrature used to approximate the integral, while δ_n is the Kronecker delta given by

$$\delta_n = \begin{cases} 1, & \text{if } n = 0, \\ 0, & \text{otherwise.} \end{cases}$$

The FFT discretisation involves partitioning the log-strike k into N points with step size Δk and similarly discretising the characteristic function variable ϕ into N points with step size $\Delta\phi$. In the original Carr–Madan formulation, the log-strike grid is defined as $k_n = k_{\min} + (n-1)\Delta k$, for $n = 1, \ldots, N$. (Here we describe it as centred at the spot s for convenience.) The integration variable values are discretised as $\varphi_j = (j-1)\Delta\phi$, for $j = 1, \ldots, N$. The discretisation constraint is

$$\Delta k = \tfrac{2\pi}{N\Delta\phi}.$$

This approach simplifies the integral calculation required for option pricing, making it computationally feasible for large datasets.

2.4 Reformulated Characteristic Function

The Carr–Madan FFT method (Carr and Madan, 1999) enables the efficient computation of option prices across a broad range of strikes. However, its application to Heston model calibration poses several numerical challenges. One common issue is the appearance of numerical artefacts, as truncating the integration domain or selecting a suboptimal damping factor (α) can lead to oscillatory pricing behaviour, particularly for deep out-of-the-money or long-dated options. Additionally, the method is highly sensitive to grid configuration parameters such as the number of points N, spacing η, and integration bound b (Boyarchenko and Levendorskii, 2024). These parameters often require manual tuning to avoid aliasing, which complicates automated calibration routines. Furthermore, when computing gradients via finite differences on FFT-based prices, the underlying numerical noise is frequently amplified. This can distort the objective function landscape, introducing spurious local minima and reducing the reliability of the calibration process.

To address these challenges, Cui et al. (2017) proposed a semi-analytical method based on a reformulated characteristic function. Their approach eliminates discontinuities associated with branch switching and ensures smoothness across the domain. As a result, their method significantly improves both numerical stability and the robustness of calibration procedures, particularly in the context of long-maturity and illiquid options.

In particular, Cui et al. (2017) reformulated the Heston characteristic function in terms of the forward price

$$F = S_t e^{(r-q)\tau},$$

which naturally incorporates the drift adjustment. This substitution yields compact expressions for the risk-neutral probabilities P_1 and P_2, avoiding redundant exponential factors and ensuring stability across maturities. In this way, the pricing representation remains both efficient and well-suited for calibration. The reformulated characteristic function is

$$\varphi(u;\tau) = \exp\left(iu \ln \frac{F}{S_0} - \frac{\kappa\bar{v}\rho\tau iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2} D\right), \tag{2.25}$$

where

$$
\begin{aligned}
F &= S_t e^{(r-q)\tau} \quad \text{(forward price)},\\
d &= \sqrt{\xi^2 + \sigma^2(iu + u^2)},\\
\xi &= \kappa - \rho\sigma iu,\\
A &= \frac{A_1}{A_2}, \quad A_1 = (u^2 + iu)\sinh\frac{d\tau}{2}, \quad A_2 = \frac{d}{v_0}\cosh\frac{d\tau}{2} + \frac{\xi}{v_0}\sinh\frac{d\tau}{2},\\
D &= \ln\frac{d}{v_0} + \frac{(\kappa - d)\tau}{2} - \ln\left(\frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}e^{-d\tau}\right).
\end{aligned}
$$

This reformulation resolves logarithmic discontinuities via the restructured term D (equation 16e in Cui et al., 2017). The price of a vanilla European call option is given by

$$
C(\boldsymbol{\theta}; K, \tau) = \tfrac{1}{2}\left(S_t e^{-q\tau} - K e^{-r\tau}\right) \quad + \frac{e^{-r\tau}}{\pi}\left[S_t P_1 - K P_2\right], \tag{2.26}
$$

where

$$
\begin{aligned}
P_1 &= \int_0^\infty \Re\left(\frac{e^{-iu\ln\frac{K}{S_0}}\varphi(u - i, \tau)}{iu}\right)\mathrm{d}u,\\
P_2 &= \int_0^\infty \Re\left(\frac{e^{-iu\ln\frac{K}{S_0}}\varphi(u, \tau)}{iu}\right)\mathrm{d}u.
\end{aligned}
$$

The risk-neutral probabilities P_1 and P_2, which appear in the pricing formula (Equation (2.26)), are computed numerically using a Gauss–Legendre quadrature method (Elliott and Johnston, 2008). Gauss–Legendre quadrature is a numerical integration method that approximates definite integrals by optimally choosing evaluation nodes and weights to exactly integrate polynomials up to $2N-1$ degrees. Specifically, given N quadrature nodes $\{u_k\}_{k=1}^N \in (0, u_{\max})$ and their corresponding weights $\{w_k\}_{k=1}^N$, the integrals are evaluated as follows:

$$
\begin{aligned}
P_1 &\approx \sum_{k=1}^{N} \Re\left[\frac{e^{-iu_k\ln(K/S_0)}\varphi_1(u_k)}{iu_k}\right] w_k,\\
P_2 &\approx \sum_{k=1}^{N} \Re\left[\frac{e^{-iu_k\ln(K/S_0)}\varphi_2(u_k)}{iu_k}\right] w_k,
\end{aligned} \tag{2.27}
$$

where $\varphi_1(u) = \varphi(u - i; \tau)$ and $\varphi_2(u) = \varphi(u; \tau)$ denote the damped and undamped characteristic functions, respectively.

Cui et al. (2017) highlighted several notable strengths of their calibration technique. Their proposed method effectively eliminates the oscillations commonly associated with FFT-based approaches and achieves exceptional numerical accuracy, delivering a relative error of 10^{-10} for out-of-the-money options compared to 10^{-5} using traditional FFT methods. Moreover, this method offers a significant speed advantage, with calibration procedures executing approximately 9.5 times faster than finite-difference counterparts. The method also exhibits high robustness, converging to the global optimum in 98.4% of 10,000 Monte Carlo trials.

Option Pricing Algorithm

The following algorithm outlines the steps for computing the option price using the Heston model and the numerical integration method developed by Cui et al. (2017). Unlike FFT-based approaches, this method circumvents numerical instabilities and ensures both stability and precision in pricing. Depending on the required accuracy, the integration can be performed using the adaptive Simpson's rule or Gauss–Legendre quadrature.

Algorithm 2.1 Heston Pricing with Characteristic Function

1: Input: $S_0, K, \tau, r, q, \boldsymbol{\theta} = \{\kappa, \bar{v}, \rho, \sigma, v_0\}$
2: Compute the forward price: $F = S_0 e^{(r-q)\tau}$
3: Compute the characteristic functions via Equation (2.25):
- $\varphi_1(u) = \varphi(u - i; \tau, \boldsymbol{\theta})$
- $\varphi_2(u) = \varphi(u; \tau, \boldsymbol{\theta})$

4: Select the number of quadrature nodes N and the corresponding Gauss–Legendre grid $\{u_k, w_k\}_{k=1}^{N}$
5: Compute P_1 and P_2 via Equation (2.27)
6: Compute the call option (C) via Equation (2.26)
7: Compute the put option via parity: $P = C - S_0 e^{-q\tau} + K e^{-r\tau}$

Chapter 3

Calibration Methods

3.1 Introduction

The Heston model (Heston, 1993) has become a standard framework for modelling stochastic volatility in equity and foreign exchange markets due to its ability to capture volatility smiles and skews and its semi-closed-form solution for European options.

The model's effectiveness depends on its calibration. This process is typically formulated as an inverse problem, which is determining the unknown model parameters from observed market data (Cont and Tankov, 2004), by adjusting the parameter vector $\boldsymbol{\theta}$ so that the generated model prices closely match the observed market prices (Bain et al., 2019). This inverse problem is challenging due to the non-convexity of the objective function, strong interdependencies among parameters, and the significant computational cost associated with evaluating the pricing function repeatedly over a high-dimensional parameter space. As discussed in Chapter 2, notable advancements in this area include the work by Cui et al. (2017) that reformulated the characteristic function and analytical gradient derivation, significantly improving the speed and robustness of Heston calibration.

This chapter focuses on calibrating the Heston model using non-linear least squares. Least-squares minimisation provides a direct and interpretable objective. It seeks to minimise the squared differences between model and market option prices across a grid of strikes and maturities. This is a common approach in the literature (Cui et al., 2017; Gatheral, 2006) and supports efficient gradient-based optimisation when analytic derivatives are available.

Beyond the first derivatives, this chapter highlights the importance of second-order derivatives with respect to model parameters and market inputs such as the interest rate r and time to maturity τ. These higher-order derivatives, while computationally more involved, play a critical role in the local curvature of the objective function, the design of Hessian-based optimisation algorithms, and the stability and identifiability

of calibrated parameters. This chapter outlines the computation of second derivatives for all five Heston parameters, as well as for the interest rate r (Rho derivatives) and maturity τ (Theta derivatives), while the full derivations are provided in the appendix.

The chapter is structured as follows. It begins by formally defining the calibration problem and discussing the numerical and structural challenges associated with solving it. It then presents the derivation of the first and second derivatives of the option price with respect to the model parameters and key market variables. Finally, it outlines the calibration algorithm, including the analytic gradient and Hessian-based steps that form the foundation of the approach proposed in this thesis. The empirical results and implementation details are presented later in Chapter 5.

3.2 Heston Calibration Stability

The loss landscape of the parameters in the Heston calibration is known to exhibit extreme flatness in certain parameter directions, which complicates optimisation and affects convergence. As noted by Cui et al. (2017), the objective function is shaped as a "narrow valley with a flat bottom", making it highly sensitive to initial conditions. Other studies have also reported degeneracy and identifiability issues in the parameter space (e.g. Jacquier and Shi, 2019; del Baño Rollin et al., 2010; Fengler, 2009). Moreover, the Hessian is often ill-conditioned (Cui et al., 2017), contributing to numerical instability. This ill-conditioning is a well-documented challenge in stochastic volatility model calibration, where multiple local minima are frequently encountered when using gradient-based methods.

3.3 Residual Analysis in Financial Modelling

While calibration residuals are typically treated as noise to be minimised, they may contain valuable information about market structures and model adequacy (Cui et al., 2017). In econometrics and time-series analysis, examining residuals is a standard diagnostic step to test model adequacy. Yet, residual analysis is not traditionally included in the derivatives pricing literature. In chapter 5, we fill this gap by providing a metric to analyse these residuals. This novel approach transforms calibration from a mere parameter-fitting exercise into an analytical tool that can provide valuable market insights.

3.4 Problem Statement

The calibration of the Heston model is formulated as a non-linear least-squares minimisation problem. Given N observed option prices $\{C_i^{\text{obs}}\}_{i=1}^N$ with corresponding strikes K_i and maturities T_i, and the Heston model price $C^{\text{Hstn}}(\boldsymbol{\theta}; S_0, K_i, T_i, r, q)$ parameterised by $\boldsymbol{\theta} = (v_0, \bar{v}, \rho, \kappa, \sigma)$, the objective is to adjust the parameter vector to minimise the discrepancy between the theoretical and observed prices:

$$\min_{\boldsymbol{\theta}\in\Theta} \sum_{i=1}^{N} \left[C^{\text{Hstn}}(\boldsymbol{\theta}; S_t, K_i, T_i, r, q) - C_i^{\text{obs}}\right]^2, \tag{3.1}$$

where Θ denotes the admissible parameter space, typically subject to constraints $v_0 > 0$, $\bar{v} > 0$, $\sigma > 0$, $\kappa > 0$, $\rho \in [-1, 1]$. Following the BPCM (see Chapter 5 for more details), arbitrary weighting schemes are not introduced in the loss function. This approach reduces potential biases and simplifies the optimisation landscape, ensuring that the resulting minimisation reflects genuine model-to-market misfit rather than data quality artefacts.

Despite its apparent simplicity, solving Equation (3.1) poses significant theoretical and computational challenges, requiring careful consideration of the following aspects.

Non-Convexity and Local Minima. The objective function is highly non-convex due to the non-linear dependence of option prices on the parameter vector $\boldsymbol{\theta}$. Empirical studies such as Gilli and Schumann (2013) document strong sensitivity to initial guesses, often leading to entrapment in local minima. However, Cui et al. (2017) argued that the landscape of the objective function resembles a narrow valley with a flat base rather than a surface with many isolated minima. This topography implies that, although many parameter combinations may yield nearly equivalent pricing performance, the gradient direction remains informative. Therefore, precise gradient evaluation can play a central role in achieving robust convergence.

Parameter Derivatives and Identifiability. The Heston model parameters exhibit complex interdependencies. A prominent example is the compensatory relationship between κ and $\bar{v}$, where different combinations can yield similar steady-state variance dynamics, complicating parameter identifiability. This degeneracy means that multiple parameter combinations can produce similar implied volatilities, particularly in the absence of a sufficiently diverse set of maturities (Guterding and Boenkost, 2018). As a result, identifiability becomes problematic, and some parameters may

remain weakly determined without strong regularisation or extended datasets. Thus, care must be taken to avoid overfitting or numerical instability due to poorly constrained directions in the parameter space.

Analytical Gradients and Second Derivatives. A major computational bottleneck in Heston calibration arises from the need to compute the gradients of the pricing function. While finite-difference methods are easy to implement, they entail discretisation errors and high computational costs. Building on the branch-consistent formulation by del Baño Rollin et al. (2010), Cui et al. (2017) derived an analytical expression for the gradient, improving both accuracy and efficiency by an order of magnitude. This thesis extends this approach by computing not only the first-order derivatives with respect to the five Heston parameters but also the second-order derivatives. These higher-order terms are essential for understanding the local curvature of the loss function and for designing Newton-type optimisation schemes.

External Inputs: Derivatives with Respect to Interest Rate and Maturity. Although the risk-free rate r and time to maturity τ are typically treated as exogenous inputs rather than calibration parameters, they play a central role in option valuation. The Heston model's derivatives with respect to these quantities (quantified by the Greeks Rho ($\frac{\partial C}{\partial r}$) and Theta ($\frac{\partial C}{\partial \tau}$)) mean that even small misestimations can produce significant pricing errors. For example, a slight change in r alters the discounting of expected payoffs across all maturities, while τ determines the time horizon for diffusion and thus the total variance accumulated by the option. Analytical and numerical studies have shown that particularly for long-dated or deep in/out-of-the-money options, option values can be highly sensitive to these inputs, making them a potential source of calibration error when model prices are matched to market data (Gradojevic and Kukolj, 2024; Hull, 2018; Marcato and Sebehela, 2022). Therefore, even though r and τ are fixed during calibration, their impact on model outputs must be carefully considered to ensure robust and accurate pricing.

To quantify these derivatives precisely, this chapter derives exact analytical expressions for the first- and second-order partial derivatives of the call price with respect to r and τ. Although these derivatives are not used in the calibration algorithm presented here, they are included for two reasons. First, they enhance the accuracy of the Greek calculations, especially in regimes where numerical approximations become unstable or ill-conditioned. Second, they provide a theoretical basis for exploring maturity and rate-related risks within model outputs. Their symbolic form eliminates truncation

and rounding errors associated with finite-difference schemes, thereby increasing both computational robustness and interpretability.

3.5 Analytic Derivative Computation for Calibration

Efficient Heston model calibration hinges on the ability to compute the gradient (first-order derivatives) and ideally the Hessian (second-order derivatives) of the option pricing function with respect to the model parameters. These derivatives can be obtained via numerical differentiation such as the finite differences, however, this approach is slow and prone to numerical errors in high-dimensional settings (Alòs et al., 2015). Instead, this research derives closed-form expressions for the partial derivatives of the Heston call price with respect to each model parameter. The use of analytical gradients speeds up Heston calibration by roughly an order of magnitude compared to finite-difference methods (Gudmundsson and Vyncke, 2019). Having access to explicit analytical derivatives makes it easier to apply advanced optimisation strategies that rely on gradient and curvature information for stable and rapid convergence, such as Levenberg–Marquardt or trust-region methods (Andres et al., 2020; Cui et al., 2017). In the Heston model, the price of a European call option $C(S_t, K, \tau)$ (Equation (2.26)) can be expressed semi-analytically via Fourier integrals. Here, $\varphi_j(u)$ depends on the model parameters $\boldsymbol{\theta} = (v_0, \bar{v}, \kappa, \sigma, \rho)$, which means that the differentiation of φ_j with respect to any parameter is feasible. Differentiating the pricing formula under the integral sign (a technique justified by the smooth dependence of $\varphi_j(u)$ on the parameters) produces analytic expressions for the gradient of C with respect to each component of $\boldsymbol{\theta}$. In general, the analytical gradient with respect to θ (an element of $\boldsymbol{\theta}$) is:

$$\frac{\partial C}{\partial \theta} = S_t e^{-q\tau} \frac{\partial P_1}{\partial \theta} - K e^{-r\tau} \frac{\partial P_2}{\partial \theta}, \tag{3.2}$$

where

$$\begin{aligned} \frac{\partial P_1}{\partial \theta} &= \frac{1}{\pi} \int_0^\infty \Re \left[\frac{e^{-iu \ln K}}{iu} \frac{\partial \varphi_1(u - i)}{\partial \theta} \right] \mathrm{d}u, \\ \frac{\partial P_2}{\partial \theta} &= \frac{1}{\pi} \int_0^\infty \Re \left[\frac{e^{-iu \ln K}}{iu} \frac{\partial \varphi_2(u)}{\partial \theta} \right] \mathrm{d}u. \end{aligned} \tag{3.3}$$

The derivatives $\partial\varphi(u)/\partial\theta$ can be computed analytically using the stabilised form of the Heston characteristic function (see Cui et al. (2017), Section 3.2). In the pricing formulation, the characteristic function $\partial\varphi_j(u)/\partial\theta$ is required in the evaluation of the gradient. By inserting these analytic expressions for $\partial\varphi$ into the Fourier integral

formulas for $\partial P_1/\partial\theta$ and $\partial P_2/\partial\theta$ (Equation (3.3)), and then substituting the resulting values into the pricing gradient formula for C (Equation (3.2)), we obtain closed-form expressions for the first-order derivatives of the call price with respect to each model parameter: $\partial C/\partial v_0$, $\partial C/\partial\bar{v}$, $\partial C/\partial\kappa$, $\partial C/\partial\sigma$, and $\partial C/\partial\rho$.

These analytical expressions constitute the Jacobian of model prices with respect to the parameter vector and serve as essential inputs for efficient gradient-based calibration. By using Equation (3.2) in place of finite-difference approximations, the proposed calibration framework obtains exact partial derivatives for all contracts in the dataset simultaneously. This leads to substantial gains in both computational speed and numerical precision. In practical terms, the full gradient vector $\nabla_{\boldsymbol{\theta}} C$ can be evaluated with machine-level accuracy in a fraction of the time required by numerical differentiation. As a result, gradient-based optimisation becomes tractable even in the five-dimensional setting of the Heston model. Although the primary implementation in this thesis focuses on first-order optimisation, the second-order partial derivatives of the call price function (constituting the Hessian matrix $H = \nabla^2_{\boldsymbol{\theta}} C$) have also been derived. These second-order partial derivatives are obtained by differentiating the gradient expressions (Equation (3.2)) again with respect to an additional parameter. The second-order partial derivative of the call price with respect to parameters θ_l and θ_k is given by the componentwise formula:

$$\frac{\partial^2 C}{\partial\theta_l\partial\theta_k} = S_0 e^{-q\tau}\frac{\partial^2 P_1}{\partial\theta_l\partial\theta_k} - Ke^{-r\tau}\frac{\partial^2 P_2}{\partial\theta_l\partial\theta_k}.$$

Each second-order term $\partial^2 P_j/(\partial\theta_l\partial\theta_k)$ can, in turn, be expressed as a Fourier integral involving second derivatives of the characteristic function:

$$\frac{\partial^2 P_j}{\partial\theta_l\partial\theta_k} = \frac{1}{\pi}\int_0^\infty \Re\left[e^{-iu\ln K}\frac{\partial^2\varphi_j(u)}{\partial\theta_l\partial\theta_k}\mathcal{K}(u)\right]\mathrm{d}u,$$

where $\mathcal{K}(u)$ denotes the relevant Fourier kernel in the pricing representation, and $\varphi_j(u)$ represents the damped ($j = 1$) or undamped ($j = 2$) characteristic function. The next theorem unifies these expressions into a general closed-form representation for both the gradient and Hessian of the option price, laying the theoretical foundation for second-order calibration algorithms.

In the context of calibration, knowing the Hessian is not strictly necessary to employ the Levenberg–Marquardt algorithm, which uses an approximate Hessian constructed from the Jacobian. However, having the exact Hessian available is valuable for analysing the curvature of the objective function and for experimenting with full

second-order optimisation methods, such as a Newton–Raphson scheme or other trust-region solvers that require $\nabla^2_{\boldsymbol{\theta}} C$. The Hessian reveals how the gradient changes as parameters change, offering insight into the local geometry of the error surface. For example, if the Hessian is ill-conditioned or nearly singular, it may indicate directions in parameter space where the gradient is nearly flat (i.e. regions where second derivatives are small and the option price is weakly dependent on certain parameters, suggesting possible identifiability issues). This thesis derived the Hessian analytically to ensure that the proposed theoretical framework is complete and can support such analyses.

The Heston model gradient and Hessian structure are formulated as a non-linear least squares minimisation problem as follows:

$$f(\boldsymbol{\theta}) = \frac{1}{2} \sum_{i=1}^{n} \left[C^{\mathrm{mdl}}(\boldsymbol{\theta}; K_i, \tau_i) - C_i^{\mathrm{mkt}} \right]^2 . \tag{3.4}$$

Here, C^{mdl} is the Heston price and C_i^{mkt} is the observed market price. The residual vector is defined as $\boldsymbol{r} = (r_1, r_2, \ldots, r_n)$, where $r_i = C_i^{\mathrm{mdl}} - C_i^{\mathrm{mkt}}$. The gradient of f is given by

$$\nabla f(\boldsymbol{\theta}) = \boldsymbol{J}^\top \boldsymbol{r},$$

where $\boldsymbol{\theta} \in \mathbb{R}^m$ and $\boldsymbol{J} \in \mathbb{R}^{n \times m}$ is the Jacobian matrix with elements

$$J_{ij} = \frac{\partial r_i}{\partial \theta_j} = \frac{\partial C^{\mathrm{mdl}}(\boldsymbol{\theta}; K_i, \tau_i)}{\partial \theta_j}.$$

This equation is the basis for gradient-based optimisation. Cui et al. (2017) computed this gradient analytically by differentiating under the integral sign of the Heston pricing formula. Their method avoids finite differences and improves speed and accuracy.

The Hessian of f is given by

$$\nabla^2 f(\boldsymbol{\theta}) = \boldsymbol{J}^\top \boldsymbol{J} + \sum_{i=1}^{n} r_i \boldsymbol{H}(r_i),$$

where $\boldsymbol{H}(r_i)$ is the Hessian of the residual r_i, that is, the matrix of second-order derivatives

$$[\boldsymbol{H}(r_i)]_{lk} = \frac{\partial^2 r_i}{\partial \theta_l \partial \theta_k} = \frac{\partial^2 C^{\mathrm{mdl}}(\boldsymbol{\theta}; K_i, \tau_i)}{\partial \theta_l \partial \theta_k}.$$

Cui et al. (2017) used only the first term, $\boldsymbol{J}^\top \boldsymbol{J}$, and ignored the second derivative

term. This is called the Gauss–Newton or Cauchy–Newton approximation:

$$\nabla^2 f(\boldsymbol{\theta}) \approx \boldsymbol{J}^\top \boldsymbol{J}.$$

It is important to clarify that this thesis does not claim empirical improvements in convergence, as no direct benchmarking was performed. However, the availability of analytically derived derivatives enables the application of second-order optimisation methods, such as trust-region or Newton-type algorithms, without relying on numerical approximation. The primary focus is on the exact symbolic derivation of gradients and Hessians, and on constructing a Hessian approximation using the Jacobian matrix (e.g. $\boldsymbol{J}^\top \boldsymbol{J}$) for practical calibration routines.

This thesis also goes beyond such approximations by computing the second derivatives $\boldsymbol{H}(r_i)$ analytically for each model parameter. The aim is to improve accuracy and provide deeper insight into the curvature of the objective function. The next section derives the second-order partial derivatives of the Heston pricing function with respect to each parameter in $\boldsymbol{\theta}$.

3.6 Analytical Gradient Derivation

The derivation of analytical gradients for the Heston model hinges on two essential components: the expression of the option pricing formula in terms of characteristic functions and the stabilised formulation of the Heston characteristic function $\varphi(u)$ (Cui et al., 2017). While using analytical gradients accelerates calibration and improves numerical stability, extending this approach to second-order derivatives introduces substantial complexity. Since there are five model parameters, the full Hessian matrix includes 25 second-order partial derivatives, including all cross-derivatives. Due to symmetry, only 15 of these are distinct. For each residual r_i, computing the Hessian $\boldsymbol{H}(r_i)$ means evaluating every $\frac{\partial^2 r_i}{\partial \theta_l \partial \theta_k}$ entry, and since this must be done across all observed option prices, the total computational cost scales rapidly with the number of instruments. This renders full analytic Hessian computation expensive and tedious, particularly for cross-partial derivatives with intricate derivations.

To mitigate these issues, this thesis proposes a hybrid approach that combines the Gauss–Newton approximation with selective analytical corrections. The standard Cauchy–Newton approximation to the Hessian is given by $\nabla^2 f \approx \boldsymbol{J}^\top \boldsymbol{J}$, which omits the second-order curvature terms $\sum_{i=1}^{n} r_i \boldsymbol{H}(r_i)$. While this approximation is typically adequate near the optimum, where residuals r_i are small, it may fail to capture the

full curvature of the objective surface in directions where second-order parameter derivatives are significant. For example, the initial variance parameter v_0 is known to have a significant and nonlinear impact on option prices; as a result, the second-order derivative $\frac{\partial^2 C}{\partial v_0^2}$ is expected to contribute substantially to the curvature of the objective function.

This research proposes a diagonally-corrected hybrid scheme to improve accuracy without incurring the full cost of a complete analytic Hessian. Specifically, the Gauss–Newton approximation $\boldsymbol{J}^\top \boldsymbol{J}$ is computed for the full matrix while selected diagonal entries are replaced or augmented using exact analytical second derivatives:

$$\left[\nabla^2 f\right]_{jj} \leftarrow \boldsymbol{J}^\top \boldsymbol{J} + \sum_{i=1}^{n} r_i \frac{\partial^2 r_i}{\partial \theta_l^2}, \quad l = 1, 2, \ldots, m.$$

This correction incorporates exact curvature information along key directions (e.g. for v_0 and σ), while cross-partials are set to zero. Such a hybrid method balances computational tractability with improved accuracy, particularly for parameters that induce high local curvature in the objective function and exhibit strong second-order effects. As a result, the calibration algorithm gains more reliable curvature information where it is most needed, enhancing numerical stability and accelerating convergence.

The gradient of the option price with respect to the model parameters follows directly from differentiating Equation (2.26). The expression is given by

$$\nabla_{\boldsymbol{\theta}} C = e^{-r\tau} \left(S_t \nabla_{\boldsymbol{\theta}} P_1 - K e^{-q\tau} \nabla_{\boldsymbol{\theta}} P_2 \right).$$

Here, the core computational challenge lies in evaluating the derivatives $\nabla_{\boldsymbol{\theta}} P_j$, which requires differentiation under the integral sign in Equation (2.26). This, in turn, reduces to the computation of $\nabla_{\boldsymbol{\theta}} \varphi(u)$ (the gradient of the Heston characteristic function with respect to the model parameters).

The key innovation proposed by Cui et al. (2017) is the decomposition of the gradient of the characteristic function as

$$\nabla_{\boldsymbol{\theta}} \varphi(u) = \varphi(u) h(u),$$

where the vector $h(u)$ contains the partial derivatives of the logarithm of the characteristic function with respect to each model parameter. This decomposition, based on a stabilised representation of $\varphi(u)$, avoids numerical instabilities such as branch-switching discontinuities and enables efficient evaluation via symbolic differentiation.

The components of $h(u)$ are derived analytically using the parameter-separable structure of the characteristic function.

The following auxiliary terms are introduced to formulate the gradient:

$$\varphi(u;\tau) = \exp\left(iu \ln \frac{F}{S_t} - \frac{\kappa\,\bar{v}\,\rho\,\tau\,i\,u}{\sigma} - A + \frac{2\,\kappa\,\bar{v}}{\sigma^2} D \right), \tag{3.5}$$

where

$$F = S_t e^{(r-q)\tau} \quad \text{(forward price)},$$

$$\xi \coloneqq \kappa - \sigma\rho i u, \qquad d \coloneqq \sqrt{\xi^2 + \sigma^2(u^2 + iu)},$$

$$A_1 \coloneqq (u^2 + iu)\, v_0 \sinh\!\left(\tfrac{d\tau}{2}\right), \qquad A_2 \coloneqq d \cosh\!\left(\tfrac{d\tau}{2}\right) + \xi \sinh\!\left(\tfrac{d\tau}{2}\right),$$

$$A \coloneqq \frac{A_1}{A_2}, \quad B \coloneqq \frac{d\, e^{\kappa\tau/2}}{v_0 A_2}, \quad L \coloneqq \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0} e^{-d\tau},$$

$$D \coloneqq \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L.$$

Derivatives with respect to the initial variance v_0 and the long-run mean variance $\bar{v}$ are crucial for constructing the gradient and Hessian of the stabilised characteristic function. In particular, the derivative with respect to v_0 exhibits a singularity in the limit $v_0 \to 0$:

$$h_{v_0} = -\frac{A}{v_0} + \frac{2\,\kappa\,\bar{v}}{\sigma^2} \frac{\partial D}{\partial v_0}, \qquad \text{where} \quad \frac{\partial D}{\partial v_0} = -\frac{1}{v_0}.$$

Since A_2 is independent of v_0, the differentiation of A is considerably simplified.

The second-order derivative of the characteristic function with respect to v_0 is required for the Hessian:

$$\frac{\partial^2 \varphi}{\partial v_0^2} = \frac{A^2}{v_0^2}\,\varphi.$$

The complete derivation is provided in Appendix C.1.

The derivative with respect to the long-run mean variance $\bar{v}$ is

$$h_{\bar{v}} = \frac{2\,\kappa}{\sigma^2} D - \frac{\kappa\,\rho\,\tau\,i\,u}{\sigma},$$

with vanishing partial derivative

$$\frac{\partial h_{\bar{v}}}{\partial \bar{v}} = 0,$$

which implies

$$\frac{\partial^2 \varphi}{\partial \bar{v}^2} = h_{\bar{v}}^2\,\varphi.$$

The correlation parameter ρ enters the model non-linearly through the terms ξ and d, which propagate into the stabilised characteristic function. The first-order derivative of φ with respect to ρ is

$$\frac{\partial \varphi}{\partial \rho} = \varphi \, h_\rho,$$

where

$$h_\rho = -\frac{\partial A}{\partial \rho} + \frac{2\,\kappa\,\bar{v}}{\sigma^2 d}\left(\frac{\partial d}{\partial \rho} - \frac{d}{A_2}\frac{\partial A_2}{\partial \rho}\right) - \frac{\kappa\,\bar{v}\,\tau\,i\,u}{\sigma}.$$

Differentiating h_ρ with respect to ρ gives the second-order derivative of the characteristic function:

$$\frac{\partial^2 \varphi}{\partial \rho^2} = \varphi \, h_\rho^2 + \varphi \, \frac{\partial h_\rho}{\partial \rho},$$

with

$$\frac{\partial h_\rho}{\partial \rho} = -\frac{\partial^2 A}{\partial \rho^2} + \frac{2\,\kappa\,\bar{v}}{\sigma^2}\left[-\frac{1}{d^2}\frac{\partial d}{\partial \rho}\left(\frac{\partial d}{\partial \rho} - \frac{d}{A_2}\frac{\partial A_2}{\partial \rho}\right) + \frac{1}{d}\left(\frac{\partial^2 d}{\partial \rho^2} - \frac{\partial}{\partial \rho}\left(\frac{d}{A_2}\frac{\partial A_2}{\partial \rho}\right)\right)\right].$$

The evaluation of $\partial h_\rho / \partial \rho$ requires higher-order derivatives of the auxiliary functions A, A_1, A_2, g, and d. Since these expressions are algebraically extensive, the complete derivations are presented in Appendix C.2.

Second Derivative of φ with respect to κ

$$\frac{\partial \varphi}{\partial \kappa} = \varphi h_\kappa,$$

$$h_\kappa = \frac{1}{\sigma\, iu}\frac{\partial A}{\partial \rho} + \frac{2\bar{v}}{\sigma^2} D + \frac{2\kappa\bar{v}}{\sigma^2 B}\frac{\partial B}{\partial \kappa} - \frac{\bar{v}\,\rho\,\tau\,iu}{\sigma}.$$

The second derivative of the characteristic function φ with respect to the mean-reversion parameter κ takes the form

$$\frac{\partial^2 \varphi}{\partial \kappa^2} = \frac{\partial}{\partial \kappa}\left(\varphi h_\kappa\right) = \varphi h_\kappa^2 + \varphi \frac{\partial h_\kappa}{\partial \kappa}.$$

where the derivative of h_κ is expressed as

$$\frac{\partial h_\kappa}{\partial \kappa} = \frac{1}{\sigma\, iu}\frac{\partial^2 A}{\partial \rho\, \partial \kappa} + \frac{2\bar{v}}{\sigma^2}\left(\frac{\partial D}{\partial \kappa} + \frac{1}{B}\frac{\partial B}{\partial \kappa} - \frac{\kappa}{B^2}\left(\frac{\partial B}{\partial \kappa}\right)^2 + \frac{\kappa}{B}\frac{\partial^2 B}{\partial \kappa^2}\right). \tag{3.6}$$

The computation of (3.6) requires higher-order derivatives of the auxiliary functions

A, B, D, the key expressions are:

$$\begin{aligned}
\frac{\partial^2 A}{\partial \kappa^2} &= \frac{i}{\sigma u}\left(\frac{1}{A_2}\frac{\partial^2 A_1}{\partial \rho \partial \kappa} - \frac{A_1}{A_2^2}\frac{\partial^2 A_2}{\partial \rho \partial \kappa} - \frac{1}{A_2^2}\Big(\frac{\partial A_2}{\partial \kappa}\frac{\partial A_1}{\partial \rho} + \frac{\partial A_1}{\partial \kappa}\frac{\partial A_2}{\partial \rho}\Big) + \frac{2A_1}{A_2^3}\frac{\partial A_2}{\partial \kappa}\frac{\partial A_2}{\partial \rho}\right), \\
\frac{\partial D}{\partial \kappa} &= \frac{\xi}{d^2} + \frac{\tau}{2}\left(1 - \frac{\xi}{d}\right) - \frac{1}{L}\frac{\partial L}{\partial \kappa}, \\
\frac{\partial L}{\partial \kappa} &= \frac{1}{2v_0}\left[\left(\frac{\xi}{d} + 1\right) + \left(\frac{\xi}{d} - 1\right)e^{-d\tau} - \tau(d - \xi)\frac{\xi}{d}e^{-d\tau}\right], \\
\frac{\partial B}{\partial \kappa} &= \frac{i}{\sigma u}\frac{\partial B}{\partial \rho} + \frac{B\tau}{2}.
\end{aligned}$$

To evaluate $\partial^2 h_\kappa / \partial \kappa^2$, the required first and second order derivatives, including $\frac{\partial^2 B}{\partial \kappa^2}$, $\frac{\partial^2 d}{\partial \rho \partial \kappa}$, and those for A_1 and A_2, are presented in Appendix C.3 together with the intermediate terms on which they depend.

The sensitivity of the characteristic function φ with respect to the volatility-of-volatility parameter σ is given by

$$\frac{\partial \varphi}{\partial \sigma} = \varphi\, h_\sigma,$$

where the auxiliary function h_σ is defined as

$$h_\sigma = -\frac{\partial A}{\partial \sigma} - \frac{4\,\kappa\,\bar{v}}{\sigma^3}\,D + \frac{2\,\kappa\,\bar{v}}{\sigma^2\,d}\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\,\frac{\partial A_2}{\partial \sigma}\right) + \frac{\kappa\,\bar{v}\,\rho\,\tau\,i\,u}{\sigma^2}.$$

Applying the product rule, the second derivative of φ with respect to σ is expressed as

$$\frac{\partial^2 \varphi}{\partial \sigma^2} = \varphi\, h_\sigma^2 + \varphi\,\frac{\partial h_\sigma}{\partial \sigma}.$$

Then, the full derivative of h_σ is given by Then, the full derivative of h_σ is given by

$$\begin{aligned}
\frac{\partial h_\sigma}{\partial \sigma} = &-\frac{\partial^2 A}{\partial \sigma^2} + \frac{12\,\kappa\,\bar{v}}{\sigma^4}D - \frac{4\,\kappa\,\bar{v}}{\sigma^3}\frac{\partial D}{\partial \sigma} - \left(\frac{4\,\kappa\,\bar{v}}{\sigma^3 d} + \frac{2\,\kappa\,\bar{v}}{\sigma^2 d^2}\frac{\partial d}{\partial \sigma}\right)\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\frac{\partial A_2}{\partial \sigma}\right) \\
&+ \frac{2\,\kappa\,\bar{v}}{\sigma^2 d}\left[\frac{\partial^2 d}{\partial \sigma^2} - \frac{A_2\left(\frac{\partial d}{\partial \sigma}\right)^2 - d\,\frac{\partial d}{\partial \sigma}\,\frac{\partial A_2}{\partial \sigma}}{A_2^2} - \frac{d}{A_2}\frac{\partial^2 A_2}{\partial \sigma^2}\right] - \frac{2\,\kappa\,\bar{v}\,\rho\,\tau\,i\,u}{\sigma^3}.
\end{aligned}$$

The evaluation of $\frac{\partial h_\sigma}{\partial \sigma}$ requires several nested first and second order derivatives of the auxiliary functions A, A_1, A_2, d, D, and L. Since these expressions are algebraically extensive, the full derivations are provided in Appendix C.4.

3.7 Cross Derivatives Implementation and Interpretation

Appendix E presents the complete derivations of all second-order cross derivatives for the Heston characteristic function, including $\frac{\partial^2\varphi}{\partial v_0 \partial \bar{v}}$, $\frac{\partial^2\varphi}{\partial v_0 \partial \rho}$, $\frac{\partial^2\varphi}{\partial v_0 \partial \kappa}$, $\frac{\partial^2\varphi}{\partial v_0 \partial \sigma}$, $\frac{\partial^2\varphi}{\partial \bar{v} \partial \rho}$, $\frac{\partial^2\varphi}{\partial \bar{v} \partial \kappa}$, $\frac{\partial^2\varphi}{\partial \bar{v} \partial \sigma}$, $\frac{\partial^2\varphi}{\partial \rho \partial \kappa}$, $\frac{\partial^2\varphi}{\partial \rho \partial \sigma}$, and $\frac{\partial^2\varphi}{\partial \kappa \partial \sigma}$.

Successful numerical implementation hinges on several important considerations. When d becomes small, the Taylor expansions of the hyperbolic functions $\sinh(x) \approx x + x^3/6$ and $\cosh(x) \approx 1 + x^2/2$ maintain stability. Complex arithmetic requires careful handling of branch cuts in logarithmic terms, particularly in the D component. Computational efficiency improves by precomputing common terms such as $\frac{\sigma(u^2+iu)-\rho\xi\, iu}{d}$ and reusing them across derivative calculations.

Although algebraically complex, these derivatives are essential for advanced quantitative analysis, particularly in volatility surface modelling. They reveal critical phenomena such as the persistence of volatility skew and the speed at which volatility reverts to its long-run mean by capturing nuanced interactions between model parameters. These interactions influence the dynamics of option pricing across a variety of maturities and strikes. The second-order effects quantified by these derivatives are especially valuable in three applications: in calibration routines, they enhance the robustness of gradient-based optimisation; in risk management, they offer precise evaluation of joint parameter dependencies; and in model validation, they provide tools for testing the internal consistency of the Heston framework.

Concerning implementation, the vector $h(u)$ consists of interdependent components, which allows common intermediate terms such as $\partial A/\partial\rho$ and $\partial D/\partial\sigma$ to be computed once per integration node. This structure enables a fully vectorised numerical quadrature scheme, as proposed in Algorithm 3.1 of Cui et al. (2017), which evaluates all required gradients simultaneously for the entire parameter set. Additionally, the analytical gradient enables efficient calibration of the Heston model using methods such as Levenberg–Marquardt, typically converging in 10–15 iterations. This enhances parameter identifiability, particularly in flat regions of the objective surface, which are caused by parameter interdependencies such as the well-known κ–σ trade-off (Andersen and Piterbarg, 2010). The empirical findings in Section 5.8 indicate that the estimated correlation between the gradients of κ and σ is -0.56.

The proposed analytical method offers significant computational advantages in calibrating the Heston model. By circumventing the need for repeated function evaluations along individual parameter directions, it substantially reduces computational complexity and enables efficient gradient computation. In Table 3 of Cui et al. (2017),

the computational cost is reduced from 15.8 to 1.0 units, and the number of integral evaluations is reduced from 800 to 80, resulting in a speed increase. Beyond these performance gains, the method eliminates truncation and rounding errors inherent in numerical differentiation, thereby enhancing accuracy. It also preserves the smoothness and continuity of the stabilised Heston characteristic function, which contributes to improved numerical stability. Furthermore, the availability of analytic gradients enhances robustness in pricing frameworks, particularly by reducing the method's sensitivity to damping parameters in FFT-based calibration routines. Collectively, these properties render the method well-suited for real-time applications in trading systems, where both speed and precision are critical.

The following theorems and lemmas establish the theoretical foundation of the proposed analytical method, specifically the closed-form gradient and Hessian expressions of the Heston call price with respect to model parameters. These expressions enable efficient and stable calibration by providing exact curvature information, and are consistent with the formulation of Cui et al. (2017). The closed-form derivatives support the use of second-order optimisation methods, such as trust-region schemes. The following theorems formalise these results and offer convergence guarantees for the resulting algorithm.

Theorem 2 (Analytical Gradient and Hessian)**.** *Under the Heston model with parameters* $\boldsymbol{\theta} = (v_0, \bar{v}, \rho, \kappa, \sigma)^T$*, the gradient* $\nabla_{\boldsymbol{\theta}} C(\boldsymbol{\theta}; K, \tau)$ *and Hessian* $\nabla^2_{\boldsymbol{\theta}} C(\boldsymbol{\theta}; K, \tau)$ *of the call price* C *admit closed-form expressions:*

1. ***Gradient****:*
$$\frac{\partial C}{\partial \theta_l} = \frac{e^{-r\tau}}{2\pi} \int_0^\infty \Re \left[e^{-iu \ln K} \frac{\partial \varphi(u;\tau)}{\partial \theta_l} \mathcal{K}(u) \right] \mathrm{d}u,$$
where $\mathcal{K}(u)$ *is the Fourier kernel and* $\varphi(u;\tau)$ *is the characteristic function of the log-asset price.*

2. ***Hessian****:*
$$\frac{\partial^2 C}{\partial \theta_l \partial \theta_k} = \frac{e^{-r\tau}}{2\pi} \int_0^\infty \Re \left[e^{-iu \ln K} \left(\frac{\partial^2 \varphi}{\partial \theta_l \partial \theta_k} + \frac{\partial \varphi}{\partial \theta_l} \frac{\partial \varphi}{\partial \theta_k} \right) \mathcal{K}(u) \right] \mathrm{d}u.$$

Having derived the necessary analytic expressions for both first- and second-order derivatives, we next examine the convergence properties of a trust-region calibration algorithm that leverages this information. The following result guarantees global and quadratic convergence when curvature information is incorporated precisely.

Theorem 3 (Convergence of Trust-Region Method (Nocedal and Wright, 2006; Conn et al., 2000))**.** *Let $f(\boldsymbol{\theta})$ denote the Heston model calibration objective function, and assume its Hessian $\nabla^2 f(\boldsymbol{\theta})$ is Lipschitz-continuous in a neighbourhood of a local minimiser $\boldsymbol{\theta}^*$. Suppose a trust-region method is applied such that:*

1. *The gradient $\nabla f(\boldsymbol{\theta}_k)$ and Hessian $\nabla^2 f(\boldsymbol{\theta}_k)$ are computed using Theorem 2;*
2. *At each iteration k, the trial step p_k is obtained by solving the subproblem:*

$$\min_{\|p\| \leq \Delta_k} \quad m_k(p),$$

where the local model $m_k(p)$ is defined as

$$m_k(p) \coloneqq f(\boldsymbol{\theta}_k) + \nabla f(\boldsymbol{\theta}_k)^\top p + \tfrac{1}{2} p^\top \nabla^2 f(\boldsymbol{\theta}_k) p;$$

3. *The trust-region radius is adapted according to*

$$\Delta_k = \min\left(\Delta_0, \gamma \left\|\nabla f(\boldsymbol{\theta}_k)\right\|\right),$$

for constants $\Delta_0 > 0$ and $\gamma > 0$;

4. *The step p_k is accepted only if the ratio*

$$\rho_k \coloneqq \frac{f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}_k + p_k)}{m_k(0) - m_k(p_k)}$$

satisfies $\rho_k > \eta$, for some acceptance threshold $\eta \in (0, 1)$.

Then the sequence $\{\boldsymbol{\theta}_k\}$ generated by the algorithm satisfies

$$\lim_{k \to \infty} \inf \left\|\nabla f(\boldsymbol{\theta}_k)\right\| = 0,$$

and converges quadratically to a local minimiser $\boldsymbol{\theta}^*$ whenever $\boldsymbol{\theta}_k$ enters a sufficiently small neighborhood around $\boldsymbol{\theta}^*$. Specifically, there exists a constant $\beta > 0$ such that

$$\left\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}^*\right\| \leq \beta \left\|\boldsymbol{\theta}_k - \boldsymbol{\theta}^*\right\|^2.$$

Remark. The model agreement ratio ρ_k compares the actual reduction $f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}_k + p_k)$ with the model-predicted reduction $m_k(0) - m_k(p_k)$. Since $m_k(0) = f(\boldsymbol{\theta}_k)$, the denominator represents the predicted decrease of the quadratic model at iteration k.

A vanishing denominator occurs only when the model predicts no reduction along p_k, that is,

$$m_k(0) - m_k(p_k) = 0 \quad \Longleftrightarrow \quad \nabla f(\boldsymbol{\theta}_k)^\top p_k + \tfrac{1}{2} p_k^\top \nabla^2 f(\boldsymbol{\theta}_k) p_k = 0.$$

In this case, ρ_k is undefined and the step is treated as unproductive, typically leading to rejection of the step and contraction of the trust region.

Algorithm 3.1 Hessian-Aware Trust-Region Calibration

Initialize: $\boldsymbol{\theta}_0$, initial trust-region radius $\Delta_0 > 0$, maximum radius $\Delta_{\max} > \Delta_0$, thresholds $0 < \eta_1 \leq \eta_2 < 1$, scaling parameters $\gamma_1 < 1 < \gamma_2$, convergence tolerance $\varepsilon > 0$

for $k = 0, 1, 2, \ldots$ until $\|\nabla f(\boldsymbol{\theta}_k)\| < \varepsilon$ **do**

Compute gradient $\nabla f(\boldsymbol{\theta}_k)$ and Hessian $\nabla^2 f(\boldsymbol{\theta}_k)$ via Theorem 2

Define quadratic model:

$$m_k(p) \coloneqq f(\boldsymbol{\theta}_k) + \nabla f(\boldsymbol{\theta}_k)^\top p + \tfrac{1}{2} p^\top \nabla^2 f(\boldsymbol{\theta}_k) p$$

Solve trust-region subproblem:

$$p_k \coloneqq \arg\min_{\|p\| \leq \Delta_k} m_k(p)$$

Evaluate model agreement ratio:

$$\rho_k \coloneqq \frac{f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}_k + p_k)}{m_k(0) - m_k(p_k)}, \qquad \text{where } m_k(0) = f(\boldsymbol{\theta}_k).$$

if $\rho_k \geq \eta_1$ **then**

Accept step: $\boldsymbol{\theta}_{k+1} \coloneqq \boldsymbol{\theta}_k + p_k$

else

Reject step: $\boldsymbol{\theta}_{k+1} \coloneqq \boldsymbol{\theta}_k$

end if

Update trust-region radius:

$$\Delta_{k+1} \coloneqq \begin{cases} \gamma_1 \Delta_k & \text{if } \rho_k < \eta_1 \\ \min(\gamma_2 \Delta_k, \Delta_{\max}) & \text{if } \rho_k \geq \eta_2 \text{ and } \|p_k\| = \Delta_k \\ \Delta_k & \text{otherwise} \end{cases}$$

end for

Algorithm 3.1 integrates the exact derivatives provided by Theorem 2, and its convergence behaviour is guaranteed by Theorem 3. Together, these results form the theoretical foundation for curvature-aware optimisation strategies in Heston model cal-

ibration. Although the empirical implementation in later chapters utilises first-order derivatives only, this framework supports future extensions to second-order schemes with provable numerical properties.

3.8 Derivatives with respect to r, S_0, and τ

While the primary focus of this chapter is the analytical differentiation of the Heston pricing function with respect to model parameters $\boldsymbol{\theta}$, we also derive closed-form expressions for the first- and second-order derivatives of the call price with respect to time to maturity τ, spot, and the risk-free interest rate r. These quantities correspond to four key option Greeks. Delta, which measures the sensitivity of the option value to the underlying asset, Gamma, which captures the curvature with respect to the underlying, Theta, which measures sensitivity to time to maturity τ, and Rho, which quantifies sensitivity to interest rate changes. Accurate computation of these derivatives is essential for risk-neutral valuation, hedging strategies, and regulatory stress testing.

The analytic derivation of these derivatives enhances both precision and interpretability compared to finite-difference approximations, particularly in regions where time decay or interest rate shifts exert nonlinear effects on prices. Furthermore, the second-order derivatives $\partial^2 C/\partial\tau^2$ and $\partial^2 C/\partial r^2$ provide curvature information relevant for higher-order risk management and convexity adjustment in interest rate modelling. Although not used in the calibration algorithm, these results enhance the analytical toolkit for model-based pricing and risk assessment and are included for completeness in the following sections.

Derivatives with respect to risk-free rate r

The call price C is given by

$$C = \frac{1}{2}\left(Se^{-q\tau} - Ke^{-r\tau}\right) + \frac{e^{-r\tau}}{\pi}\left[SI_1 - KI_2\right], \tag{3.7}$$

where

$$I_1 = \int_0^\infty \operatorname{Re}\left[e^{-iu\ln(K/S)}\frac{\varphi(u-i)}{iu}\right]\mathrm{d}u, \tag{3.8}$$

$$I_2 = \int_0^\infty \operatorname{Re}\left[e^{-iu\ln(K/S)}\frac{\varphi(u)}{iu}\right]\mathrm{d}u. \tag{3.9}$$

First derivative with respect to r (rho).

Starting from Equation (3.7), differentiate term by term with respect to r (note $\frac{\partial}{\partial r}e^{-r\tau} = -\tau e^{-r\tau}$):

$$\begin{aligned}\frac{\partial C}{\partial r} &= \tfrac{1}{2}\,\frac{\partial}{\partial r}\big(Se^{-q\tau} - Ke^{-r\tau}\big) + \frac{\partial}{\partial r}\left(\frac{e^{-r\tau}}{\pi}\,\big(SI_1 - KI_2\big)\right)\\ &= \tfrac{1}{2}\,\big(0 - (-\tau K e^{-r\tau})\big) + \frac{1}{\pi}\,\big(-\tau e^{-r\tau}(SI_1 - KI_2) + e^{-r\tau}\big(S\tfrac{\partial I_1}{\partial r} - K\tfrac{\partial I_2}{\partial r}\big)\big)\\ &= \boxed{\tfrac{1}{2}\,\tau K e^{-r\tau} - \frac{\tau e^{-r\tau}}{\pi}\,(SI_1 - KI_2) + \frac{e^{-r\tau}}{\pi}\left(S\frac{\partial I_1}{\partial r} - K\frac{\partial I_2}{\partial r}\right).}\end{aligned} \tag{3.10}$$

Derivatives of I_1, I_2 with respect to r. Using Equation (3.9) and the (3.5) (forward–moneyness) characteristic function, we have

$$\boxed{\frac{\partial \varphi(u;\tau)}{\partial r} = (iu\,\tau)\,\varphi(u;\tau), \qquad \frac{\partial \varphi(u-i;\tau)}{\partial r} = (i(u-i)\,\tau)\,\varphi(u-i;\tau).}$$

Hence,

$$\begin{aligned}\frac{\partial I_1}{\partial r} &= \int_0^\infty \mathrm{Re}\left\{e^{-\,iu\ln(K/S)}\,\frac{1}{iu}\,\frac{\partial \varphi(u-i;\tau)}{\partial r}\right\}\,\mathrm{d}u\\ &= \tau\int_0^\infty \mathrm{Re}\left\{e^{-\,iu\ln(K/S)}\,\frac{u-i}{u}\,\varphi(u-i;\tau)\right\}\,\mathrm{d}u.\end{aligned}$$

i.e.

$$\boxed{\frac{\partial I_1}{\partial r} = \tau\int_0^\infty \Re\left\{e^{-iu\ln(K/S)}\,\frac{u-i}{u}\,\varphi(u-i;\tau)\right\}du.} \tag{3.11}$$

Similarly,

$$\frac{\partial I_2}{\partial r} = \int_0^\infty \Re\left\{e^{-iu\ln(K/S)}\frac{1}{iu}\,\frac{\partial \varphi(u;\tau)}{\partial r}\right\}du = \tau\int_0^\infty \Re\big\{e^{-iu\ln(K/S)}\,\varphi(u;\tau)\big\}\,du,$$

i.e.

$$\boxed{\frac{\partial I_2}{\partial r} = \tau \int_0^\infty \Re\big\{e^{-iu\ln(K/S)}\,\varphi(u;\tau)\big\}\,du.} \tag{3.12}$$

Second derivative with respect to r ($\partial^2 C/\partial r^2$). Starting from Equation (3.10), differentiate once more w.r.t. r (using $\frac{\partial}{\partial r}e^{-r\tau} = -\tau e^{-r\tau}$):

$$\begin{aligned}
\frac{\partial^2 C}{\partial r^2} &= \tfrac{1}{2}\,\tau\,\frac{\partial}{\partial r}\big(Ke^{-r\tau}\big) - \frac{\partial}{\partial r}\Big(\frac{\tau e^{-r\tau}}{\pi}\,(SI_1 - KI_2)\Big) + \frac{\partial}{\partial r}\Big(\frac{e^{-r\tau}}{\pi}\,(SI_1^{(r)} - KI_2^{(r)})\Big) \\
&= \tfrac{1}{2}\,\tau\big(-\tau Ke^{-r\tau}\big) - \frac{1}{\pi}\Big(-\tau^2 e^{-r\tau}(SI_1 - KI_2) + \tau e^{-r\tau}(SI_1^{(r)} - KI_2^{(r)})\Big) \\
&\quad + \frac{1}{\pi}\Big(-\tau e^{-r\tau}(SI_1^{(r)} - KI_2^{(r)}) + e^{-r\tau}(SI_1^{(rr)} - KI_2^{(rr)})\Big)
\end{aligned}$$

$$\boxed{\begin{aligned}
&= -\tfrac{1}{2}\,\tau^2 K\,e^{-r\tau} + \frac{\tau^2 e^{-r\tau}}{\pi}\,(S\,I_1 - K\,I_2) \\
&\quad - \frac{2\tau e^{-r\tau}}{\pi}\,\big(S\,I_1^{(r)} - K\,I_2^{(r)}\big) + \frac{e^{-r\tau}}{\pi}\,\big(S\,I_1^{(rr)} - K\,I_2^{(rr)}\big)
\end{aligned}}.$$

Here, we used the shorthand

$$I_j^{(r)} \coloneqq \frac{\partial I_j}{\partial r}, \qquad I_j^{(rr)} \coloneqq \frac{\partial^2 I_j}{\partial r^2}.$$

Inner derivatives of I_1, I_2 with respect to r. we have the first derivatives Equation (3.11) and (3.12) (from first derivative). Differentiating once more (the prefactors $(u-i)/u$ and 1 are r-independent):

$$\boxed{I_1^{(rr)} = \tau^2 \int_0^\infty \Re\Big\{\,i\,\frac{(u-i)^2}{u}\,e^{-iu\ln(K/S)}\,\varphi(u-i;\tau)\Big\}\,du,}$$

$$\boxed{I_2^{(rr)} = \tau^2 \int_0^\infty \Re\big\{\,i\,u\,e^{-iu\ln(K/S)}\,\varphi(u;\tau)\big\}\,du.} \tag{3.13}$$

Derivatives with respect to S_0: Delta Δ and Gamma Γ

We consider the Heston call price in Equation (3.7) (forward–moneyness) representation. In this formulation the characteristic function depends on $\ln(F/S)$ with

$F = S\,e^{(r-q)\tau}$, hence $F/S = e^{(r-q)\tau}$ is independent of S and

$$\frac{\partial \varphi}{\partial S} = 0.$$

Therefore, all S-dependence arises from the explicit factor S in front of I_1 and from the phase $e^{-iu\ln(K/S)}$ inside I_1, I_2.

First derivative (Δ). Differentiating term by term,

$$\begin{aligned}\frac{\partial C}{\partial S} &= \frac{1}{2}\frac{\partial}{\partial S}\big(Se^{-q\tau} - Ke^{-r\tau}\big) + \frac{e^{-r\tau}}{\pi}\frac{\partial}{\partial S}\big(S\,I_1(S) - K\,I_2(S)\big) \\ &= \frac{1}{2}\,e^{-q\tau} + \frac{e^{-r\tau}}{\pi}\Big(\underbrace{\frac{\partial}{\partial S}(S\,I_1)}_{=\,I_1 + S\,\partial I_1/\partial S} - \underbrace{\frac{\partial}{\partial S}(K\,I_2)}_{=\,K\,\partial I_2/\partial S}\Big) \\ &= \boxed{\frac{1}{2}\,e^{-q\tau} + \frac{e^{-r\tau}}{\pi}\Big[\,I_1(S) + S\,\frac{\partial I_1}{\partial S} - K\,\frac{\partial I_2}{\partial S}\Big]}.\end{aligned}$$

To express the second derivatives of I_j explicitly, differentiate the first-derivative formulas once more. Define for brevity

$$E(u; K, S) = e^{-iu\ln(K/S)}.$$

Then

$$\frac{\partial I_1}{\partial S} = \frac{1}{S}\int_0^\infty \Re\{E\,\varphi(u-i)\}\,du, \qquad \frac{\partial I_2}{\partial S} = \frac{1}{S}\int_0^\infty \Re\{E\,\varphi(u)\}\,du,$$

and, using $\frac{\partial E}{\partial S} = \frac{iu}{S}E$, The inner derivatives (using $\partial\varphi/\partial S = 0$ and $\frac{\partial}{\partial S}e^{-iu\ln(K/S)} = \frac{iu}{S}e^{-iu\ln(K/S)}$) are

$$\boxed{\begin{aligned}\frac{\partial I_1}{\partial S} &= \frac{1}{S}\int_0^\infty \mathrm{Re}\big\{e^{-\,iu\ln(K/S)}\,\varphi(u-i)\big\}\,\mathrm{d}u, &(3.14)\\ \frac{\partial I_2}{\partial S} &= \frac{1}{S}\int_0^\infty \mathrm{Re}\big\{e^{-\,iu\ln(K/S)}\,\varphi(u)\big\}\,\mathrm{d}u. &(3.15)\end{aligned}}$$

Second derivative (Γ). Starting from the boxed Δ expression above, differentiate once more:

$$\begin{aligned}\Gamma \;=\; \frac{\partial^2 C}{\partial S^2} &= \frac{e^{-r\tau}}{\pi}\left[\frac{\partial}{\partial S} I_1 + \frac{\partial}{\partial S}\Big(S\frac{\partial I_1}{\partial S}\Big) - K\,\frac{\partial^2 I_2}{\partial S^2}\right] \\ &= \frac{e^{-r\tau}}{\pi}\left[\frac{\partial I_1}{\partial S} + \frac{\partial I_1}{\partial S} + S\,\frac{\partial^2 I_1}{\partial S^2} - K\,\frac{\partial^2 I_2}{\partial S^2}\right] \\ &= \boxed{\frac{e^{-r\tau}}{\pi}\left[2\,\frac{\partial I_1}{\partial S} + S\,\frac{\partial^2 I_1}{\partial S^2} - K\,\frac{\partial^2 I_2}{\partial S^2}\right]}.\end{aligned}$$

$$\boxed{\frac{\partial^2 I_1}{\partial S^2} = -\frac{1}{S^2}\int_0^\infty \Re\{E\,\varphi(u-i)\}\,du + \frac{1}{S^2}\int_0^\infty \Re\{\,iu\,E\,\varphi(u-i)\}\,du,}$$

$$\boxed{\frac{\partial^2 I_2}{\partial S^2} = -\frac{1}{S^2}\int_0^\infty \Re\{E\,\varphi(u)\}\,du + \frac{1}{S^2}\int_0^\infty \Re\{\,iu\,E\,\varphi(u)\}\,du.}$$

Substituting these into the boxed Γ formula yields the fully explicit second-derivative (gamma) in terms of the same Fourier kernels used for Δ.

Analytical Derivatives of Maturity τ

Setup and notation. We work with the Equation (3.7) (forward–moneyness) representation

$$C(\tau) = \tfrac{1}{2}\big(Se^{-q\tau} - Ke^{-r\tau}\big) + \frac{e^{-r\tau}}{\pi}\,\big(SI_1(\tau) - KI_2(\tau)\big),$$

$$\begin{aligned}I_1(\tau) &= \int_0^\infty \mathrm{Re}\left\{e^{-\,iu\ln(K/S)}\,\frac{\varphi(u-i;\tau)}{i\,u}\right\}\,\mathrm{d}u, \\ I_2(\tau) &= \int_0^\infty \mathrm{Re}\left\{e^{-\,iu\ln(K/S)}\,\frac{\varphi(u;\tau)}{i\,u}\right\}\,\mathrm{d}u.\end{aligned}$$

$$\varphi(u;\tau) = \exp\left(iu\ln\frac{F}{S_0} - \frac{\kappa\bar{v}\rho\,\tau\,iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right), \qquad F = S_0\,e^{(r-q)\tau}.$$

In this form, the *phase* $e^{-iu\ln(K/S)}$ is τ-independent; all τ-dependence of I_1, I_2 enters via $\varphi(\cdot;\tau)$.

Differentiate C with respect to τ:

$$\frac{\partial C}{\partial \tau} = \tfrac{1}{2}\left(-q\,Se^{-q\tau} + r\,Ke^{-r\tau}\right) + \frac{\partial}{\partial \tau}\left(\frac{e^{-r\tau}}{\pi}\right)\left(SI_1 - KI_2\right) + \frac{e^{-r\tau}}{\pi}\left(S\frac{\partial I_1}{\partial \tau} - K\frac{\partial I_2}{\partial \tau}\right)$$

$$= \boxed{\tfrac{1}{2}\left(-q\,Se^{-q\tau} + r\,Ke^{-r\tau}\right) + \frac{e^{-r\tau}}{\pi}\left(S\,\frac{\partial I_1}{\partial \tau} - K\,\frac{\partial I_2}{\partial \tau} - r\left(SI_1 - KI_2\right)\right)}. \quad (3.16)$$

Inner τ-derivatives of the integrals. Since the phase is τ-independent,

$$\boxed{\frac{\partial I_1}{\partial \tau} = \int_0^\infty \Re\left\{e^{-iu\ln(K/S)}\,\frac{1}{iu}\,\frac{\partial \varphi(u-i;\tau)}{\partial \tau}\right\} du,}$$

$$\boxed{\frac{\partial I_2}{\partial \tau} = \int_0^\infty \Re\left\{e^{-iu\ln(K/S)}\,\frac{1}{iu}\,\frac{\partial \varphi(u;\tau)}{\partial \tau}\right\} du.}$$

Derivative of the characteristic function. Let

$$X(u;\tau) = iu\ln\frac{F}{S_0} - \frac{\kappa\bar{v}\rho\,\tau\,iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D, \qquad \varphi = \exp(X).$$

Then $\dfrac{\partial \varphi}{\partial \tau} = \dfrac{\partial X}{\partial \tau}\,\varphi$ with

$$\frac{\partial}{\partial \tau}\left(iu\ln\frac{F}{S_0}\right) = iu\,\frac{\partial}{\partial \tau}\big((r-q)\tau\big) = iu\,(r-q), \qquad \frac{\partial}{\partial \tau}\left(-\frac{\kappa\bar{v}\rho\,\tau\,iu}{\sigma}\right) = -\frac{\kappa\bar{v}\rho\,iu}{\sigma}.$$

Hence

$$\boxed{\frac{\partial \varphi}{\partial \tau} = \left(iu(r-q) - \frac{\kappa\bar{v}\rho\,iu}{\sigma} - \frac{\partial A}{\partial \tau} + \frac{2\kappa\bar{v}}{\sigma^2}\,\frac{\partial D}{\partial \tau}\right)\varphi(u;\tau).}$$

Components $\partial_\tau A$ and $\partial_\tau D$. Write $A = A_1/A_2$ with

$$A_1 = (u^2 + iu)\,\sinh\left(\frac{d\tau}{2}\right), \qquad A_2 = \frac{d}{v_0}\cosh\left(\frac{d\tau}{2}\right) + \frac{\xi}{v_0}\sinh\left(\frac{d\tau}{2}\right),$$

where $\xi = \kappa - \sigma\rho\,iu$ and $d = \sqrt{\xi^2 + \sigma^2(u^2+iu)}$ (constant in τ). Then

$$A_1' = \frac{\partial A_1}{\partial \tau} = (u^2+iu)\,\frac{d}{2}\cosh\left(\frac{d\tau}{2}\right), \quad A_2' = \frac{\partial A_2}{\partial \tau} = \frac{d^2}{2v_0}\sinh\left(\frac{d\tau}{2}\right) + \frac{d\,\xi}{2v_0}\cosh\left(\frac{d\tau}{2}\right),$$

$$\boxed{\frac{\partial A}{\partial \tau} = \frac{A_1' A_2 - A_1 A_2'}{A_2^2}.}$$

For D, use the stabilized form

$$D(\tau) = \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L(\tau), \qquad L(\tau) = \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}\, e^{-d\tau}.$$

Then

$$\frac{\partial D}{\partial \tau} = \frac{\kappa - d}{2} - \frac{1}{L}\frac{\partial L}{\partial \tau}, \qquad \frac{\partial L}{\partial \tau} = -\frac{d(d-\xi)}{2v_0}\, e^{-d\tau},$$

and therefore

$$\boxed{\frac{\partial D}{\partial \tau} = \frac{\kappa - d}{2} + \frac{d(d-\xi)\, e^{-d\tau}}{(d+\xi) + (d-\xi)\, e^{-d\tau}}.}$$

Combine the pieces to obtain the final expression

$$\boxed{\frac{\partial C}{\partial \tau} = \tfrac{1}{2}\left(-q\, S e^{-q\tau} + r\, K e^{-r\tau}\right) + \frac{e^{-r\tau}}{\pi}\Big[S\,\frac{\partial I_1}{\partial \tau} - K\,\frac{\partial I_2}{\partial \tau} - r\,(S I_1 - K I_2)\Big],}$$

with $\partial_\tau I_j$ given above and $\partial_\tau \varphi$ supplied by the boxed formulas for $\partial_\tau A$ and $\partial_\tau D$.

Second derivative with respect to τ. Starting from the compact first–derivative expression (3.16) differentiate again with respect to τ. Term by term:

$$\begin{aligned}\frac{\partial^2 C}{\partial \tau^2} &= \tfrac{1}{2}\left(q^2 S e^{-q\tau} - r^2 K e^{-r\tau}\right) \\ &\quad + \frac{\partial}{\partial \tau}\left(\frac{e^{-r\tau}}{\pi}\right)\left(S\tfrac{\partial I_1}{\partial \tau} - K\tfrac{\partial I_2}{\partial \tau} - r(S I_1 - K I_2)\right) \\ &\quad + \frac{e^{-r\tau}}{\pi}\Big(S\tfrac{\partial^2 I_1}{\partial \tau^2} - K\tfrac{\partial^2 I_2}{\partial \tau^2} - r\,(S\tfrac{\partial I_1}{\partial \tau} - K\tfrac{\partial I_2}{\partial \tau})\Big).\end{aligned}$$

Using $\frac{\partial}{\partial \tau}(e^{-r\tau}) = -r e^{-r\tau}$, this becomes

$$\boxed{\begin{aligned}\frac{\partial^2 C}{\partial \tau^2} &= \tfrac{1}{2}\left(q^2 S e^{-q\tau} - r^2 K e^{-r\tau}\right) - \frac{r e^{-r\tau}}{\pi}\Big(S\tfrac{\partial I_1}{\partial \tau} - K\tfrac{\partial I_2}{\partial \tau} - r(S I_1 - K I_2)\Big) \\ &\quad + \frac{e^{-r\tau}}{\pi}\Big(S\tfrac{\partial^2 I_1}{\partial \tau^2} - K\tfrac{\partial^2 I_2}{\partial \tau^2} - r(S\tfrac{\partial I_1}{\partial \tau} - K\tfrac{\partial I_2}{\partial \tau})\Big).\end{aligned}}$$

Since I_1, I_2 depend on τ only via the characteristic function,

$$\boxed{\frac{\partial^2 I_1}{\partial \tau^2} = \int_0^\infty \Re\left\{ e^{-iu\ln(K/S)}\, \frac{1}{iu}\, \frac{\partial^2 \varphi(u-i;\tau)}{\partial \tau^2} \right\} du,}$$

$$\boxed{\frac{\partial^2 I_2}{\partial \tau^2} = \int_0^\infty \Re\left\{ e^{-iu\ln(K/S)}\, \frac{1}{iu}\, \frac{\partial^2 \varphi(u;\tau)}{\partial \tau^2} \right\} du.}$$

The explicit forms of $\frac{\partial^2 \varphi}{\partial \tau^2}$ were derived in the previous subsection, combining $\frac{\partial^2 X}{\partial \tau^2}$ and $\left(\frac{\partial X}{\partial \tau}\right)^2$.

For the second derivatives we proceed analogously. Recall that

$$\frac{\partial \varphi}{\partial \tau} = \frac{\partial X}{\partial \tau}\,\varphi, \qquad X(u;\tau) = iu \ln \frac{F}{S_0} - \frac{\kappa \bar{v} \rho\, \tau\, iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2} D.$$

Thus,

$$\frac{\partial^2 \varphi}{\partial \tau^2} = \left(\frac{\partial^2 X}{\partial \tau^2} + \left(\frac{\partial X}{\partial \tau}\right)^2 \right) \varphi.$$

Second derivatives of the auxiliary functions. For $A = A_1/A_2$ with

$$A_1 = (u^2 + iu)\, \sinh\left(\tfrac{d\tau}{2}\right), \qquad A_2 = \tfrac{d}{v_0} \cosh\left(\tfrac{d\tau}{2}\right) + \tfrac{\xi}{v_0} \sinh\left(\tfrac{d\tau}{2}\right),$$

we have

$$A_1'' = \tfrac{d^2}{4}(u^2 + iu)\, \sinh\left(\tfrac{d\tau}{2}\right),$$
$$A_2'' = \tfrac{d^3}{4v_0} \cosh\left(\tfrac{d\tau}{2}\right) + \tfrac{\xi d^2}{4v_0} \sinh\left(\tfrac{d\tau}{2}\right).$$

Hence,

$$\boxed{\frac{\partial^2 A}{\partial \tau^2} = \frac{1}{A_2^3}\Big[(A_1'' A_2 - A_1 A_2'') A_2 - 2(A_1' A_2 - A_1 A_2') A_2' \Big].}$$

For $D(\tau) = \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L(\tau)$ with

$$L(\tau) = \tfrac{d+\xi}{2v_0} + \tfrac{d-\xi}{2v_0} e^{-d\tau},$$

we already had

$$\frac{\partial D}{\partial \tau} = \tfrac{\kappa - d}{2} + \frac{d(d-\xi)e^{-d\tau}}{(d+\xi) + (d-\xi)e^{-d\tau}}.$$

Differentiating again yields

$$\boxed{\frac{\partial^2 D}{\partial \tau^2} = \left(\frac{\partial D}{\partial \tau} - \tfrac{\kappa - d}{2}\right)\left(-d + \frac{d(d-\xi)e^{-d\tau}}{(d+\xi) + (d-\xi)e^{-d\tau}}\right).}$$

Second derivative of X. Putting these pieces together,

$$\boxed{\frac{\partial^2 X}{\partial \tau^2} = -\frac{\partial^2 A}{\partial \tau^2} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial^2 D}{\partial \tau^2}.}$$

Therefore,

$$\boxed{\frac{\partial^2 \varphi}{\partial \tau^2} = \left(\frac{\partial^2 X}{\partial \tau^2} + \left[iu(r-q) - \frac{\kappa\bar{v}\rho\, iu}{\sigma} - \frac{\partial A}{\partial \tau} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \tau}\right]^2\right)\varphi(u;\tau).}$$

The full derivations of $\partial_\tau A$, $\partial_\tau^2 A$, $\partial_\tau D$, and $\partial_\tau^2 D$ are provided in Appendix D.

Chapter 4

Data Collection and Preparation

4.1 Introduction and Data Overview

This chapter describes the dataset and data acquisition pipeline utilised in this research, including the methodology employed to collect and prepare real market data. The option price data used in this thesis are end-of-day quotes for all listed Apple Inc. (AAPL) call and put options obtained from `https://www.marketdata.app`, with complementary inputs such as U.S. Treasury yield curves to represent the term structure of risk-free interest rates. Dividend payments, which affect the pricing of American and European options differently, were sourced from official financial records and incorporated into the calibration framework.

The final dataset comprises approximately 1.5 million call and put option observations for Apple Inc. from 2 January 2018 to 3 July 2024. This dataset forms the empirical foundation for all subsequent model development and calibration. The dataset spans about 2,000 different strike prices and multiple maturities each trading day, capturing a wide variety of market conditions including volatility peaks caused by COVID-19 (e.g. March 2020) and periods of elevated uncertainty in 2022 (e.g. Russia–Ukraine conflict).

In the dataset, each option contract is defined by a fixed strike price and expiration date and remains active in the market until expiry. Contracts vary in their lifespan from a few days to as long as two years, but all may be present simultaneously across trading days, contributing to a dense and overlapping structure of market activity. To manage this complexity, the data are organised into a matrix structure where rows correspond to strike prices, columns to expiration dates, and each matrix cell contains the complete trading history of the associated option. This structure facilitates both individual contract analysis and cross-sectional comparisons across trading days.

The dataset was used to compute the primary option sensitivities and the implied volatility. These derivatives, commonly referred to as the Greeks, describe the re-

sponsiveness of option prices to underlying market variables. Next, the relationship between implied volatility and log-moneyness (the natural logarithm of the strike price divided by the spot price) was plotted to further explore the structure of the data. Implied volatilities were derived using a MATLAB implementation of the Black–Scholes model that numerically inverts observed market prices. The resulting plot (Figure 4.1) reveals a clustering of options around the at-the-money region and provides insight into how volatility behaves across moneyness levels.

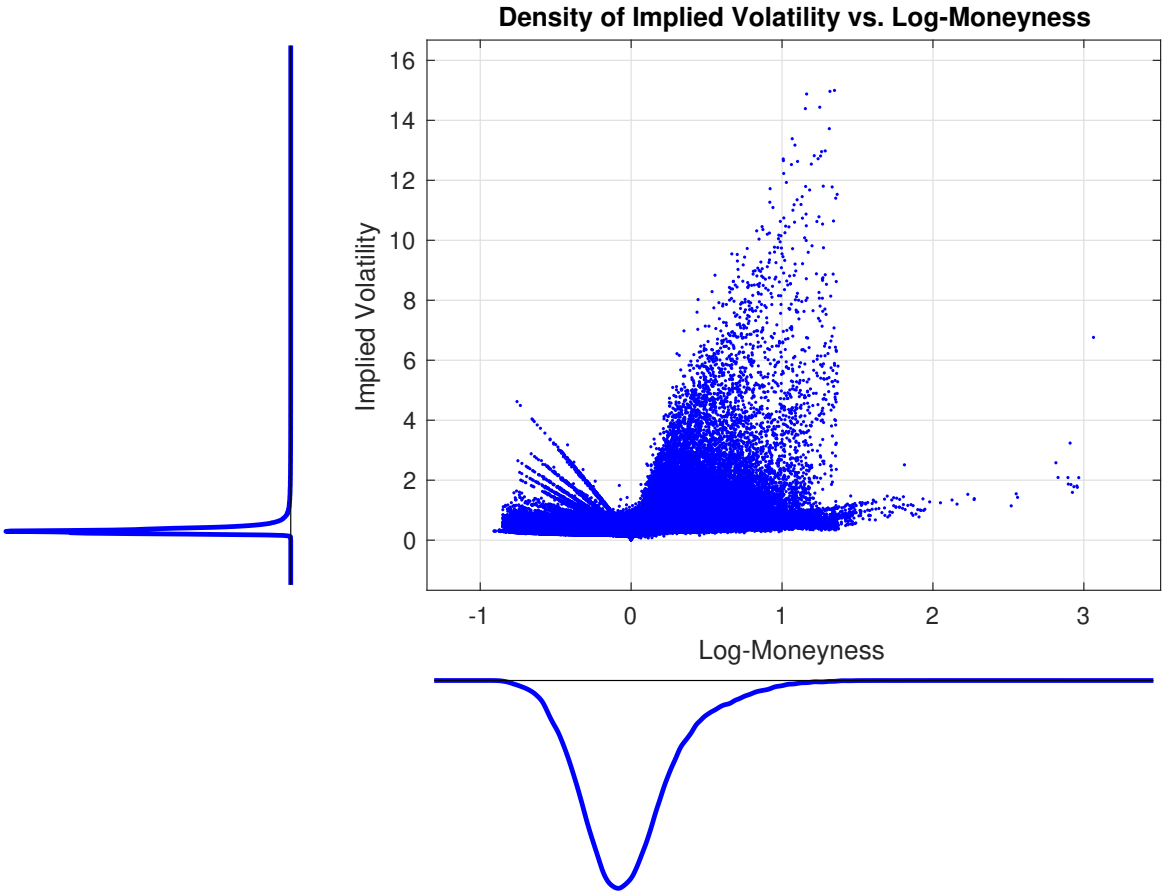


Figure 4.1: Density of implied volatility vs. log-moneyness.

4.2 Supplementary Data Inputs

In addition to options data, Apple stock data were collected from Macrotrends (2025), including daily high and low prices from 2018 to 2024. These supplementary data supported the volatility analysis and helped contextualise the behaviour of out-of-the-money options. Furthermore, quarterly dividend payments (typically issued in February, May, August, and November) were obtained from Apple Inc. (2025) and aggregated to compute the total annual dividend per share for each year. These an-

nualised values were then used to estimate the dividend yield, calculated as the ratio of the total dividend to the average stock price for the respective year. Incorporating accurate dividend yields is essential for risk-neutral modelling frameworks such as the Heston model, where dividends adjust the expected return of the underlying asset. Table 4.1 summarises Apple's annual dividends per share and the corresponding dividend yields from 2018 to 2024.

Table 4.1: Annual dividends per share and dividend yields for Apple Inc. over 2018–2024.

Year	*Annual Dividend per Share ($)*	*Average Stock Price ($)*	*Dividend Yield (%)*
2018	0.7050	39.43	1.79
2019	0.7600	73.08	1.04
2020	0.8075	92.84	0.87
2021	0.8650	138.19	0.63
2022	0.9100	152.61	0.60
2023	0.9500	171.09	0.56
2024	0.9900	206.54	0.48

4.3 Data Challenges and Market Microstructure

A critical characteristic of this dataset is its complexity. For each trading day, it is necessary to filter and match active contracts by date, strike price, and maturity while addressing data gaps and noisy observations. Such data management is vital to ensure the reliability of subsequent calibration and filtering procedures.

Another important aspect of the dataset is the pricing information available for each contract. Four prices are recorded for each option: the bid price, the ask price, the mid-price (the average of the bid and ask prices), and the last-traded price. The bid and ask prices represent the best available quotations from buyers and sellers, respectively. Moreover, the bid–ask spread is a key market indicator: narrower spreads imply higher liquidity and competitive trading environments, whereas wider spreads suggest uncertainty or diminished market interest. The mid price is a commonly used reference value for calibration. The last price reflects the most recent transaction, which may be outdated or noisy during periods of low trading activity. In addition to pricing data, each option contract includes a record of the daily trading volume, representing the number of contracts exchanged on a given day. This measure reflects the level of market activity and reliability, which can be used to identify and exclude illiquid contracts that may otherwise distort the calibration process.

Building on these empirical components, this research developed a calibration framework and investigated the temporal dynamics of the underlying volatility pro-

cess. A central challenge in calibrating option pricing models lies in reconciling theoretical constructs with complex market behaviours. Classical frameworks, such as Black–Scholes and Heston models, rely on idealised assumptions (frictionless markets, continuous trading, and perfect liquidity) that enable tractable solutions but often fail to represent how options are quoted and traded in practice. As noted by Sepp (2008), even widely accepted models must be extended with features like jumps in returns and volatility to capture observed pricing patterns, particularly for instruments that reflect market skews and discontinuities (e.g. VIX options). Moreover, theoretical constructs based on continuous-time dynamics, such as quadratic variation, tend to underestimate the realised variance in short-term contracts when confronted with discrete sampling (Sepp, 2012).

Calibration is further complicated by the fragmented nature of market data, where available quotes are often sparse and inconsistent, comprising bid, ask, and last prices across various strikes and maturities. As emphasised by Andreasen and Huge (2010), interpolating these discrete observations into an arbitrage-consistent implied volatility surface is non-trivial and requires sophisticated numerical methods. In alignment with the concerns raised by Bauer (2017), who highlighted the instability of model parameters in thin markets, this research adopted a pre-processing approach to enhance the robustness of the proposed calibration framework. Enforcing no-arbitrage conditions during preprocessing leads to reduced parameter variance and more stable model performance (Nelis, 2022). Similarly, Ruf and Wang (2020) highlight the importance of structured data preparation in financial modelling, particularly for option pricing.

Many existing approaches rely on exclusion rules or ad-hoc filtering, which risk discarding valuable market information. For instance, Ivanovas (2015) and Meier (2015) survey practical filtering criteria such as liquidity thresholds and moneyness bounds. In contrast, arbitrage-free smoothing of the volatility surface provides a method to preserve more data while enhancing calibration stability (Fengler, 2009; Sridi and Bilokon, 2023). An additional and often underappreciated complication in model calibration arises from microstructure noise, which is particularly pronounced in tick-level or daily market data. As outlined by Aït-Sahalia and Jacod (2014), observed market quotes inherently include microstructure noise, such as bid–ask bounce, execution latency, and quote asynchrony, rather than purely reflecting asset fundamentals. Despite this, many calibration methods in the literature proceed as if such noise were negligible, attempting to fit models precisely to market prices that are inherently unreliable. This disregard for data imperfections undermines the robustness of the resulting model parameters and calls for more nuanced approaches that account for the noisy nature of

financial data. To better account for the noise inherent in option market quotes, this research treated observed prices as noisy signals subject to microstructure effects such as bid–ask spreads, asynchronous trading, and low liquidity. Inspired by the principles of data assimilation discussed by Evensen et al. (2022), stochastic perturbations to the market prices were introduced during calibration. This probabilistic framing allowed the optimiser to account for uncertainty, improving numerical stability and reducing sensitivity to outliers and illiquid quotes.

This approach to noisy market data also informed the projection of the underlying asset level. Recognising that discrepancies between call and put prices complicate direct spot inference, this research considered call and put quotes to reflect a consistent underlying level, factoring in deviations from put–call parity. While Cui et al. (2017) explored this idea in synthetic contexts, this research implemented it using actual market data, leveraging the observed spreads to inform calibration decisions. The effect was twofold, providing stable parameter estimates and structural consistency across trading days. This noise-aware calibration approach also set the stage for downstream modelling. Sequential filtering techniques, such as the ensemble Kalman filter, rely on the accurate characterisation of observation uncertainty. By adopting this perspective from the outset, the proposed calibration strategy takes into account the empirical nature of market data. It effectively bridges the gap between theoretical models and practical implementation by design.

4.4 Data Structure and Challenges

Given the complexity and scale of modern financial markets, market data collection, filtering, and analysis play a crucial role in calibrating and validating financial models. Option prices are viewed as derivatives of the underlying asset price. Early studies, such as Manaster and Rendleman (1982), found that option prices seemed to move before the prices of the underlying stocks. Later, Stephan and Whaley (1990) reported the opposite, showing that stock prices moved first. The difference was explained by Chan et al. (1993), who showed that it was not a real economic effect but rather the result of market frictions, such as the minimum tick size and the fact that many options did not trade frequently. These studies suggest that information flows both ways between options and stocks, making it difficult to say that one market always leads the other.

Two key challenges in addressing this question are the volume and heterogeneity of option contracts. For example, Apple alone may have more than $2,000$ option con-

tracts active on a trading day. These contracts are disseminated across platforms that utilise custom data formatting. Each option is defined by key parameters, including the asset ticker, expiration date, contract type (call or put), and strike price. Furthermore, each contract has a unique historical record that evolves from its issuance until its expiration, typically on a Friday, though contracts may be initiated on any business day.

Option prices fluctuate due to several factors, including changes in the underlying asset price and shifts in supply and demand. Market participants with opposing outlooks (call option buyers anticipating price increases and put buyers expecting declines) create competing pressures. Additionally, extrinsic influences such as news events and media sentiment can affect price dynamics. Corporate actions such as stock splits or reverse splits can also alter contract terms. These events affect both the option price and its strike price, which in turn affect the contract's symbol. For instance, Apple's stock split on August 28, 2020, prompted such adjustments. These data issues were addressed during calibration through the enforcement of economic constraints, notably put–call parity, as detailed in Chapter 5.

The behaviour of an option over time can be examined from two temporal perspectives:

1. *Time-series evolution*: analysing how the option price (or its underlying S_t) changes over calendar time t, for a fixed contract.

2. *Cross-sectional maturity progression*: comparing how option prices behave for contracts of different maturities $T_1, T_2, \ldots, T_F$, and different strikes $K_1, \ldots, K_F$, beginning from a fixed reference point.

To illustrate the second perspective, we fix a valuation date (pricing day) and examine a panel of options with varying strikes and maturities. Each contract is characterised by its strike price K and time to maturity $\tau = T - t_0$, where t_0 is the pricing day and T the expiration date. The payoff is realised at maturity as a function of the terminal stock price S_T, which is unknown at the time of pricing.

This structure aligns with the dataset used in this thesis, which contains daily closing prices for options with a range of strikes and maturities, all quoted on the same trading day. Since continuous-time data is rarely available at high frequency across the full option surface, the analysis focuses on end-of-day observations. Moreover, capturing intraday quotes at every minute or hour would dramatically increase the data volume and computational burden, especially when pricing thousands of contracts

across multiple strikes and maturities. End-of-day prices therefore offer a practical balance between data availability and analytical tractability.

Chapter 5

Balanced Premium Calibration Method

5.1 Introduction

This chapter introduces the Balanced Premium Calibration Method (BPCM), a three-layer framework designed to address instability and noise in option model calibration. Following a detailed description of the methodology and implementation details, the chapter presents empirical evaluations of performance, error analysis, diagnostic metrics, and sensitivity structures. BPCM represents a fundamental shift in how calibration is conceptualised. Rather than a traditional one-step optimisation to fit market prices, BPCM is a multi-layer data assimilation and model refinement process. This approach addresses several key challenges (highlighted below) in calibrating the Heston stochastic volatility model to noisy market data.

Signal vs. Noise Separation. Classic calibration approaches attempt to fit the model directly to observed option prices, effectively treating all market quotes (however noisy or inconsistent) as exact targets. By contrast, BPCM assumes the Heston model provides a reasonable structural description of the dynamics and interprets most discrepancies as arising from observational noise. This perspective, inspired by strong-constraint data assimilation techniques (Evensen et al., 2022), helps disentangle genuine market signals from microstructure noise, arbitrage violations, and asynchronous quote errors. While some degree of model misspecification cannot be ruled out, by filtering out noise and focusing on the dominant pricing signal, BPCM achieves a more robust calibration.

Layered Refinement Strategy. BPCM employs a structured three-layer procedure to refine the data and the model parameters iteratively. The first layer cleanses and reconstructs market data for consistency, while the second layer fits model parameters to the cleansed data. Then, the third layer fine-tunes the model by distributing

residual errors into known noise sources. This progressive, layered refinement ensures that the primary calibration is driven by reliable information and is not distorted by peripheral market imperfections.

Enforcing Economic Coherence. BPCM explicitly rejects the notion that any single observed quote is a perfect reflection of the actual price. Instead, it constructs an internally consistent and economically plausible set of prices that respects fundamental no-arbitrage conditions such as put–call parity and bid–ask bounds. By maintaining these structural constraints throughout the calibration, the Heston model's outputs remain statistically well-fitted, economically meaningful, and arbitrage-free.

The proposed BPCM builds on and significantly extends prior calibration methods. One such example is the efficient calibration scheme presented by Cui et al. (2017), which leverages analytical gradients but encounters difficulties when applied to inconsistent data. These approaches, while valuable, generally treat data cleaning and model fitting as distinct or isolated steps. By contrast, BPCM tightly integrates data adjustment into the calibration process, continually guiding the model toward internally consistent solutions. While Cui et al. (2017) showed that applying put–call parity in controlled settings improves performance, BPCM applies such structural constraints dynamically on real, imperfect market data throughout the calibration. As a result, the calibrated model aligns with an economically coherent pricing surface rather than a scatter of noisy points.

Table 5.1, highlighted in recent surveys (e.g. Gupta (2025)), summarises how BPCM addresses common limitations in Heston calibration.

This chapter is organised as follows. Section 5.2 introduces the BPCM architecture and its three-layer calibration structure. Section 5.3 describes the practical implementation of the framework, including parameter initialisation, optimisation settings, and numerical procedures. Section 5.4 assesses calibration accuracy across stable and unstable market regimes. Section 5.5 examines the computational cost and scaling behaviour of the calibration framework. Section 5.6 reports empirical performance results and estimated model parameters, followed by Section 5.7, which discusses the treatment of interest rates and dividend yields. Section 5.8 analyses gradient diagnostics, and Section 5.9 examines model sensitivities (Greeks). Finally, Section 5.10 explores parameter interrelationships and their implications for calibration stability and interpretability.

Table 5.1: BPCM solutions to common Heston calibration limitations.

Limitation	*Addressed by BPCM?*	*BPCM Approach*
Assumption of independent and identically distributed Gaussian noise in option prices.	Yes	Explicitly **filters** quotes (parity + bid–ask) into a parity-consistent cleaned panel and calibrates to it, avoiding i.i.d.-noise assumptions on raw quotes.
Ambiguity between model error and data error.	Yes	Treats the Heston model as structurally correct and attributes inconsistencies to noisy data inputs, similar to strong-constraint filtering, thereby clarifying the source of errors.
High computational cost in advanced calibration (e.g. using ensembles or adjoint methods).	Partially	Reduces the parameter space by fixing the reconstructed spot and parity-consistent option prices, and uses analytical gradients for Heston parameters, achieving faster calibration without costly ensemble or adjoint PDE methods.
Model misspecification (e.g. volatility structure not rich enough).	No	BPCM assumes the Heston model is structurally valid and focuses on data quality; addressing true misspecification would require extending the model itself.
Calibration instability from illiquid or inconsistent quotes.	Yes	Uses reconstructed spot and parity-consistent option prices; bid/ask quotes act as bounds, yielding stable calibration even with illiquid or inconsistent data.

5.2 The BPCM Framework

BPCM incrementally improves data quality and model fit in three stages: Layer 1 filters noisy quotes, Layer 2 calibrates model parameters, and Layer 3 redistributes residual errors under bounded constraints through an adaptive refinement procedure.

Layer 1: Data Filtering via Put–Call Parity

The first BPCM layer focuses on cleansing the market data and extracting reliable pricing signals. It examines each option's quotes (bid and ask for calls and puts) along with its strike K, time to maturity τ, risk-free interest rate r, and dividend yield q. Rather than taking these quotes at face value, this layer seeks an implied spot price S and the corresponding call (C) and put (P) premiums that best satisfy put–call parity, i.e. the no-arbitrage relationship between European call and put prices on the

same strike and maturity (Su and Guo, 2006).

$$C - P = S\,e^{-q\tau} - K\,e^{-r\tau}. \tag{5.1}$$

Although call–put parity provides a convenient adjustment between call and put prices with the same strike and maturity, it imposes only one necessary condition for the absence of static arbitrage. A broader no-arbitrage characterisation follows from the martingale property of the discounted asset price under the risk–neutral measure, which implies additional requirements such as monotonicity and convexity in strike and the absence of calendar-spread arbitrage across maturities. In this thesis, parity adjustments are used solely to correct local inconsistencies between paired call–put quotes, and basic filters remove prices that violate trivial bounds or yield economically implausible spreads across neighbouring strikes and maturities. A full enforcement of static no-arbitrage conditions, including convex-optimisation approaches that project prices onto arbitrage-free surfaces, is beyond the scope of this study and may be considered in future work.

Notation for daily cross-section. Let $d = 1, \ldots, D$ index the trading days in the sample. On each day d, denote by N_d the number of available option contracts and let $\mathcal{I}_d = \{1, \ldots, N_d\}$ be their index set. Each contract $i \in \mathcal{I}_d$ is defined by a strike K_i and time to maturity τ_i, and is observed through bid and ask quotes for both calls and puts, $C_i^{\text{bid}}, C_i^{\text{ask}}$ and $P_i^{\text{bid}}, P_i^{\text{ask}}$. The optimisation for contract i adjusts these quotes within their bid–ask bounds and infers a consistent spot price S, by solving

$$\min_{\substack{S\in[S_l,S_h]\\ P\in[P_i^{\text{bid}},P_i^{\text{ask}}]\\ C\in[C_i^{\text{bid}},C_i^{\text{ask}}]}} \Big[\, C - P - \big(Se^{-q\tau_i} - K_i e^{-r\tau_i}\big) \,\Big]^2, \tag{5.2}$$

where $[S_l, S_h]$ is the admissible range for the spot. The solution yields the adjusted triplet $\big(C_i^{\text{adj}}, P_i^{\text{adj}}, S_i^{\text{adj}}\big)$ that is internally consistent with put–call parity.

Layer 1 Error. For each option pair i (consisting of call and put quotes at strike K_i and maturity τ_i), we define the *Layer 1 error* as the absolute deviation from put–call parity:

$$\mathbf{r}_i = \Big|\, C_i^{\text{adj}} - P_i^{\text{adj}} - \big(S_i^{\text{adj}} e^{-q\tau_i} - K_i e^{-r\tau_i}\big) \,\Big|. \tag{5.3}$$

By construction, both the call and the put in pair i share the same error value,

$$E_i^{\mathrm{L1,call}} = E_i^{\mathrm{L1,put}} = \mathbf{r}_i. \tag{5.4}$$

Contracts are classified as `stable` if $\mathbf{r}_i \leq \alpha$ and `unstable` otherwise, where $\alpha > 0$ is a noise tolerance threshold that determines how closely the observed option quotes must satisfy put–call parity; its determination is described in Section 5.3. This classification divides the dataset into two subsets: stable contracts, whose adjusted quotes are consistent with parity within tolerance, and unstable contracts, which likely reflect data errors, extreme illiquidity, or temporary arbitrage. Both subsets are retained, and the labels provide insight into data quality as well as a benchmark for calibration performance.

Outputs. Layer 1 therefore produces (i) an implied spot price S^{adj} for each day, (ii) adjusted option premiums $(C^{\mathrm{adj}}, P^{\mathrm{adj}})$ for each contract, and (iii) a stability label. These outputs ensure that subsequent layers operate on data filtered for basic arbitrage consistency.

The above procedure can be summarised as Algorithm 5.1, which makes explicit the optimisation and classification steps.

Algorithm 5.1 Layer 1: Put–Call Parity Filtering

Input: For day d, bid/ask quotes $\{(C_i^{\mathrm{ask}}, C_i^{\mathrm{bid}}, P_i^{\mathrm{ask}}, P_i^{\mathrm{bid}}, K_i, \tau_i, r_i, q)\}_{i=1}^{N_d}$
Output: Stored adjusted contracts $\{(C_i^{\mathrm{adj}}, P_i^{\mathrm{adj}}, S_i^{\mathrm{adj}}, \mathrm{label}_i)\}_{i=1}^{N_d}$

for each contract $i = 1, \ldots, N_d$ **do**
  Solve optimisation (5.2) for $(C_i^{\mathrm{adj}}, P_i^{\mathrm{adj}}, S_i^{\mathrm{adj}})$
  Compute Layer 1 error $\mathbf{r}_i$ using (5.3)
  if $\mathbf{r}_i \leq \alpha$ **then**
    Label contract as `stable`
  else
    Label contract as `unstable`
  end if
  Store $(C_i^{\mathrm{adj}}, P_i^{\mathrm{adj}}, S_i^{\mathrm{adj}}, \mathrm{label}_i)$
end for
return $\{(C_i^{\mathrm{adj}}, P_i^{\mathrm{adj}}, S_i^{\mathrm{adj}}, \mathrm{label}_i)\}_{i=1}^{N_d}$

Layer 2: Daily Parameter Calibration

Layer 2 calibrates the Heston model parameters to the parity-adjusted option data produced in Layer 1. The essential feature of this layer is that calibration is performed

separately for the `stable` and `unstable` subsets identified in Layer 1, so that we can assess how data quality affects the estimated parameters. Within each trading day, the calibration proceeds in two steps: first on the stable contracts, then on the unstable contracts.

Notation for daily cross-section. Let $d = 1, \ldots, D$ index the trading days, and let $\mathcal{I}_d = \{1, \ldots, N_d\}$ denote the contracts available on day d. From Layer 1, each contract $i \in \mathcal{I}_d$ provides

$$\mathbf{y}_i^{\text{adj}} = \big(C_i^{\text{adj}},\, P_i^{\text{adj}},\, \text{label}_i\big),$$

where $\text{label}_i \in \{\texttt{stable}, \texttt{unstable}\}$ indicates the subset membership. We also define the per-contract input vector

$$\mathbf{u}_i = \big(S_i^{\text{adj}},\, K_i,\, \tau_i,\, r_i,\, q_i\big),$$

consisting of the Layer 1. Thus, the calibration dataset for day d is

$$\mathcal{U}_d = \big\{(\mathbf{y}_i^{\text{adj}},\, \mathbf{u}_i)\big\}_{i=1}^{N_d}.$$

Calibration objective. Let $\boldsymbol{\theta} = (v_0, \bar{v}, \kappa, \sigma, \rho)$ denote the Heston parameter vector, and let $\mathbf{u}_i$ collect the contract-specific inputs. For contract i, denote by $C_i^{\text{mod}}(\boldsymbol{\theta}, \mathbf{u}_i)$ and $P_i^{\text{mod}}(\boldsymbol{\theta}, \mathbf{u}_i)$ the model-implied call and put prices, respectively. The calibration problem is defined in minimising the following least-squares objective

$$\min_{\boldsymbol{\theta}} \sum_{i=1}^{N_d} \Big(C_i^{\text{mod}}(\boldsymbol{\theta}, \mathbf{u}_i) - C_i^{\text{adj}}\Big)^2 + \Big(P_i^{\text{mod}}(\boldsymbol{\theta}, \mathbf{u}_i) - P_i^{\text{adj}}\Big)^2. \tag{5.5}$$

This calibration follows the approach of Cui et al. (2017), in which the implied spot S_i^{adj} obtained from Layer 1 is fixed, and parameters are estimated against the adjusted arbitrage-consistent quotes.

Layer 2 errors. Given parameters $\boldsymbol{\theta}$, the signed residuals for contract i are

$$r_i^{\text{call}}(\boldsymbol{\theta}) = C_i^{\text{mod}}(\boldsymbol{\theta}, \mathbf{u}_i) - C_i^{\text{adj}}, \tag{5.6}$$

$$r_i^{\text{put}}(\boldsymbol{\theta}) = P_i^{\text{mod}}(\boldsymbol{\theta}, \mathbf{u}_i) - P_i^{\text{adj}}. \tag{5.7}$$

Their absolute values are the *Layer 2 errors*,

$$E_i^{\mathrm{L2,call}} = \left|r_i^{\mathrm{call}}(\boldsymbol{\theta})\right|, \qquad E_i^{\mathrm{L2,put}} = \left|r_i^{\mathrm{put}}(\boldsymbol{\theta})\right|. \tag{5.8}$$

The optimisation (5.5) thus minimises the aggregate squared Layer 2 errors.

Sequential initialisation. On each day d, calibration proceeds in two steps. For $d = 1$, the procedure is initialised with a generic starting guess $\boldsymbol{\theta}_0$ (the first-day prior). For $d > 1$, the stable subset is initialised with the previous day's stable parameters $\boldsymbol{\theta}_{d-1}^{\mathrm{stable}}$. After calibrating the stable subset, its output $\boldsymbol{\theta}_d^{\mathrm{stable}}$ is then used as the starting guess for the unstable subset, yielding $\boldsymbol{\theta}_d^{\mathrm{unstable}}$. Finally, $\boldsymbol{\theta}_d^{\mathrm{stable}}$ is carried forward to initialise day $d + 1$. This sequential structure enforces temporal smoothness in the estimates and stabilises the daily calibration procedure.

Outputs. For each day d, Layer 2 produces two calibrated parameter vectors,

$$\boldsymbol{\theta}_d^{\mathrm{stable}}, \qquad \boldsymbol{\theta}_d^{\mathrm{unstable}}.$$

Together, these form two parallel time series of parameter estimates. The stable sequence reflects calibration against the liquid, well-behaved subset of quotes, providing a benchmark for economic plausibility. The unstable sequence, by contrast, captures how the calibration responds when noisier or less reliable data are included, thereby highlighting the sensitivity of inference to data quality and market frictions. Maintaining both series is therefore informative: the stable path anchors the analysis, while the unstable path offers a diagnostic perspective on robustness across market conditions. In this way, Layer 2 delivers not only parameter estimates for subsequent refinement but also a broader view of calibration reliability. The calibration steps are formalised in Algorithm 5.2.

Algorithm 5.2 Layer 2: Daily Calibration of Heston Parameters (Single Day)

Input: For day d, Adjusted dataset $\{(\mathbf{y}_i^{\mathrm{adj}}, \mathbf{u}_i)\}_{i=1}^{N_d}$ from Layer 1, partitioned into `stable` and `unstable` subsets; initial guess $\boldsymbol{\theta}_0$ (taken as $\boldsymbol{\theta}_{d-1}^{\mathrm{stable}}$ for $d > 1$)
Output: Calibrated parameters $\boldsymbol{\theta}^{\mathrm{stable}}, \boldsymbol{\theta}^{\mathrm{unstable}}$

1: Solve optimisation (5.5) on the `stable` subset initialised at $\boldsymbol{\theta}_0$
2: Store result as $\boldsymbol{\theta}^{\mathrm{stable}}$
3: Solve optimisation (5.5) on the `unstable` subset initialised at $\boldsymbol{\theta}^{\mathrm{stable}}$
4: Store result as $\boldsymbol{\theta}^{\mathrm{unstable}}$
5: **return** $\boldsymbol{\theta}_d^{\mathrm{stable}}$, $\boldsymbol{\theta}_d^{\mathrm{unstable}}$

Layer 3: Refinement via Joint Error Redistribution

Even after Layer 2 calibration, residual discrepancies remain between model-implied prices and the adjusted market quotes. These misfits may arise from microstructure noise (e.g. bid–ask bounce, discrete ticks; see Aït-Sahalia and Lo, 1998; Aït-Sahalia et al., 2005) or structural limitations of the Heston model. Layer 3 introduces a refinement step that treats both the model and the adjusted market inputs as uncertain and redistributes the residual errors in a disciplined manner. The goal is not to enforce an exact match to noisy quotes, but to strike a balance: minor, bounded adjustments to both data and parameters yield an economically interpretable and numerically stable fit.

Notation. We retain the notation of Layers 1–2 for trading days; thus, Layer 3 extends the optimisation space by treating both data and parameters as adjustable within predefined bounds. Collecting all optimisation variables into a joint vector

$$\mathbf{z} = \big(\boldsymbol{\theta}, \mathbf{u}, \mathbf{y}^{\mathrm{adj}}\big),$$

Objective. For a single call–put pair $(C^{\mathrm{adj}}, P^{\mathrm{adj}})$ with maturity τ, the Layer 3 calibration minimises the joint squared error between model and adjusted prices:

$$\min_{\mathbf{z}} \quad \Big(C^{\mathrm{mod}}(\boldsymbol{\theta}, \mathbf{u}) - C^{\mathrm{adj}}\Big)^2 + \Big(P^{\mathrm{mod}}(\boldsymbol{\theta}, \mathbf{u}) - P^{\mathrm{adj}}\Big)^2, \tag{5.9}$$

subject to dynamic bounds on $\mathbf{z}$.

Layer 3 errors. As in the previous layers, we define pricing residuals as the difference between model-implied premiums and their adjusted counterparts. Given parameters $\boldsymbol{\theta}$ and inputs $\mathbf{u}$ (both now treated as variables within $\mathbf{z}$), the call and put residuals are

$$r^{\mathrm{call}} = C^{\mathrm{mod}}(\boldsymbol{\theta}, \mathbf{u}) - C^{\mathrm{adj}}, \qquad r^{\mathrm{put}} = P^{\mathrm{mod}}(\boldsymbol{\theta}, \mathbf{u}) - P^{\mathrm{adj}}. \tag{5.10}$$

Their absolute values are the *Layer 2 errors*,

$$E_i^{\mathrm{L3,call}} = \big|r_i^{\mathrm{call}}(\boldsymbol{\theta})\big|, \qquad E_i^{\mathrm{L3,put}} = \big|r_i^{\mathrm{put}}(\boldsymbol{\theta})\big|. \tag{5.11}$$

Unlike in Layer 2, the residuals in Layer 3 are computed with respect to premiums, while allowing the inputs themselves to act as optimisation variables. This means that

misfits are no longer attributed solely to model parameters: instead, the procedure enables minor, controlled adjustments in both the model (via $\boldsymbol{\theta}$) and the market inputs (spot, rates, and quotes). In this way, residual errors are redistributed jointly across all sources of uncertainty, rather than being forced entirely onto the parameter estimates.

Bound refinement. Each optimisation variable $z \in \mathbf{z}$ is constrained by lower and upper bounds $(l_z^{(k)}, u_z^{(k)})$ that evolve across refinement iterations k. The initial values, i.e. for $k = 1$, of the lower and upper bounds for each variable are given by

$$l_z^{(1)} = z - \delta_z^{(1)}, \qquad u_z^{(1)} = z + \delta_z^{(1)}. \tag{5.12}$$

The quantity $\delta_z^{(1)}$ represents the width of the initial bound around each optimisation variable z. It is best understood as a proxy for the *noise* present in the input data carried forward from Layer 2. Option datasets are known to contain several sources of noise — bid–ask bounce, recording delays, and inconsistencies across calls and puts — which distort the underlying signals used for calibration (see e.g. Jackwerth, 2000; Cont, 2001). Accordingly, $\delta_z^{(1)}$ is set as a conservative margin that reflects the typical scale of such market noise: for directly observed quantities (e.g. adjusted premiums or implied spots), it is anchored to bid–ask or daily high–low ranges; for latent model parameters, where no direct market interval exists, it is taken as a small relative perturbation of the previous estimate. In both cases, the intent is not to predict the exact magnitude of noise, but to provide a stable trust region within which the optimisation can explore plausible adjustments without overfitting spurious fluctuations. After solving (5.9), the maximum absolute error across call and put is computed as

$$R_{\max} = \max\left(|r^{\text{call}}|,\ |r^{\text{put}}|\right), \tag{5.13}$$

that measures the pricing error between model outputs and adjusted market quotes. Let $\alpha_{\text{res}} > 0$ denote the residual tolerance, interpreted as the maximum acceptable error. In practice, α_{res} acts as a tuning parameter: a smaller value enforces very tight fits to the data (risking overfitting to noise), while a larger value allows more flexibility and tolerates noisier quotes. This role is closely related to regularisation in other calibration frameworks (e.g. (Coleman et al., 1999; Fengler, 2009)). Whenever $R_{\max}$ exceeds α_{res}, the bounds for each optimisation variable z are relaxed by a fixed increment δ_z:

$$l_z^{(k+1)} = l_z^{(k)} - \delta_z, \qquad u_z^{(k+1)} = u_z^{(k)} + \delta_z, \tag{5.14}$$

and the optimisation is repeated until either $R_{\max} \leq \alpha_{\text{res}}$ or the maximum iteration limit $K_{\max}$ is reached, an approach akin to trust-region strategies in numerical optimisation (Conn et al., 2000; Lagnado and Osher, 1997).

Termination. The refinement procedure begins with conservative, market consistent bounds (5.12), and only expands them if residual errors remain above tolerance.In this way, Layer 3 enforces a disciplined redistribution of errors: the model is not forced to explain away all noise, but neither are data adjustments unconstrained.To formalise this process, Algorithm 5.3 specifies the iterative refinement scheme, showing how residuals are evaluated and bounds progressively expanded until a satisfactory solution is obtained or iteration limits are reached.

Algorithm 5.3 Layer 3: Joint Refinement Procedure

Input: Initial guess $\mathbf{z} = (\boldsymbol{\theta}, \mathbf{u}, \mathbf{y}^{\text{adj}})$; tolerance α_{res}; maximum iterations $K_{\max}$
Output: Refined optimisation vector $\mathbf{z}^*$

Initialise bounds for each $z \in \mathbf{z}$ using rule (5.12)
repeat
    Solve optimisation (5.9) subject to current bounds
    Compute residual $R_{\max}$ using (5.13)
    if $R_{\max} > \alpha_{\text{res}}$ **and** $k < K_{\max}$ **then**
        Expand bounds via rule (5.14); increment k
    else
        Terminate
    end if
until $R_{\max} \leq \alpha_{\text{res}}$ or $k = K_{\max}$
return $\mathbf{z}^*$

Summary. Layer 3 complements Layers 1–2 by addressing residual errors that persist after calibration. It implements a joint optimisation over parameters, market inputs, and adjusted quotes, all constrained within economically meaningful ranges. The staged refinement design ensures stability: bounds begin narrow and only expand when residuals remain high. As a result, the final outputs $\mathbf{z}^*$ merge Layer 2 parameter estimates with minimally adjusted data, producing a calibration that is simultaneously robust to market noise and consistent with observed option information.

5.3 Implementation Details

We now detail the practical implementation of the three-layer BPCM framework. Although the methodology of Layers 1, 2, and 3 was presented abstractly, deploying

the framework on real market data requires specific parameter choices, data preprocessing steps, and stopping criteria. We outline these implementation details for each layer below, emphasising how data is processed and how the calibration algorithms are executed to ensure robustness and economic consistency.

Implementation of Layer 1

Layer 1 is implemented as a constrained least-squares adjustment applied *pairwise* to matching call and put contracts, consistent with the theoretical setup in (5.1)–(5.3) and Algorithm 5.1. The preprocessing begins with the removal of duplicate records, defined as contracts with identical dates, strikes, expirations, and option types. When last-traded prices are missing or recorded as zero, they are replaced with the mid-quote average of bid and ask to avoid distortions. Finally, for each trading day, only those strike–maturity combinations for which both call and put quotes exist are retained, while solitary quotes are discarded.

Timestamp alignment. The raw dataset records only calendar dates (with timestamps fixed at 00:00), so maturities cannot be measured directly at intraday resolution. To reconstruct realistic timing, we anchor each contract to two session-specific points:

$$t_{\text{quote}} = \text{QuoteDate} + 9.4 \text{ hours (09:24 ET)},$$

$$t_{\text{exp}} = \text{ExpiryDate} + 16.1 \text{ hours (16:06 ET)}.$$

The time to maturity is then the ACT/365 year fraction between these anchors,

$$\tau \;=\; \frac{t_{\text{exp}} - t_{\text{quote}}}{365 \text{ days}},$$

which naturally includes hours and minutes, not just whole days. Anchoring quotes just before the U.S. cash open (09:30 ET) and expiries just after the cash close (16:00 ET) ensures that maturities are consistent with trading hours, rather than being biased by midnight-based timestamps.[1]

Parity fit per pair. For each call–put pair, we solve the optimisation problem (5.2) using a Levenberg–Marquardt least-squares algorithm. The variables (C, P, S)

[1]The slight offsets (09:24 and 16:06) follow implementation practice: the quote time is set a few minutes before the open to capture the session start uniformly, and the expiry time a few minutes after the close to ensure contracts are fully expired.

are initialised from observed last trades and the recorded underlying price, and are constrained by

$$C \in [C^{\text{bid}}, C^{\text{ask}}], \quad P \in [P^{\text{bid}}, P^{\text{ask}}], \quad S \in [S_l, S_h].$$

The spot bounds $[S_l, S_h]$ are defined by the daily low and high. A strict convergence tolerance (10^{-8}) ensures accurate enforcement of (5.1). In stressed market conditions, however, bid and ask quotes in the raw data occasionally appear inverted or inconsistent. To guard against such anomalies, the feasible intervals are always taken as the minimum and maximum of the recorded quotes, ensuring that optimisation remains well-posed. Contracts affected by such inconsistencies are subsequently classified as unstable.

Adjusted outputs and error. The solution $\left(C^{\text{adj}}, P^{\text{adj}}, S^{\text{adj}}\right)$ yields parity consistent premiums and an implied spot. The Layer 1 error E^{L1} is then computed using (5.3). These values are stored per contract, matching the outputs of Algorithm 5.1. Contracts are classified as `stable` if $E^{\text{L1}} \leq \alpha$ and `unstable` otherwise. Additionally, contracts are discarded if bid–ask quotes are inverted ($C^{\text{bid}} > C^{\text{ask}}$ or $P^{\text{bid}} > P^{\text{ask}}$). We set $\alpha = 1$ (currency units), corresponding to the average bid–ask spread (~ 1.08) in the dataset, ensuring that only genuine mispricings are filtered. All retained contracts are enriched with adjusted values $\left(C^{\text{adj}}, P^{\text{adj}}, S^{\text{adj}}\right)$ and the Layer 1 error E^{L1}, stored as additional fields.

Implementation of Layer 2

Layer 2 calibration is implemented as a nonlinear least-squares optimisation, consistent with the theoretical formulation (5.5)–(5.8) and the framework of Algorithm 5.2.

Residual formulation. For each trading day d, the adjusted dataset $\mathcal{U}_d$ from Layer 1 provides, per contract, the parity-consistent premium $y_i^{\text{adj}} \in \{C_i^{\text{adj}}, P_i^{\text{adj}}\}$ and inputs $(S_i^{\text{adj}}, K_i, \tau_i, r_i, q_i)$. To avoid ill-conditioning, option prices below 10^{-3} are floored at that level. All contract residuals (5.6)–(5.7) are stacked into a single vector $\mathbf{r}(\boldsymbol{\theta})$, reducing the calibration to the least-squares problem

$$\min_{\boldsymbol{\theta}} \ \|\mathbf{r}(\boldsymbol{\theta})\|_2^2.$$

Jacobian and LM step. At each iteration, the Jacobian $\mathbf{J}(\boldsymbol{\theta}) = \partial \mathbf{r} / \partial \boldsymbol{\theta}$ is assembled from analytic option-price derivatives, following the formulation of Cui et al. (2017).

The Levenberg–Marquardt (LM) update is then obtained by solving

$$\Delta\boldsymbol{\theta} \;=\; \big(\mathbf{J}^\top\mathbf{J} + \mu\boldsymbol{I}\big)^{-1}\mathbf{J}^\top\mathbf{r}, \tag{5.15}$$

where $\mu > 0$ is the damping parameter and $\boldsymbol{I}$ denotes the identity matrix of appropriate dimension.

Algorithm 5.4 Levenberg–Marquardt solver with analytic Jacobian

Input: Residual function $\mathbf{r}(\boldsymbol{\theta})$, Jacobian $\mathbf{J}(\boldsymbol{\theta})$, initial guess $\boldsymbol{\theta}_0$, damping μ_0, tolerances $(\epsilon_{\text{res}}, \epsilon_{\text{grad}}, \epsilon_\theta)$, maximum iterations $K_{\max}$
Output: Calibrated parameters $\boldsymbol{\theta}^*$

Initialise $\boldsymbol{\theta} \leftarrow \Pi(\boldsymbol{\theta}_0)$, compute $\mathbf{r}(\boldsymbol{\theta})$, set $f \leftarrow \|\mathbf{r}(\boldsymbol{\theta})\|^2$, and $\mu \leftarrow \mu_0$
for $k = 1$ to $K_{\max}$ **do**
    Assemble $\mathbf{J}(\boldsymbol{\theta})$
    Solve for $\Delta\boldsymbol{\theta}$ using (5.15)
    $\boldsymbol{\theta}_{\text{new}} \leftarrow \Pi(\boldsymbol{\theta} + \Delta\boldsymbol{\theta})$
    Compute $\mathbf{r}_{\text{new}} = \mathbf{r}(\boldsymbol{\theta}_{\text{new}})$ and $f_{\text{new}} = \|\mathbf{r}_{\text{new}}\|^2$
    Compute $\rho \approx (f - f_{\text{new}}) / \big(\Delta\boldsymbol{\theta}^\top(\mu\Delta\boldsymbol{\theta} + \mathbf{J}^\top\mathbf{r}(\boldsymbol{\theta}))\big)$
    if $f_{\text{new}} < f$ **and** $\rho > 0$ **then**
        Accept step: $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_{\text{new}}$, $f \leftarrow f_{\text{new}}$
        Update $\mu \leftarrow \mu \cdot \max\big(1/3,\, 1 - (2\rho - 1)^3\big)$
    else
        Reject step: increase $\mu \leftarrow 2\mu$
    end if
    if any condition in (5.16) holds **then**
        break ▷ convergence achieved
    end if
end for
return $\boldsymbol{\theta}^* = \boldsymbol{\theta}$

Projection and constraints. Each trial update $\boldsymbol{\theta}^{\text{new}} = \boldsymbol{\theta} + \Delta\boldsymbol{\theta}$ is mapped back to the admissible domain by the projection operator $\Pi(\cdot)$. This enforces positivity of $(v_0, \bar{v}, \kappa, \sigma)$, bounds $|\rho| < 1$, and applies a safeguard on the Feller index.

Feller regularisation. The Feller index is defined as

$$\mathcal{F}(\boldsymbol{\theta}) \;=\; \frac{2\kappa\bar{v}}{\sigma^2}.$$

Instead of enforcing the hard constraint $\mathcal{F} \geq 1$, we adopt a *soft regularisation* scheme. If $\mathcal{F} < \mathcal{F}_{\min}$, the parameters are nudged as

$$\sigma \leftarrow 0.9\sigma, \quad \bar{v} \leftarrow 1.1\bar{v}, \quad \kappa \leftarrow 1.1\kappa,$$

while if $\mathcal{F} > \mathcal{F}_{\max}$, the adjustments are

$$\sigma \leftarrow 1.1\sigma, \quad \bar{v} \leftarrow 0.9\bar{v}, \quad \kappa \leftarrow 0.9\kappa.$$

These updates are repeated until the index lies within the admissible range $\mathcal{F}_{\min} \leq \mathcal{F} \leq \mathcal{F}_{\max}$, which stabilises the variance process without artificially constraining the optimisation.

Adaptive damping. The damping parameter is updated according to the agreement ratio ρ between the actual and predicted residual decrease. Accepted steps reduce μ (Gauss–Newton behaviour), while rejected steps increase μ (gradient-descent behaviour). If $\mathbf{J}$ or $\mathbf{J}^\top\mathbf{J}$ is ill-conditioned, the parameter vector is perturbed and reprojected before recomputation.

Convergence. Iterations terminate once any of the following stopping conditions is satisfied:

$$\|\mathbf{r}(\boldsymbol{\theta})\|_2 \leq \varepsilon_{\text{res}}, \qquad \|\mathbf{J}^\top\mathbf{r}(\boldsymbol{\theta})\|_\infty \leq \varepsilon_{\text{grad}}, \qquad \frac{\|\Delta\boldsymbol{\theta}\|_2}{\|\boldsymbol{\theta}\|_2} \leq \varepsilon_\theta. \tag{5.16}$$

These criteria correspond respectively to: (i) a small residual norm, (ii) a slight projected gradient, or (iii) a minor relative parameter update. The overall procedure is summarised in Algorithm 5.4. In our implementation, we set $\mathcal{F}_{\min} = 0.001$ and $\mathcal{F}_{\max} = 15$. For the stopping tolerances we use $\varepsilon_{\text{res}} = \varepsilon_{\text{grad}} = \varepsilon_\theta = 10^{-12}$, with a maximum of 100 iterations.

Implementation of Layer 3

The Layer 3 refinement is implemented as a constrained nonlinear least-squares procedure, extending Layer 2 calibration by allowing bounded adjustments to both parameters and market inputs, in line with the framework of Eqs. (5.9)–(5.14) and Algorithm 5.3. The key practical elements are summarised below.

Joint optimisation setup. For each trading day, inputs consist of (i) the Layer 2 parameter vector for both stable and unstable and (ii) the complete set of parity-adjusted option data from Layer 1. Options are grouped in call–put pairs with common (K, τ). For each such pair, the optimisation variables are collected in a joint block

$$\mathbf{z} = (\boldsymbol{\theta},\, S^{\text{adj}},\, r,\, q,\, \tau,\, C^{\text{adj}},\, P^{\text{adj}}).$$

Solver and tolerances. The optimisation is solved using the Levenberg–Marquardt algorithm, a damped least-squares method that combines the Gauss–Newton scheme with gradient descent Levenberg (1944); Marquardt (1963); Moré (1978). This approach employs tight tolerances (10^{-8}) and 100 iterations per refinement step. Analytical Jacobians are supplied for Heston prices, while numerical differences are used for adjusted inputs.

Initial bounds. Bounds are defined according to (5.12):

- **Spot and option premiums ($S^{\text{adj}}, C^{\text{adj}}, P^{\text{adj}}$):** Lower and upper bounds are anchored to observed daily ranges: for S, the daily low and high of the underlying; for C and P, the bid and ask quotes. Layer 1 adjusted values outside the raw range by tolerance α are included by extending the interval.

- **Heston parameters:** Each is restricted to a narrow window around the Layer 2 solution. Typical bounds are $\pm 4\%$ for $(v_0, \bar{v}, \kappa, \sigma)$, while ρ is clipped to $[-0.99, 0.99]$.

- **Rates and dividend yields (r, q):** Allowed to vary only within a few basis points on the first iteration (to correct interpolation or rounding noise), then held fixed.

- **Time to maturity τ.** The time to maturity τ, measured in years . For contracts expiring on the same trading day, we introduce a small adjustment ϵ to capture the remaining fraction of the trading session, with $0 \leq \epsilon \leq 7.5 \times 10^{-4}$ years (≈ 6.6 hours). In the dataset, all intraday maturities are aligned to the day's closing time, so the longest possible adjustment corresponds precisely to the length of a regular U.S. trading session.

Refinement loop. If the maximum residual $R_{\max}$ (5.13) exceeds the tolerance $\alpha_{\text{res}} = 0.1$ (10 cents, well below the mean bid–ask spread of \$1.07 (stable) and \$1.28 (unstable) in our dataset), bounds are expanded according to (5.14):

- $S^{\text{adj}}, C^{\text{adj}}, P^{\text{adj}}$: widened by $\delta = 0.1$ in price units ($\approx$ one tick, i.e. the minimum price increment).
- Heston parameters: windows widened modestly, e.g. to $\pm 5\%$ for $(v_0, \bar{v}, \kappa, \sigma)$; ρ kept within $\pm 1\%$ of its previous value.
- Exogenous inputs (τ, r, q): only adjusted in the first iteration, then a fixed boundary for $k \geq 2$.

At each iteration, parameter updates are projected back into the admissible set (enforcing positivity and $|\rho| < 1$). For the Feller condition, we do not enforce it strictly, but apply a soft correction: if the ratio

$$\frac{2\kappa\bar{v}}{\sigma^2} < 0.1,$$

then σ is rescaled as $\sigma \leftarrow \sqrt{0.95 \cdot 2\kappa\bar{v}}$. This prevents the variance process from collapsing while still allowing the optimisation to explore regions near the boundary of the admissible domain. Hence, the Feller constraint is treated not as a hard feasibility condition but as a stabilisation device to avoid degenerate variance paths. The loop continues until $R_{\max} \leq \alpha_{\text{res}}$ or the iteration limit $K_{\max}$ is reached ($K_{\max} = 5$ for stable, $K_{\max} = 20$ for unstable).

Outputs and conclusion. Upon termination, the refined solution $\mathbf{z}^*$ delivers final Heston parameters $\boldsymbol{\theta}^*$, adjusted spot S^*, option premiums C^*, P^*, and the effective maturity τ^* (computed from the optimised quotation time t^*_{quote} to the fixed expiration date). Minor corrections to r^*, q^* are also recorded. Residuals are confirmed to lie well within bid–ask spreads, and diagnostic outputs include the final Feller index $2\kappa^*\bar{v}^*/{\sigma^*}^2$, the number of refinement steps, and stored per-contract residuals. These define the day's `FitWell` (the output of the third BPCM layer) set and a consistent adjusted dataset.

The empirical dataset comprises 1,523,842 Apple Inc. option quotes (2018–2024), of which 1,431,656 ($\approx 94\%$) are classified as stable and 92,186 ($\approx 6\%$) as unstable, exhibiting parity breaks or bid–ask anomalies. Layer 3 ensures that both subsets can be calibrated coherently: the stable subset yields smooth parameter paths with minimal residuals, while the unstable subset still admits economically plausible solutions without implausible jumps.

This concludes the implementation section. The following sections will evaluate performance in detail, focusing on the time-series behaviour of parameters, compliance

with structural conditions, and residual magnitudes relative to bid–ask spreads.

5.4 Calibration Accuracy and Robustness

BPCM's calibration accuracy and stability were evaluated using multiple error metrics across both (*stable*) and (*unstable*) market regimes. Each calibration layer produces its own error measure: Layer 1 addresses put–call parity (yielding a parity violation error for each call–put pair), Layer 2 fits the Heston model (yielding model fitting residuals), and Layer 3 performs a joint refinement (yielding final residual errors). We assess how well the calibrated model prices align with market quotes in each regime, using both static error distributions and temporal analysis.

Error metrics by layer

Calibration accuracy is evaluated using standard error statistics. For each BPCM layer, we compute the mean squared error (MSE: average squared pricing deviation), the variance of absolute errors (dispersion in the magnitude of mispricings), and the root mean squared error (RMSE: square root of MSE, in price units). The results for stable and unstable datasets are reported in Table 5.2. For the **stable** subset, pricing

Table 5.2: Calibration errors (MSE, variance of absolute errors, RMSE) for stable and unstable datasets across BPCM layers and the traditional calibration.

Calibration Stage	MSE		Var. of abs. error		RMSE	
	Stable	Unstable	Stable	Unstable	Stable	Unstable
Layer 1 (Parity fit)	0.0123	23.2376	0.0118	7.9095	0.1108	4.8205
Layer 2 (Heston fit)	0.0890	6.1130	0.0559	2.2135	0.2984	2.4724
Traditional	0.4198	19.7951	0.2640	8.3732	0.6479	4.4492
Layer 3 (Final fit)	0.0011	1.4436	0.0010	0.9576	0.0328	1.2015

errors are very small. Layer 1 (parity adjustment) yields an MSE of ≈ 0.0123 (RMSE ≈ 0.11), confirming that the raw quotes were already close to arbitrage-consistent. Layer 2 (Heston fit) increases the RMSE modestly to ≈ 0.30, reflecting the model's finite ability to fit all prices exactly. Layer 3 (joint refinement) then reduces the error substantially, reaching an RMSE of only ≈ 0.033, with absolute-error variance close to zero. This progression shows that BPCM achieves near-perfect calibration under normal market conditions.

For the **unstable** subset, initial errors are much larger. Layer 1 begins with severe parity violations (MSE ≈ 23.2, RMSE ≈ 4.82). After enforcing parity, the Layer 2

Heston fit reduces the RMSE by roughly half, to ≈ 2.47. Layer 3 further lowers it to ≈ 1.20, a substantial improvement but still an order of magnitude larger than the stable case. These higher residuals are calculated because BPCM leaves errors that are not zero, instead of overfitting unexpected quotes or breaking no-arbitrage constraints, which maintains stability.

For reference, we also report results from a **traditional calibration**, i.e. a direct Heston fit to observed option prices at market mid-quotes and underlying levels, using exactly the same fitting procedure as in Layer 2. To ensure comparability, this calibration was initialised with the Layer 2 parameter estimates (stable output for the stable subset and unstable output for the unstable set). Hence, the traditional approach is best compared against Layer 2, and the results in Table 5.2 show that BPCM adjustments yield markedly lower errors, demonstrating the benefit of enforcing arbitrage consistency before model fitting.

The overall results show that errors decline consistently from one layer to the next in both regimes. Layer 1 removes the most evident structural inconsistencies, Layer 2 enhances the internal coherence of the calibration, and Layer 3 resolves most of the remaining discrepancies. The result is an almost perfect fit for stable quotes and a calibration for unstable data that finds the right balance between accuracy and robustness.

Spread-based error assessment

A key indicator of calibration quality is whether model-implied prices lie within the market bid–ask envelope. For each option with bid y^{bid}, ask y^{ask}, mid y^{mid}, and model price y^{mod}, we measure the *relative pricing position*

$$R_{\text{pos}} \;=\; \frac{2\,(y^{\text{mod}} - y^{\text{mid}})}{y^{\text{ask}} - y^{\text{bid}}}, \tag{5.17}$$

so that $R_{\text{pos}} = -1$ at the bid, 0 at the mid, and $+1$ at the ask. Values in $[-1, 1]$ indicate that the model price lies inside the quoted spread, while $|R_{\text{pos}}| > 1$ indicates an out-of-spread violation. The spread-based diagnostics show how calibrated prices align with bid–ask envelopes. Within the spread ($|R_{\text{pos}}| \leq 1$), the conditional expectation $\mathbb{E}[R_{\text{pos}} \mid |R_{\text{pos}}| \leq 1]$ is negative in both regimes (e.g. -0.159 in the stable subset, -0.306 in the unstable subset under Layer 3), indicating that fitted prices lean towards bids rather than asks. This is in line with the results from the work of Hagströmer (2019).

To quantify the magnitude of violations, define the *excess distance*

$$D_{\text{exc}} = \begin{cases} 0, & R_{\text{pos}} \in [-1, 1], \\ y^{\text{mod}} - y^{\text{ask}}, & R_{\text{pos}} > 1, \\ y^{\text{bid}} - y^{\text{mod}}, & R_{\text{pos}} < -1, \end{cases} \tag{5.18}$$

which vanishes when the model price is inside the spread, and takes strictly positive values proportional to the distance by which the model price exceeds the bid–ask bounds.

Let $\mathcal{I}$ index all observations ($N = |\mathcal{I}|$) and let $S = \{\, i \in \mathcal{I} : |R_{\text{pos},i}| > 1 \,\}$ denote the set of violations ($N_{\text{out}} = |S|$). We report two complementary sample means for the violation size:

$$\overline{D}_{\text{exc}}^{\text{overall}} = \frac{1}{N} \sum_{i \in \mathcal{I}} D_{\text{exc}}^{(i)}, \qquad \overline{D}_{\text{exc}}^{\text{viol.}} = \frac{1}{N_{\text{out}}} \sum_{i \in S} D_{\text{exc}}^{(i)}.$$

Table 5.3 reports these spread-based diagnostics across methods and regimes, showing both the conditional mean position within the spread and the average magnitude of excess distances, computed over all contracts and over violations only. The first average is calculated over the full sample (including zeros for in-spread quotes); the second average is calculated only over out-of-spread observations. For violations, the mean excess distance is negligible in the stable subset (0.011 overall, 0.133 conditional) and substantially larger in the unstable subset; however, it is substantially reduced by Layer 3 (0.691 overall, 0.703 conditional). These diagnostics confirm that BPCM delivers near-perfect alignment in stable conditions and keeps deviations within bounds in unstable periods. Figure 5.1 shows how option quotes are distributed relative to

Table 5.3: Spread-based diagnostics by method and regime: conditional mean position inside the spread, mean excess distance over all observations, and mean excess distance over violations only.

Method	$\mathbb{E}[R_{\text{pos}} \mid \lvert R_{\text{pos}}\rvert \le 1]$		$\overline{D}_{\text{exc}}^{\text{overall}}$		$\overline{D}_{\text{exc}}^{\text{viol.}}$	
	Stable	Unstable	Stable	Unstable	Stable	Unstable
Traditional	−0.1285	−0.4992	0.1481	2.5930	0.2923	2.6263
Layer 2	−0.1517	−0.2786	0.0606	1.9730	0.1438	2.0093
Layer 3	−0.1590	−0.3059	0.0112	0.6910	0.1325	0.7033

bid–ask spreads under each calibration method.

In the **stable subset**, the effect of the layers is very clear. As expected, Layer 1

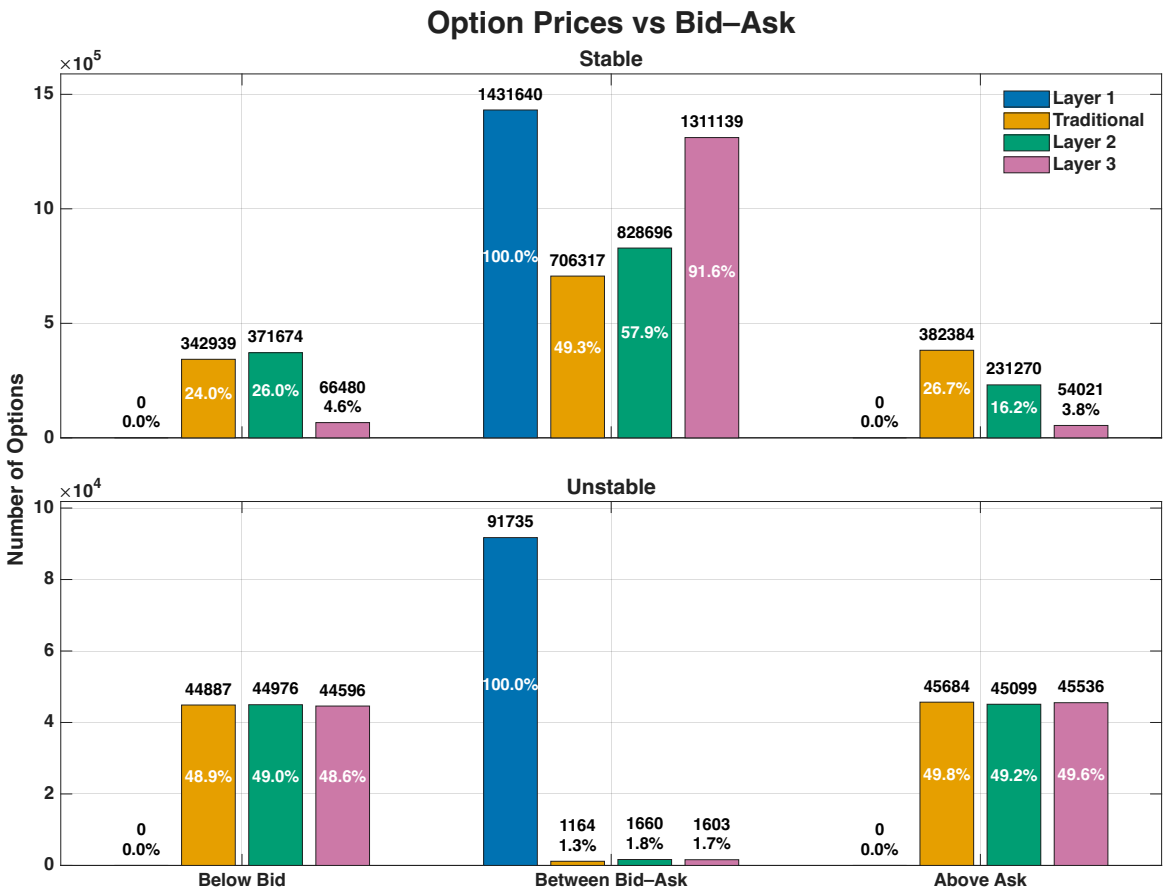


Figure 5.1: Quote assignment relative to bid–ask spreads under different calibration methods.

(parity enforcement) places all quotes inside their bid–ask envelopes by construction. The traditional direct Heston calibration performs poorly, with only about half (49%) of quotes inside the spread and large fractions classified as above or below. Layer 2, which applies the same Heston fit but to parity-cleaned data, raises the in-spread share to nearly 58%. Layer 3 achieves the strongest performance, with over 91% of prices lying between the bid and ask. This progression shows that cleaning (Layer 1) and layered fitting (Layers 2–3) materially increase the consistency of calibrated prices with market bounds.

In the **unstable subset**, by contrast, the vast majority of quotes fall outside the spread under all methods. Less than 2% of model prices lie within their original bid–ask ranges, regardless of calibration, reflecting the fact that spreads in this regime are often too narrow to accommodate noisy or inconsistent quotes. Here, the value of the layered approach lies not in pulling prices inside the envelope—which is rarely possible—but in bounding violations symmetrically and avoiding extreme misclassifications. Compared with the traditional calibration, the BPCM layers provide

a more disciplined adjustment that respects no-arbitrage while still tracking market dynamics.

Time-Series of Median Residual

To assess stability through time, we compute **median absolute residuals** in non-overlapping five-day bins. That is, for each five-day window, we take the absolute residuals of all contracts and record their median. This aggregation smooths day-to-day noise while preserving the temporal structure of calibration performance, making it easier to identify persistent improvements or stress-related spikes. Figure 5.2 shows these median residuals, comparing Traditional vs. Layer 2 (top row) and Layer 1 vs. Layer 3 (bottom row) for both stable (left column) and unstable (right column) subsets. Panels use logarithmic y-axes and *not* the same vertical range; comparisons of levels should therefore be made *within* a panel.

Stable (left column). In the top-left panel (Traditional vs. Layer 2), both series follow the same slow-moving pattern, reflecting that the underlying Heston calibration procedure is identical in both cases. The difference is that Traditional calibration uses raw quotes, whereas Layer 2 applies the same fit to the Layer 1 parity-adjusted data. The resulting dynamics are therefore unchanged, but the error level is uniformly lower for Layer 2, showing that the filtering step is important. In particular, Layer 1 removes arbitrage inconsistencies in the raw quotes while preserving the overall time-series pattern, so the Heston model continues to fit well once applied to the cleaned data. In the bottom-left panel (Layer 1 vs. Layer 3), both curves remain in the same very low range (10^{-14}–10^{-5} on the log scale). Layer 3 sits slightly above the parity baseline at times, but the deviations are small and do not alter the overall error scale. This indicates that in the stable regime, the Heston calibration yields residuals on the same minimal scale as those obtained under parity, demonstrating that the model aligns closely with the arbitrage-consistent structure of the data.

Unstable (right column). In the top-right panel (Traditional vs. Layer 2), both methods follow the same broad pattern, but the error levels are far higher than in the stable regime. Layer 2 lies consistently below the Traditional calibration, showing that the parity-adjusted fit reduces residuals while preserving the timing of stress-related spikes. In the bottom-right panel (Layer 1 vs. Layer 3), the two curves remain within the same broad range (10^{-10} to 10^{-1}), but Layer 3 displays noticeably greater dispersion. Residuals swing more widely—occasionally dropping near 10^{-10} and then rising

close to 10^{-1}—with an average level around 10^{-1}. This indicates that under unstable conditions, the joint refinement cannot fully suppress the variability of residuals. Instead, Layer 3 inherits the structural noise of the data, producing a more irregular pattern than in Layer 1, which reflects genuine market instability rather than a weakness of the calibration method.

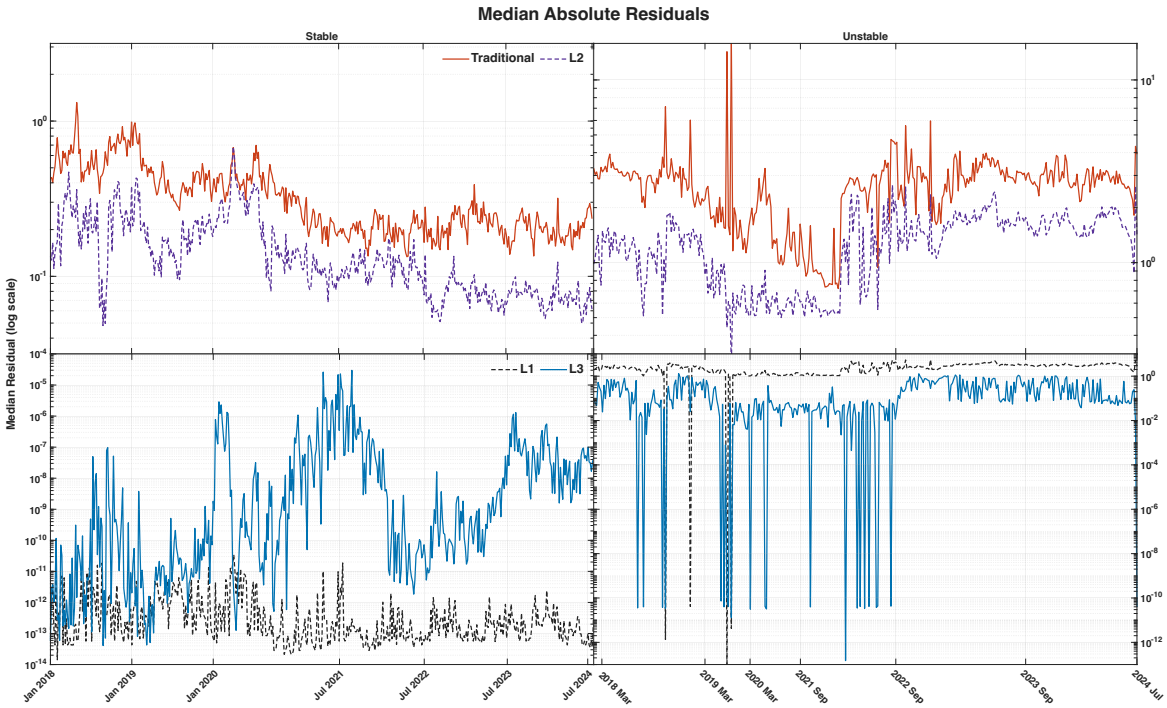


Figure 5.2: Median absolute residuals in 5-day bins.

Spikes of large residuals are extremely rare: only 773 cases, corresponding to about 0.05% of the 1,431,656 total observations, exceed a magnitude of 0.5. Almost all of these coincide with periods of stress or illiquidity (see Appendix F.2), so they are best interpreted as signals of genuine market disruption rather than calibration failure. Importantly, the final residuals never surpass the initial errors, confirming that BPCM contains mispricings even under extreme conditions.

Overall, the time-series analysis confirms that BPCM delivers accurate and robust calibrations. For stable data, the final residuals are negligible, with more than 91% of prices falling within bid–ask bounds and a near-zero RMSE. For unstable data, large initial violations are substantially reduced (RMSE falls from 4.82 to 1.20), and the calibration remains bounded despite noisy inputs. The layered structure is crucial: Layer 1 removes structural inconsistencies, Layer 2 provides a balanced fit that reduces errors for calls and puts simultaneously, and Layer 3 fine-tunes without overfitting, allowing small asymmetries to appear only on the side with tighter market constraints (e.g.

higher liquidity or narrower spreads). Residual spikes are thus diagnostic, signalling episodes of market stress rather than model failure. Hence, BPCM offers both strong calibration performance and transparency into when and why deviations persist.

Third-Layer Iteration

To evaluate how calibration effort and residuals evolve across regimes, we track the number of iterations required for convergence in the third layer for both stable and unstable datasets.

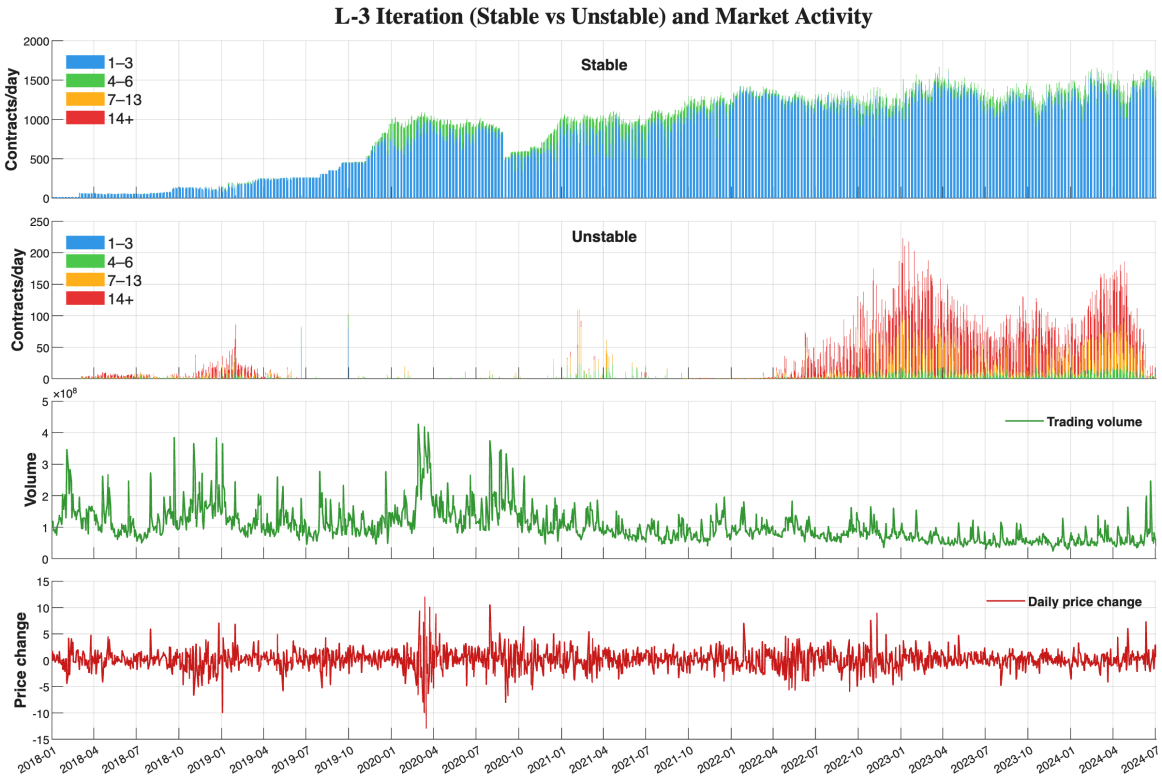


Figure 5.3: Fit diagnostics over 2018–2024.

Figure 5.3 illustrates the calibration dynamics across regimes. In stable markets, calibration converged rapidly, with most days requiring only the standard three iterations and residuals subsequently reduced to negligible levels, supported by the availability of a broad cross-section of contracts. In unstable markets, the number of third-layer iterations increased substantially, and many days required extended adjustments. Residual errors remained elevated during episodes of heightened volatility, reflecting reduced and more irregular quote availability, wider bid–ask spreads, and frequent parity violations. The market volume panel shows that unstable datasets not only featured lower average activity but also volatile surges around macroeconomic

events, while the daily price-change panel highlights large swings that amplified calibration noise.

To summarise, the third-layer iteration serves as a diagnostic of market quality. With clean and liquid data, three iterations typically suffice to achieve stability. However, under turbulent conditions, additional iterations reveal inconsistencies in quotes and underscore the increased difficulty of fitting the volatility surface. The bid–ask classification further confirms this point, as the proportions of quotes falling below the bid and above the ask are nearly identical, indicating that the calibration introduces no systematic bias. In stable regimes, this balance is achieved with negligible residual noise, whereas in unstable regimes, the percentages remain close despite larger violations, showing that the limitation lies in the data rather than the model. Thus, the Heston calibration within BPCM delivers the best attainable fit under the observed conditions, and further improvement would require either filtering noisy quotes or adopting alternative modelling assumptions.

5.5 Computational Cost Analysis

The calibration runtime is decomposed into four components. Layer 1 (**L1**) corresponds to the data-conditioning stage evaluated prior to calibration (cf. Algorithm 5.1). Layer 2 (**L2**) daily parameter calibration using analytical sensitivities (cf. Algorithm 5.2 and the analytic-Jacobian solver in Algorithm 5.4). The traditional calibration (**Trad**) corresponds to a direct, single-stage baseline calibration that applies the same Layer 2 optimisation with analytical sensitivities to the raw option quotes, *without* the Layer 1 filtering and structural repair. Layer 3 (**L3**) corresponds to the final joint refinement stage evaluated in the runtime experiments (cf. Algorithm 5.3), operating on option datasets aggregated across multiple trading days.

All timing experiments were conducted on a MacBook Pro (16-inch, November 2024) equipped with an Apple M4 Max processor and 48 GB of unified memory, running macOS Tahoe 26.1. Runtimes are reported as wall-clock elapsed time measured using built-in timing routines. The calibration routines are implemented using MATLAB's *parfor* construct, enabling parallel execution across multiple workers.

Experiments are grouped by rolling-window length $W \in \{1, 5, 16, 32, 64, 128\}$ trading days and by regime (*stable* vs. *unstable*). For each (W, regime), Table 5.4 reports the total elapsed time T_{tot} (seconds), together with normalised metrics computed as

$$T_{\text{day}} := \frac{T_{\text{tot}}}{W} \qquad \text{and} \qquad T_{\text{opt}} := \frac{T_{\text{tot}}}{N_{\text{opt}}(W)},$$

Table 5.4: Runtime summary by window length for stable and unstable option sets (seconds).

W	Regime	$N_{\text{opt}}(W)$	Metric	L1	L2	Trad	L3
1	Stable	1,142	T_{tot}	2.046	5.3781	17.516	10.213
			T_{opt}	0.00179	0.004709	0.015338	0.00894
	Unstable	368	T_{tot}	0.426	5.3873	13.380	13.909
			T_{opt}	0.00116	0.014639	0.036358	0.03780
5	Stable	6,240	T_{tot}	2.251	56.861	52.387	67.105
			T_{day}	0.450	11.372	10.477	13.421
			T_{opt}	3.61×10^{-4}	0.009112	0.008395	0.01075
	Unstable	1,744	T_{tot}	0.978	42.372	53.407	70.260
			T_{day}	0.196	8.4744	10.681	14.052
			T_{opt}	5.61×10^{-4}	0.024296	0.030624	0.04029
16	Stable	20,950	T_{tot}	2.887	109.92	226.19	252.170
			T_{day}	0.180	6.8703	14.137	15.760
			T_{opt}	1.38×10^{-4}	0.005247	0.010797	0.01204
	Unstable	4,674	T_{tot}	2.093	124.41	160.62	213.680
			T_{day}	0.131	7.7755	10.039	13.355
			T_{opt}	4.48×10^{-4}	0.026617	0.034364	0.04572
32	Stable	43,820	T_{tot}	4.248	267.75	492.28	549.970
			T_{day}	0.133	8.3672	15.384	17.186
			T_{opt}	9.69×10^{-5}	0.006110	0.011234	0.01255
	Unstable	8,614	T_{tot}	2.909	223.47	314.91	391.100
			T_{day}	0.091	6.9835	9.8409	12.222
			T_{opt}	3.38×10^{-4}	0.025943	0.036558	0.04540
64	Stable	89,878	T_{tot}	6.422	462.28	1005.1	1082.300
			T_{day}	0.100	7.2231	15.704	16.911
			T_{opt}	7.15×10^{-5}	0.005143	0.011183	0.01204
	Unstable	16,404	T_{tot}	4.503	512.72	819.97	739.300
			T_{day}	0.070	8.0112	12.812	11.552
			T_{opt}	2.74×10^{-4}	0.031256	0.049986	0.04507
128	Stable	178,860	T_{tot}	13.131	651.66	1739.3	2373.800
			T_{day}	0.103	5.0911	13.588	18.545
			T_{opt}	7.34×10^{-5}	0.003643	0.009724	0.01327
	Unstable	26,974	T_{tot}	8.299	1022.0	1752.5	1192.800
			T_{day}	0.065	7.9842	13.691	9.319
			T_{opt}	3.08×10^{-4}	0.037888	0.064969	0.04422

where T_{day} is the average runtime per trading day (seconds/day), T_{opt} is the average runtime per option quote (seconds/option), and $N_{\text{opt}}(W)$ denotes the number of option quotes aggregated within the window of length W.

Table 5.4 shows that the total runtime T_{tot} increases monotonically with the window length W for all calibration components, reflecting the growth in the number of option quotes $N_{\text{opt}}(W)$ included in the calibration set. In contrast, the normalised metrics T_{opt} and T_{day} remain relatively stable across window lengths, indicating that the computational cost per option and per day does not deteriorate as the dataset size increases. Layer 1 contributes a negligible amount of runtime at all window lengths. The main computational cost is incurred by the optimisation-based stages, namely Layer 2, the Traditional baseline, and Layer 3. As the window length W increases, the total runtime of all three components grows primarily due to the increase in the number of option quotes $N_{\text{opt}}(W)$ processed, rather than from any deterioration in per-option efficiency. This is confirmed by the fact that the normalised per-option runtimes remain of comparable magnitude across window lengths.

Among these stages, the Traditional calibration is consistently more expensive than Layer 2, reflecting the absence of prior filtering and structural repairs. Layer 3 operates at the level of individual option pairs, performing a parameter optimisation for each call–put pair in the aggregated dataset, with a per-option runtime comparable to that of Layer 2. Finally, unstable datasets exhibit systematically higher normalised runtimes than stable datasets, reflecting increased optimisation difficulty.

5.6 Empirical Performance Overview

This section evaluates the empirical performance of BPCM over the full sample period. For clarity, we report results for call options only, as calls and puts share the same underlying inputs (S, q, r) and the same calibrated parameters $\boldsymbol{\theta}$. Because put–call parity enforces a common forward price, focusing on calls illustrates model behaviour without loss of generality.

Throughout, we distinguish between stable and unstable regimes as defined earlier. The classification combines spot-path checks with parity consistency. Within this framework, we evaluate the time-series behaviour of Heston parameters and reconstructed spot prices, focusing on stability and interpretability across regimes.

Reconstructed Spot Price Trajectories

A key aim of BPCM is to reconstruct a consistent trajectory for the underlying asset price from option quotes. By inferring the implied spot S_t, we can test how well the method translates option premiums into a reasonable path that moves within the market's low and high bounds.

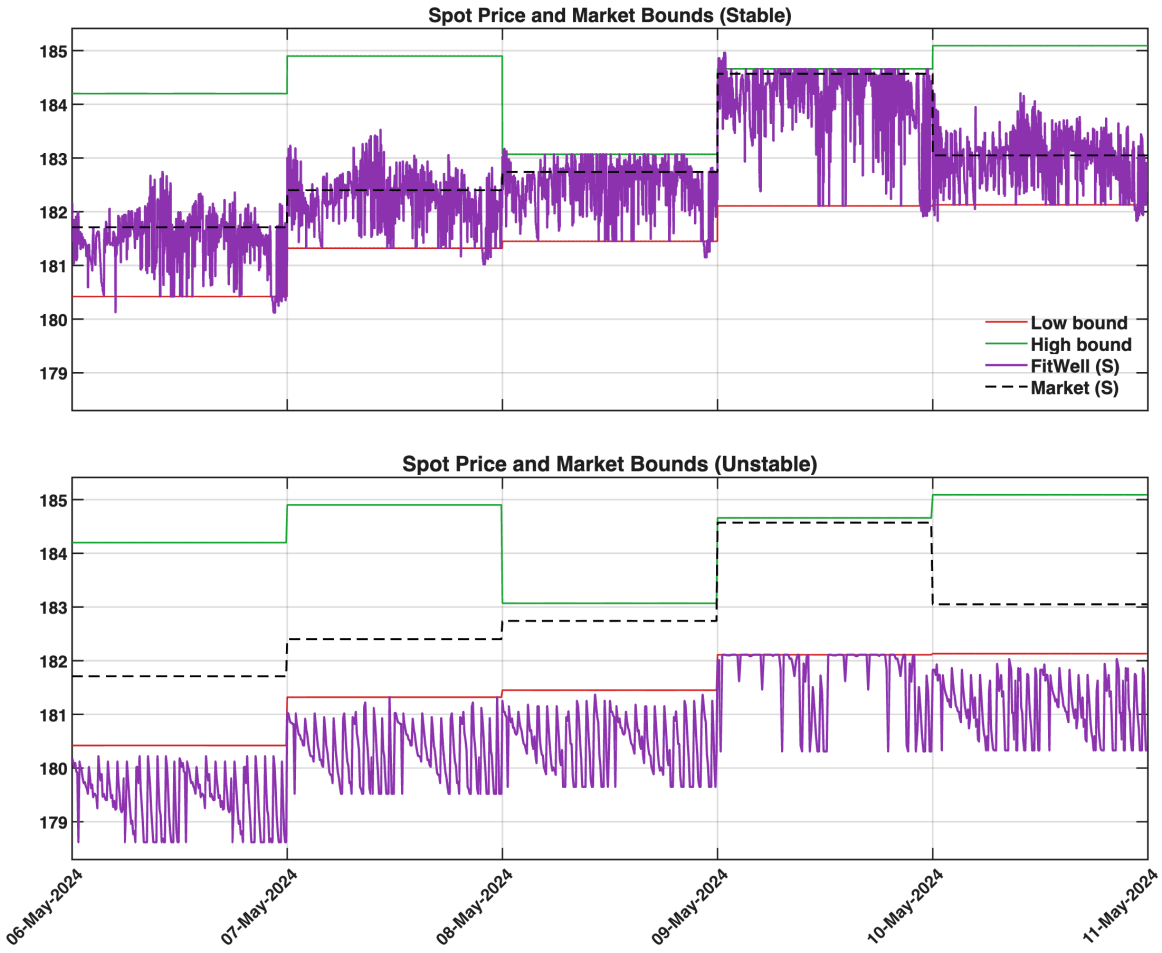


Figure 5.4: Intra-day implied spot S_t over 6–11 May 2024.

Figure 5.4 shows the intra-day spot over five trading days (6–11 May 2024). Because option quotes arrive irregularly, we rescale time to a uniform 9:30–16:00 trading day. The implied spot remains well within the actual high–low bounds and typically aligns with the market closing price. Only small jumps near price extremes are observed, reflecting individual noisy quotes rather than a structural misfit.

Figure 5.5 shows the reconstructed daily spot trajectories from 2018–2024, with stable data in the top panels and unstable data in the bottom panels. The left panel covers the period before the Apple 4-for-1 stock split (up to August 31, 2020), and the right panel covers the period after the split (from September 1, 2020). In each

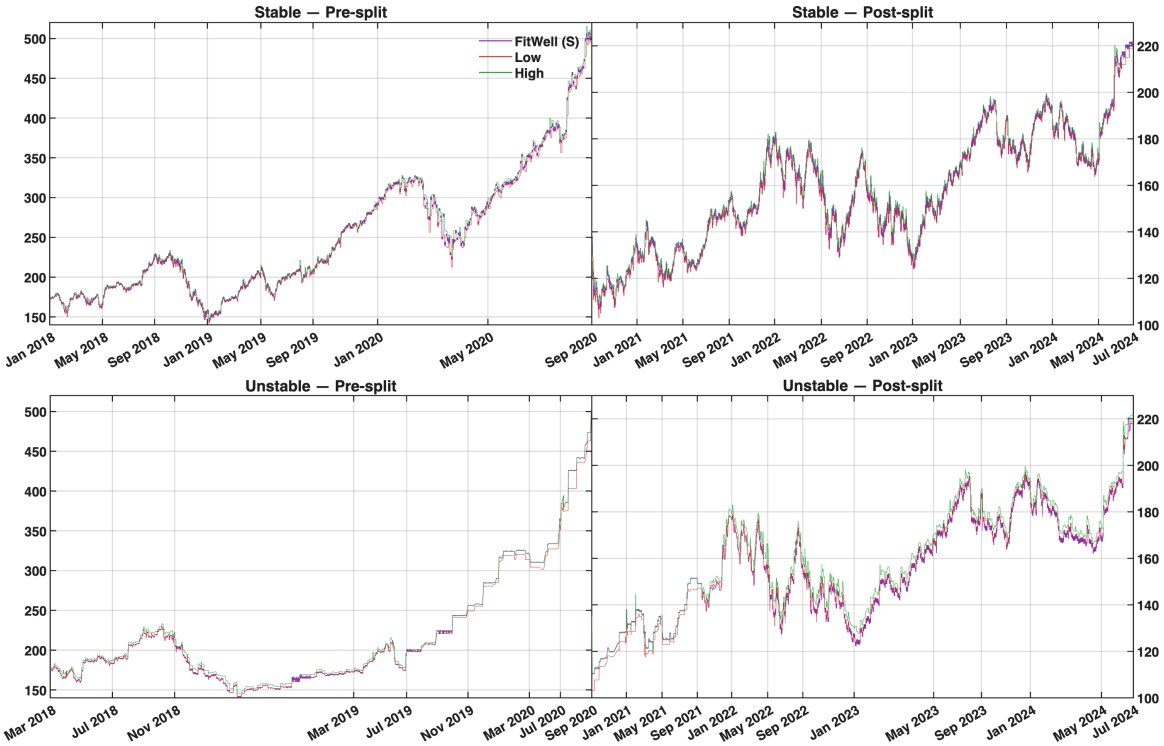


Figure 5.5: Daily implied spot S_t, 2018–2024.

panel, the green line represents the daily high, the red line represents the daily low, and the pink markers represent the model-implied spot. In the stable regime, the implied spot lies almost entirely within the observed high–low bounds, confirming that BPCM recovers the underlying asset dynamics under high-quality data. In the unstable regime, the model-implied spot varies more widely and tends to lie systematically below the observed trading range. This reflects the impact of noisy or inconsistent quotes. Rather than tracking anomalous spikes, the calibration smooths the estimate toward levels that remain consistent with no-arbitrage conditions.

Periods of heightened volatility, such as March 2020, are clearly visible in Figure 5.9. During this episode, the variance parameters (v_0 and $\bar{v}$) increased sharply, reflecting the extreme volatility observed in the market. The comparison with the stable regime confirms that deviations in unstable periods arise primarily from data imperfections rather than model failure.

Consistency of Implied Prices with Observed High–Low Bounds

An important diagnostic of calibration integrity is whether the model's implied spot is consistent with the underlying's observed trading range. Figure 5.6 reports two complementary perspectives, where the left panel classifies each implied spot relative

to the daily high–low bounds and the right panel shows the distribution of differences between market and model-implied spots.

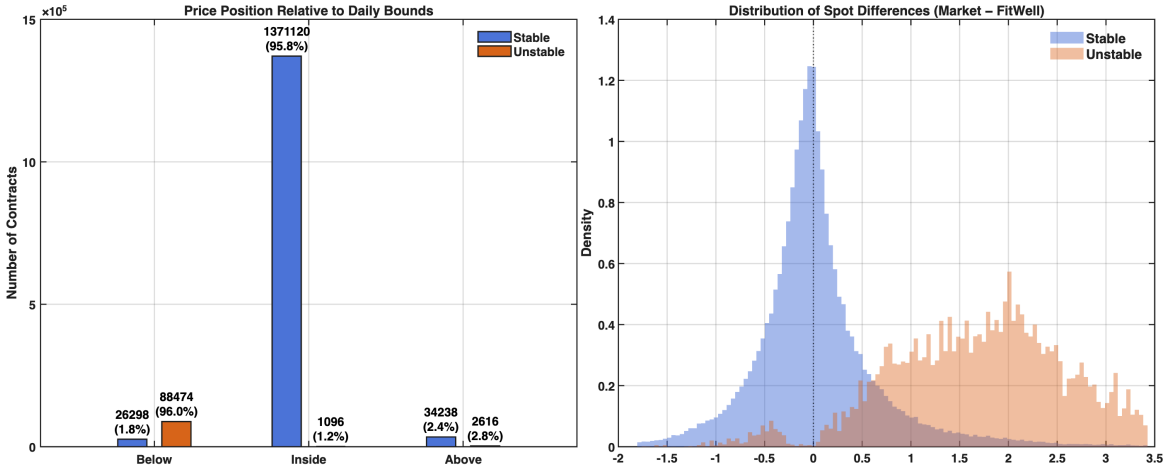


Figure 5.6: Implied spot vs. observed bounds (left) and spot differences (right).

In the stable subset (blue bars), the model's spot lies inside the observed range in 95.8% of cases, with only small fractions falling below (1.8%) or above (2.4%). Such small deviations are plausibly explained by data noise or reporting mismatches rather than model misspecification. The right panel confirms this, showing a distribution tightly centred at zero with narrow tails, i.e. negligible differences between market and model-implied spots.

In the unstable subset (orange bars), only 1.2% of implied spots fall within the high–low range, while about 96% lie below the daily lows. This systematic downward shift reflects noisy or inflated quotes in stressed conditions, as the calibration prioritises conservative spot estimates over aligning with unreliable market observations. The right panel shows a broad, right-skewed distribution of differences, consistent with this downward bias.

BPCM reconstructs spot paths that are accurate in stable regimes and well-regulated in unstable regimes. The calibration avoids overfitting noisy quotes, aligning implied spots closely with observed ranges when the data are consistent. Under stress, it emphasises robustness and economic plausibility over exact matching. This ensures that the reconstructed spot and parameter paths remain economically interpretable across different volatility regimes.

Stability and Interpretation of Heston Parameters

We now examine the calibrated Heston parameters to assess their stability over time and the impact of BPCM relative to a traditional calibration. In BPCM, Layer 2

produces daily parameter estimates after enforcing put–call parity and structural consistency. By contrast, the traditional calibration applies the same optimisation procedure directly to raw mid-quotes and underlying prices, without any adjustments. This side-by-side comparison allows us to test whether BPCM not only reduces pricing errors but also yields more stable and economically interpretable parameter dynamics across regimes.

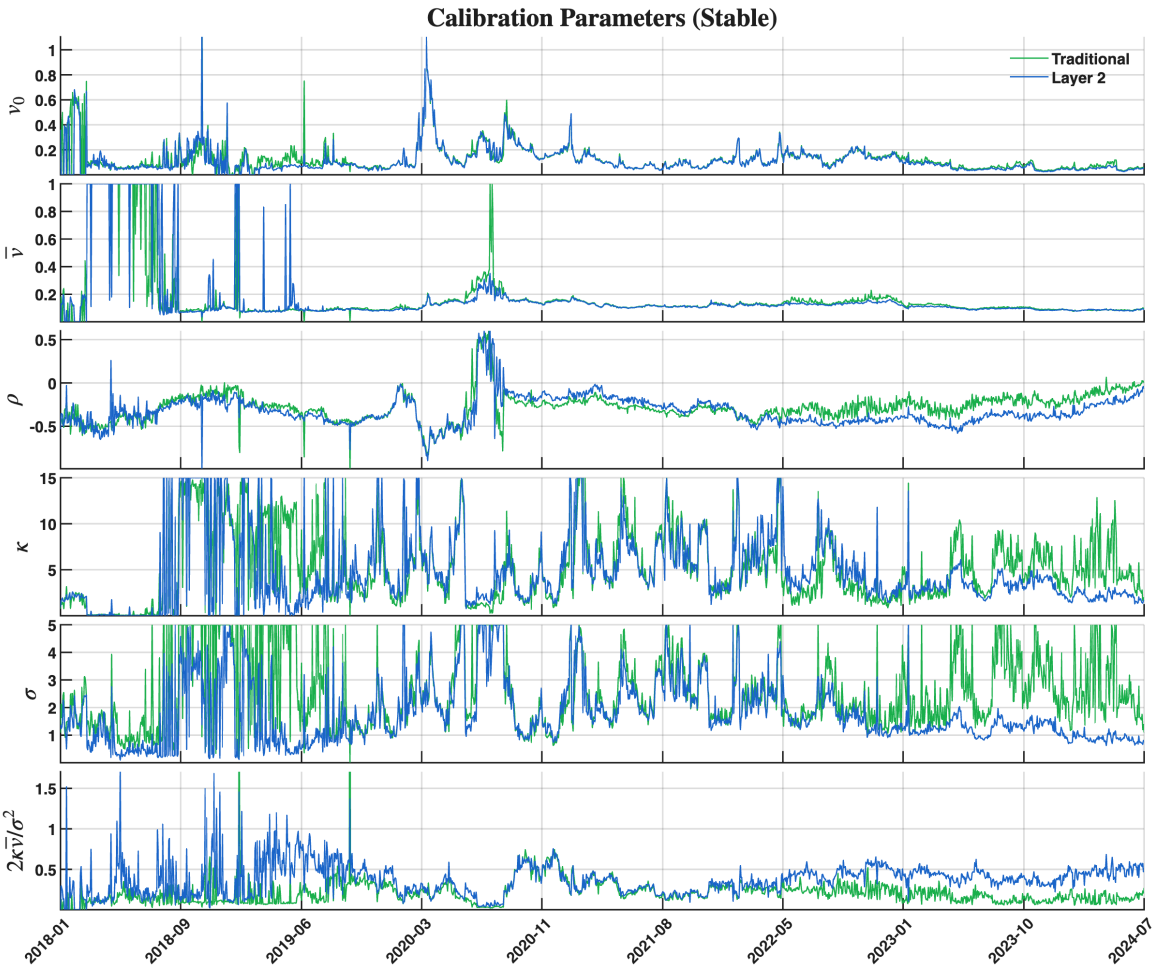


Figure 5.7: Time series of Heston parameters for the stable dataset.

Figure 5.7 reports the parameter evolution for the stable set. The evolution of Layer 2 estimates is more stable compared to that of traditional models, with v_0 and $\bar{v}$ responding plausibly to market shocks such as the COVID-19 volatility spike and the 2022 rate-hike cycle, while ρ remains consistently negative, in line with the equity leverage effect. In contrast, the traditional calibration shows visibly noisier fluctuations, with parameters such as κ and σ exhibiting jumps and reversals unrelated to market events. This confirms that BPCM's preprocessing and parity enforcement reduce spurious variability and deliver more economically interpretable dynamics.

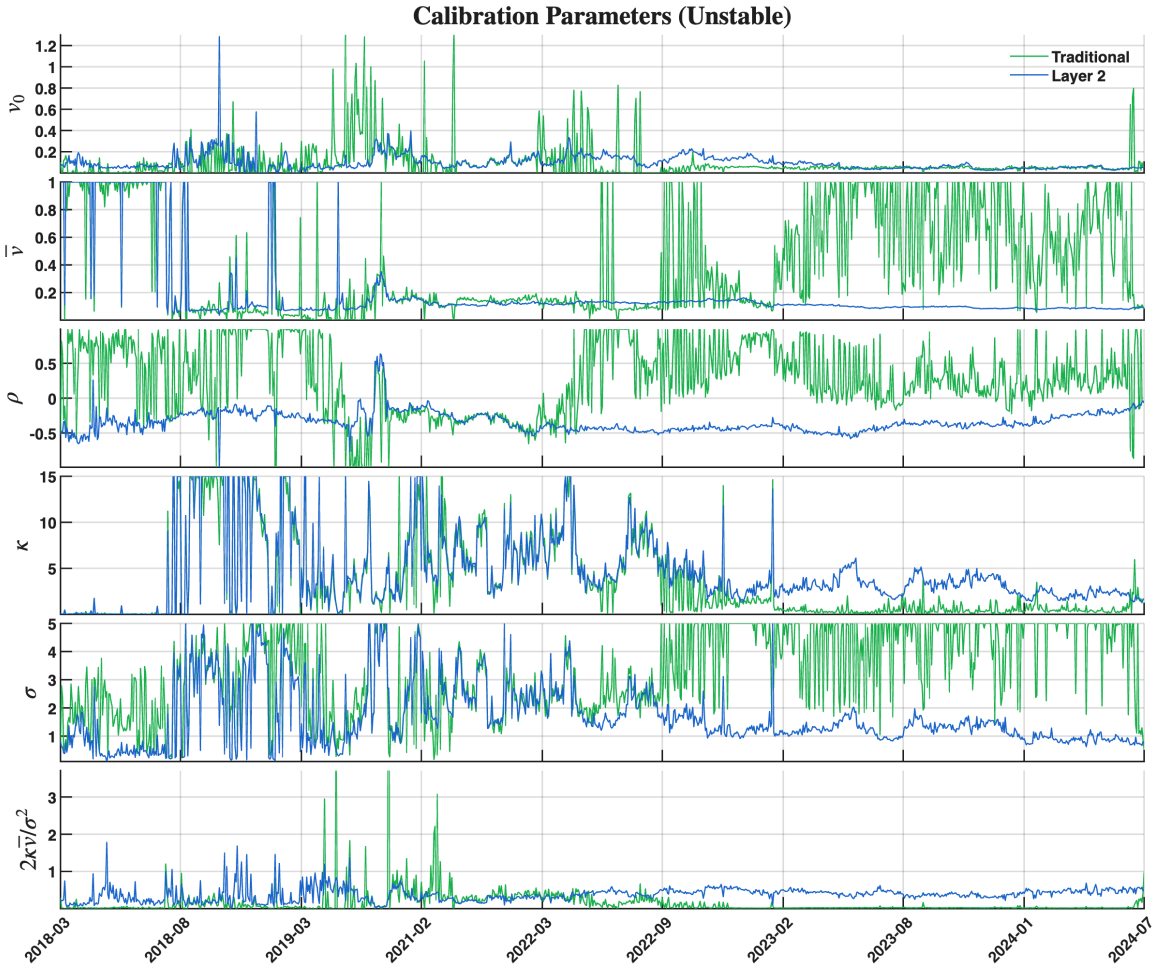


Figure 5.8: Time series of Heston parameters for the unstable dataset.

Figure 5.8 shows the corresponding results for the unstable set. Here the contrast is even sharper, as Layer 2 trajectories remain bounded and coherent despite noisy quotes, while the traditional calibration produces large swings, particularly in $\bar{v}$, ρ and σ, that are difficult to reconcile with economic plausibility. Even in stress regimes, BPCM constrains parameters within reasonable ranges and stabilises estimates, whereas the traditional fit is more erratic, confirming that the framework improves fit quality and robustness across both stable and unstable conditions.

Figure 5.9 presents the time series of Heston parameters calibrated on the stable dataset (2018–2024). The Layer 2 trajectories evolve in a stable manner and align with known market events. Notably, the instantaneous variance v_0 and long-run variance $\bar{v}$ spike sharply during the early 2020 COVID-19 volatility shock, and again around the 2022 interest-rate hike cycle, capturing the shifts in volatility regime during those periods. The σ shows a similar pattern of surges, reflecting that volatility itself became more volatile in those turbulent episodes. Meanwhile, the correlation ρ remains consis-

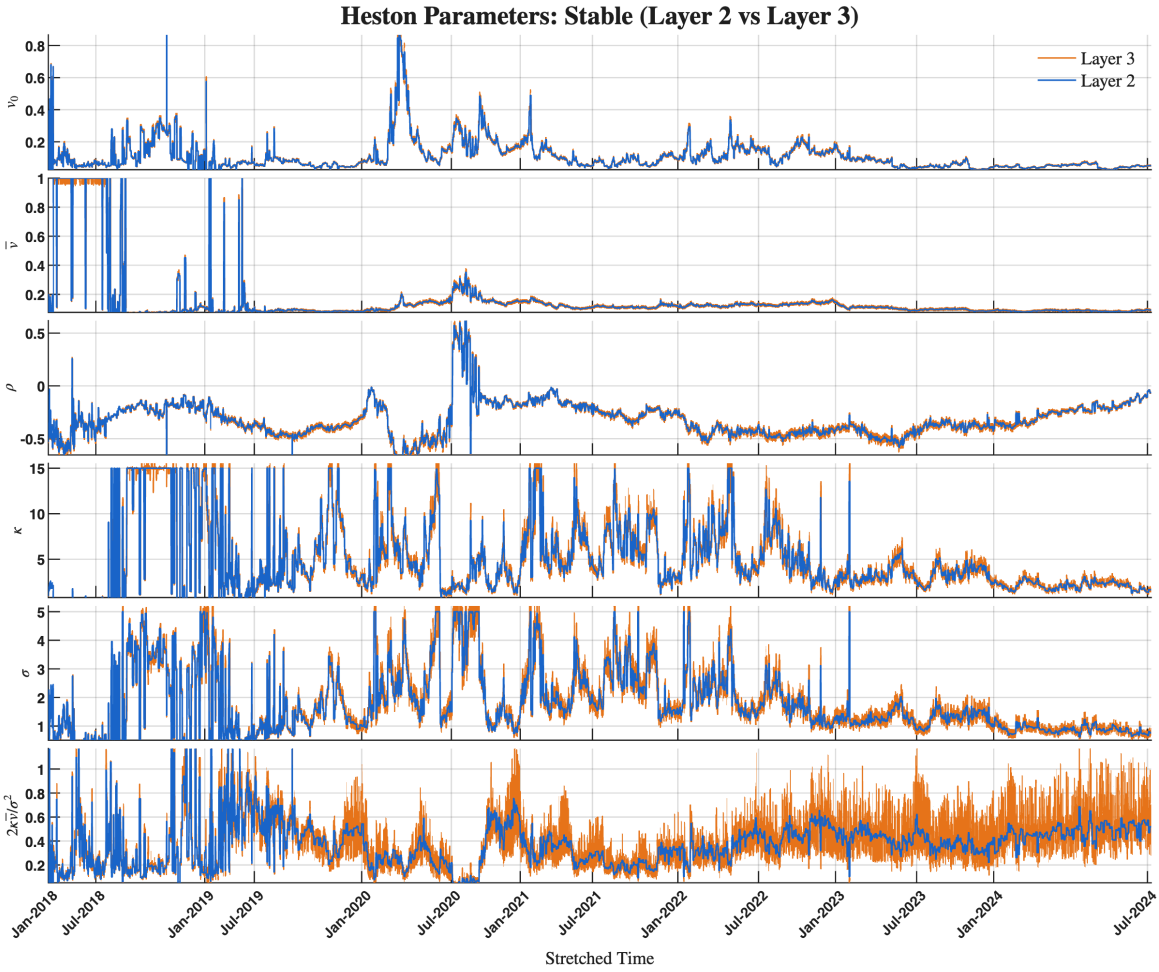


Figure 5.9: Time series of Heston parameters calibrated on stable data across 2018–2024.

tently negative throughout (consistent with the equity *leverage effect*). Crucially, the Layer 3 adjustments (visible as thin coloured bands around the blue curves) are very small. The contract-level refinements in the final layer stay tightly clustered around the daily Layer 2 values, indicating that the broad structure set by the Layer 2 calibration is largely unaltered by the final fitting. In other words, BPCM's last layer makes only minor local corrections and does not disturb the economically sensible trends in $\boldsymbol{\theta}$.

Figure 5.10 shows a five-day window of the stable-market calibration (6–11 May 2024). Layer 2 delivers smooth daily baselines for all parameters, while Layer 3 provides higher-frequency fluctuations around these baselines. In particular, ρ remains consistently negative, whereas κ and σ display visible intraday co-movements. These short-term fluctuations should be interpreted as calibration noise caused by small inconsistencies in the option quotes, rather than as true changes in market volatility. The joint movement of κ and σ reflects their structural link in the CIR variance

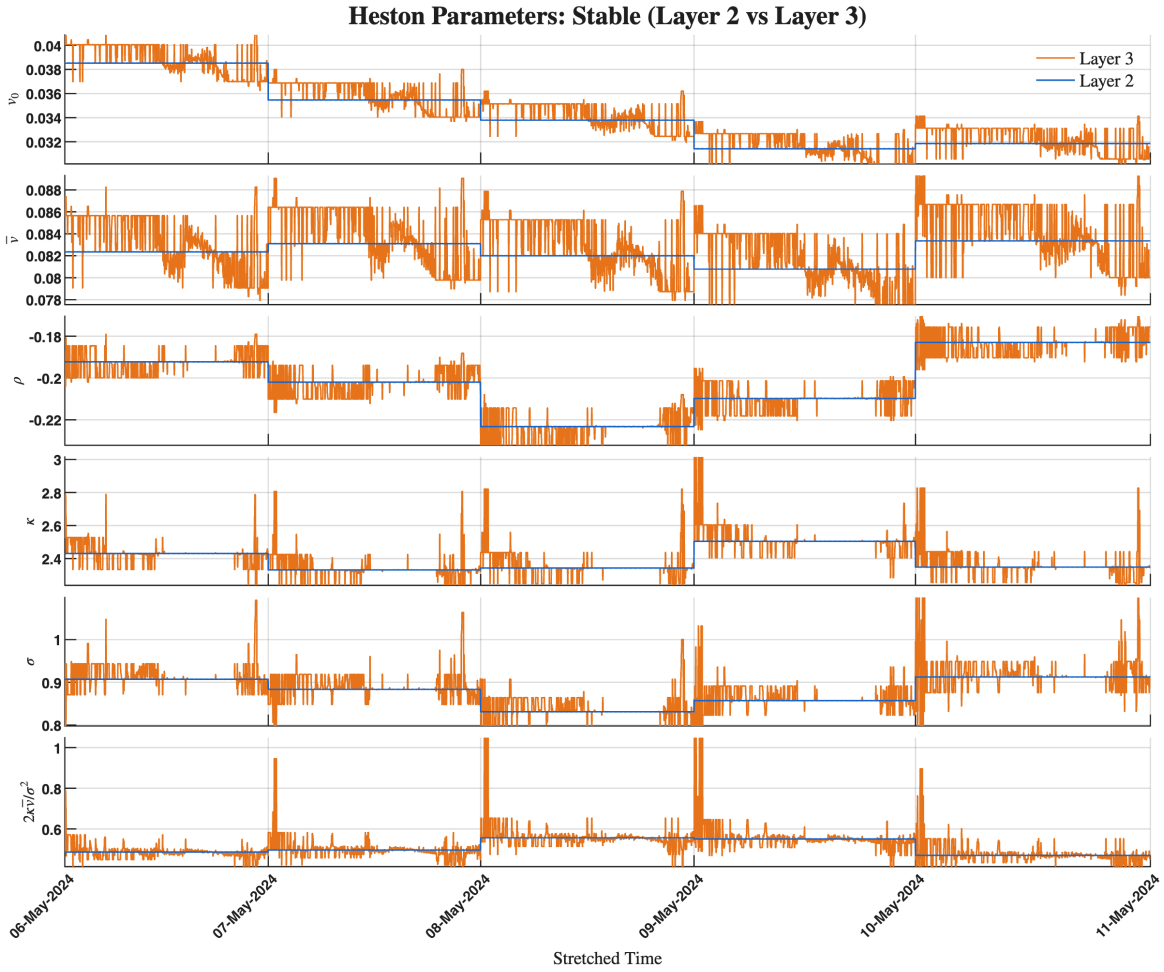


Figure 5.10: Heston parameter estimates over a five-day stable sample (6–11 May 2024).

process, where the long-run variance is given by

$$\lim_{t\to\infty} \mathrm{Var}(v_t) = \frac{\bar{v}\,\sigma^2}{2\kappa},$$

as shown by Cox et al. (1985). Because the stationary variance depends primarily on the ratio σ^2/κ, the calibration often moves κ and σ together so that this ratio stays nearly unchanged. This structural dependence explains the Layer 3 co-movement of these parameters, as the calibration surface is relatively flat along directions where σ^2/κ remains constant.

Moving to the unstable regime, Figure 5.11 shows the parameter evolution when BPCM is calibrated exclusively on the unstable data subset. The Layer 2 parameter paths still exhibit reasonable smoothness, although some parameters (notably κ and σ) fluctuate more than in the stable case, reflecting the higher noise and structural uncertainty in the underlying quotes. The Layer 3 adjustments (orange overlays) are

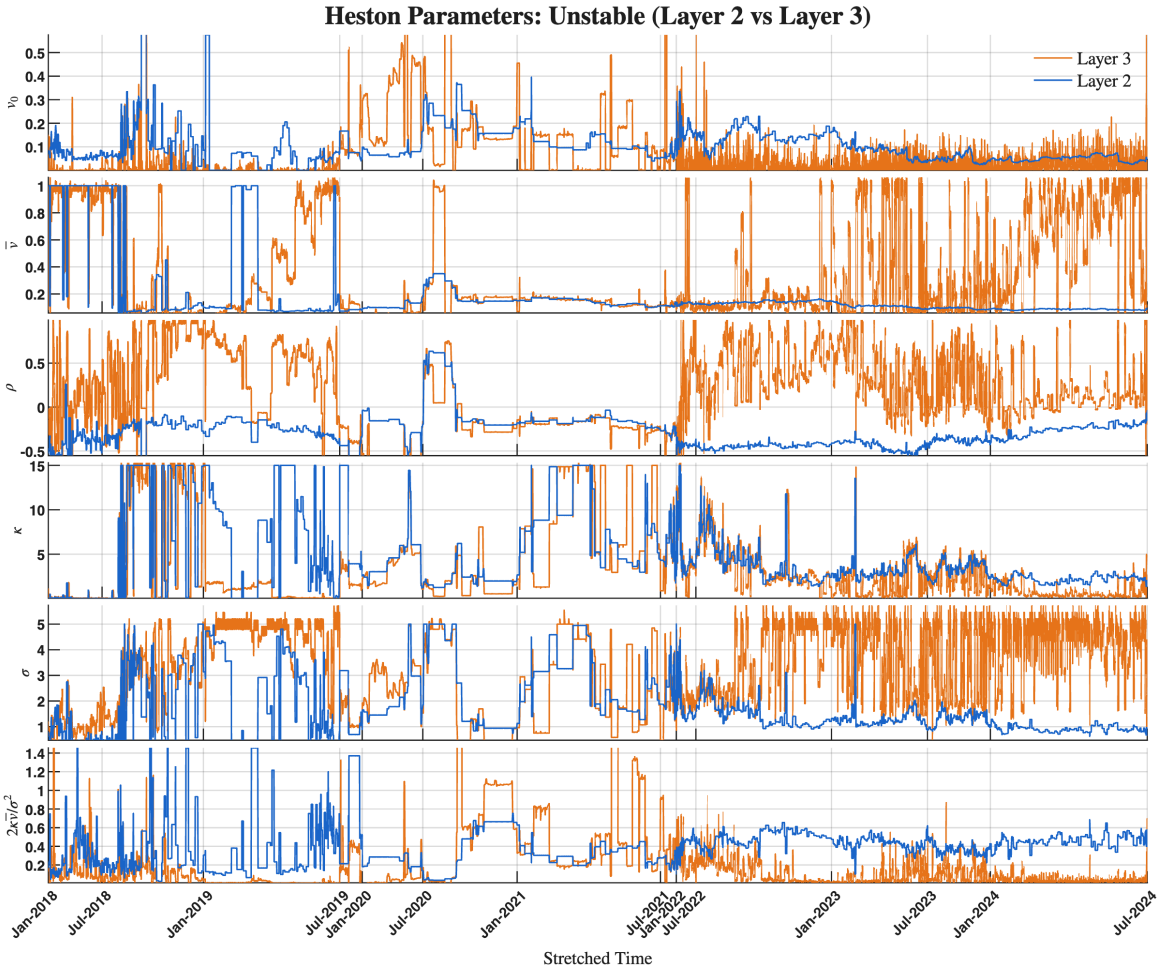


Figure 5.11: Time series of Heston parameters calibrated using unstable data.

visibly wider and more erratic here, especially for v_0, $\bar{v}$, and σ, which is unsurprising given the inconsistent input data. Despite this increased variability, the overall structure of the calibration is preserved: for instance, ρ remains negative throughout the entire period, even as the other parameters bob around.

The Feller (Cox–Ingersoll–Ross) positivity condition requires $2\kappa\bar{v} > \sigma^2$ in the Heston model to ensure that the variance process stays strictly positive. In practice, however, many calibrations do not satisfy this condition (Clevenhaus et al., 2024; Cui et al., 2017; Ribeiro and Poulsen, 2013), often during volatile market episodes. As Andersen (Andersen, 2008) emphasises, such violations are common and do not prevent the model from being applied, since the characteristic function remains well defined and option prices can still be computed. The Feller condition should therefore be regarded as a theoretical guideline, not an empirical requirement for calibration.

5.7 Interest Rate and Dividend Yield Adjustments

In addition to the Heston parameters, BPCM's layer 3 permits bounded adjustments to external market inputs, specifically the risk-free interest rate $r(t)$ and dividend yield $q(t)$. These inputs are typically derived from published yield curves or dividend schedules; however, in practice, the effective values implied by option quotes vary slightly across maturities and trading days. To reconcile such inconsistencies, BPCM allows small, bounded shifts in r and q during calibration.

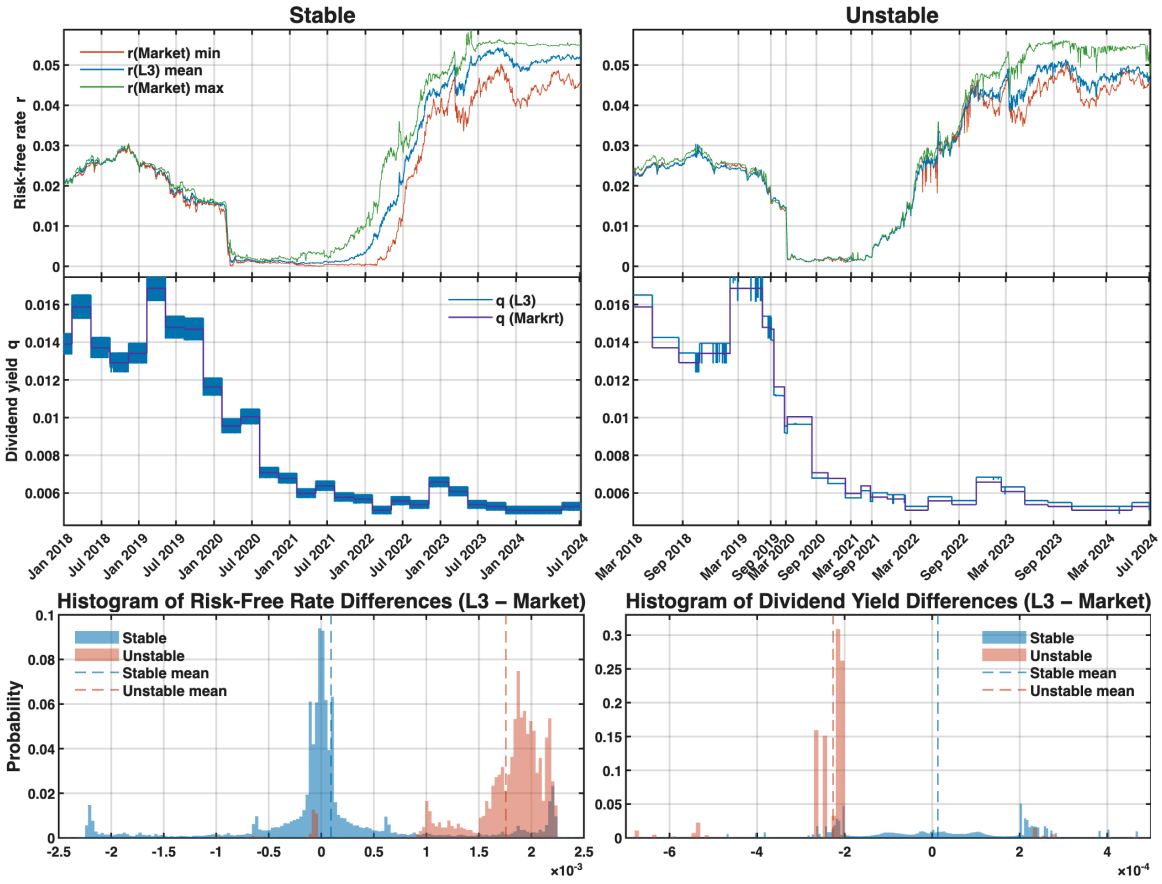


Figure 5.12: Daily interest rate and dividend yield adjustments under stable and unstable regimes.

Figure 5.12 illustrates the adjustment process. The first row shows the interest rate inputs. For each trading day, multiple maturities imply slightly different effective rates; the green and red lines plot the daily minimum and maximum across these maturities. The blue line represents the Layer 3 adjustment, defined as the daily mean of the implied rates after optimisation. It consistently lies between the bounds, confirming that the final layer does not distort the curve but rather centres it within the observed range. The second row shows the dividend yield, where the dark pink

line reflects the externally sourced dividend schedule reported in Table 4.1, and the red line is the Layer 3 estimate, which fluctuates only slightly around this baseline.

The third row shows the differences between baseline inputs and Layer 3 adjustments for both $r(t)$ and $q(t)$. Under stable conditions, the deviations are negligible (about 5×10^{-4}), while under unstable conditions, they are noisier, as the optimiser distributes the error across these inputs.

These results show that the third layer functions as a bounded error-correction mechanism. In stable markets, adjustments are minimal, while in unstable markets, they provide a practical way to reconcile fragmented data without forcing unrealistic parameter shifts. Thus, the calibrated paths for $r(t)$ and $q(t)$ remain aligned with macroeconomic and dividend realities while improving fit quality.

5.8 Gradient Dynamics of Heston Parameters

In this section we analyse the sensitivities $\partial C/\partial\theta_j$ of the Heston option price with respect to the model parameters. These derivatives are computed analytically using the closed-form expressions for the partial derivatives of the Heston characteristic function derived by Cui et al. (2017), which is substituted into the Heston pricing integral and evaluated within the standard Heston pricing framework.

From Equation (3.4), the calibration objective function is defined as

$$f(\theta) = \tfrac{1}{2}\sum_{i=1}^{n}\left[C^{\text{mod}}(\theta; K_i, \tau_i) - C_i^{\text{mkt}}\right]^2.$$

The gradient of f with respect to a model parameter θ_j is given by

$$\frac{\partial f}{\partial \theta_j} = \sum_{i=1}^{n}\left(C^{\text{mod}}(\theta; K_i, \tau_i) - C_i^{\text{mkt}}\right)\cdot\frac{\partial C^{\text{mod}}}{\partial \theta_j}.$$

This expression highlights that the calibration gradient is determined by the derivatives of the model price C^{mod} with respect to the parameters. These sensitivities scale the contribution of each pricing error in the objective function, and their variation across moneyness and maturity indicates which regions of the option surface drive the calibration. For this reason, we now compute the explicit parameter derivatives:

$$\frac{\partial C^{\text{mod}}}{\partial v_0}, \quad \frac{\partial C^{\text{mod}}}{\partial \bar{v}}, \quad \frac{\partial C^{\text{mod}}}{\partial \rho}, \quad \frac{\partial C^{\text{mod}}}{\partial \kappa}, \quad \frac{\partial C^{\text{mod}}}{\partial \sigma}.$$

We study the analytical gradients of the Heston price with respect to the parame-

ters $\boldsymbol{\theta} = (v_0, \bar{v}, \rho, \kappa, \sigma)$ as functions of log-moneyness and time to maturity. Let

$$x = \ln(K/F), \qquad F = S\,e^{(r-q)\tau}, \qquad \tau = T - t.$$

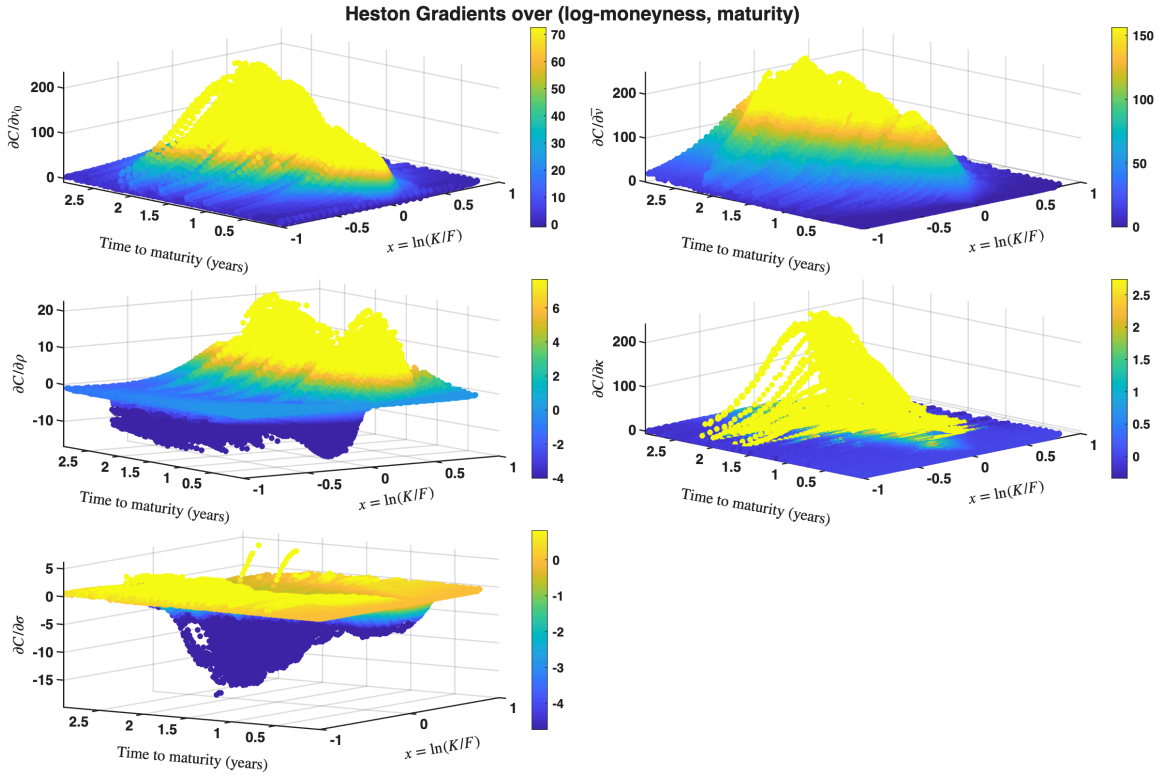


Figure 5.13: Gradients of option prices with respect to Heston parameters under stable market conditions (2018–2024).

Negative log-moneyness refers to strikes above the forward, and positive log-moneyness refers to strikes below the forward. Figure 5.13 shows that the gradients with respect to v_0, $\bar{v}$, and κ are the largest in magnitude. Both $\partial C/\partial v_0$ and $\partial C/\partial \bar{v}$ reach values above 200, whereas the mean-reversion gradient $\partial C/\partial \kappa$ exhibits only localised peaks with limited influence compared with the variance components. The influence of the initial variance v_0 is strongest at short maturities and rises further as τ increases, reaching a peak around 2.5 years. Beyond this horizon, its effect gradually decays because mean reversion reduces the memory of initial conditions and the variance process converges toward the long-run mean. In contrast, the gradient with respect to $\bar{v}$ does not decay in the same way. Its impact strengthens with maturity because the long-run variance anchors the variance path over extended horizons and therefore $\partial C/\partial \bar{v}$ remains persistently large even after the influence of v_0 has diminished. The mean-reversion rate κ controls how quickly the variance moves from v_0 to

$\bar{v}$, so its gradient $\partial C/\partial\kappa$ peaks at intermediate maturities when this growth is most active and then diminishes once the effect of v_0 has faded and the process is driven by $\bar{v}$.

The volatility-of-volatility gradient $\partial C/\partial\sigma$ is asymmetric across log-moneyness. On the left side ($x < 0$ in-the-money (ITM) call), it is slightly positive because greater volatility of volatility increases the chance of large upward moves that add value to these options. At the money (ATM), the gradient turns negative, as with an equity-like negative correlation, higher σ shifts the probability mass toward the tails and reduces the relative value of near-the-money contracts. On the right side ($x > 0$ out-of-the-money (OTM) call), the gradient decays toward zero because additional variance shocks have little impact when strike prices are far above the forward. Across maturities, the effect of σ grows at first, but after about 2.5 years, mean reversion pulls the variance process back to its long-run mean $\bar{v}$, which diminishes the sensitivity of option prices to σ.

Remark (Moneyness Conventions). Throughout this thesis, moneyness is expressed in terms of log-moneyness $x = \ln(K/F)$, where F denotes the forward price. The interpretation of in-the-money (ITM), at-the-money (ATM) and out-of-the-money (OTM) depends on the option type:

- Call: ITM when $x < 0$ (strike below forward). Deep ITM corresponds to $x \ll 0$, while OTM means $x > 0$ (strike above forward), with deep OTM as $x \gg 0$.
- ATM (calls and puts): Defined consistently as $x \approx 0$, i.e. when the strike is close to the forward.
- Put: ITM when $x > 0$ (strike above forward). Deep ITM corresponds to $x \gg 0$, while OTM means $x < 0$, with deep OTM as $x \ll 0$.

Thus, the *left side* of the log-moneyness axis ($x < 0$) corresponds to ITM calls and OTM puts, while the *right side* ($x > 0$) corresponds to OTM calls and ITM puts.

The gradient with respect to the correlation parameter, $\partial C/\partial\rho$, is small in magnitude compared to the variance-related gradients but exhibits a clear sign structure across log-moneyness. The effect arises because correlation shapes how volatility is distributed between right-tail and left-tail outcomes. For in-the-money put (positive log-moneyness, $x > 0$), a stronger correlation increases the likelihood of volatility reinforcing right-tail scenarios, which raises the option's value and produces a positive gradient. For in-the-money calls (negative log-moneyness, $x < 0$), the same adjustment reduces the probability of large left-tail scenarios, so the option value declines

and the gradient becomes negative. At the money (ATM) ($x \approx 0$), the skew adjustment has little impact and the gradient is close to zero.

We introduce the notation `FitWellGradient_i` for the analytical derivatives of the Heston option price with respect to the parameter vector $\boldsymbol{\theta} = (v_0, \bar{v}, \rho, \kappa, \sigma)$. This labelling matches the convention used in the appendix tables, enabling direct cross-reference between text and reported results.

The correlation structure among parameter sensitivities plays a central role in assessing calibration reliability. As noted by Andersen and Piterbarg (2010, Chapter 16), stable and interpretable correlations indicate a well-posed estimation problem, whereas erratic or inconsistent patterns reveal potential instability or model misspecification.

Table 5.5 reports pairwise correlations among gradients under stable and unstable market regimes. In the stable regime, economically coherent patterns emerge. For instance, $\partial C/\partial v_0$ and $\partial C/\partial \kappa$ are positively correlated (`corr` ≈ 0.54), reflecting the joint role of initial variance and mean reversion, while $\partial C/\partial v_0$ and $\partial C/\partial \sigma$ are negatively correlated (`corr` ≈ -0.77), consistent with the offsetting effect of volatility-of-volatility on variance shocks. By contrast, in the unstable regime, correlations fall toward zero or switch sign, such as between $\partial C/\partial v_0$ and $\partial C/\partial \bar{v}$ (`corr` ≈ -0.06), indicating that parameter influences become poorly identified under stress.

The evidence confirms that BPCM filtering preserves economically meaningful parameter relationships in stable conditions, while unstable markets induce widespread decorrelation. These gradient correlations thus provide a transparent diagnostic of calibration robustness across regimes. For comprehensive results, see Table F.1 in Appendix F.

Table 5.5: Pairwise gradient correlations under stable and unstable regimes.

Gradient 1	**Gradient 2**	**Stable**	**Unstable**
$\partial C/\partial v_0$	$\partial C/\partial \bar{v}$	$+0.30$	-0.06
$\partial C/\partial v_0$	$\partial C/\partial \rho$	$+0.213$	$+0.11$
$\partial C/\partial v_0$	$\partial C/\partial \kappa$	$+0.54$	$+0.61$
$\partial C/\partial v_0$	$\partial C/\partial \sigma$	-0.77	-0.25
$\partial C/\partial \bar{v}$	$\partial C/\partial \rho$	$+0.30$	$+0.18$
$\partial C/\partial \bar{v}$	$\partial C/\partial \kappa$	-0.02	-0.26
$\partial C/\partial \bar{v}$	$\partial C/\partial \sigma$	-0.50	-0.05
$\partial C/\partial \rho$	$\partial C/\partial \kappa$	$+0.19$	$+0.48$
$\partial C/\partial \rho$	$\partial C/\partial \sigma$	-0.51	-0.72
$\partial C/\partial \kappa$	$\partial C/\partial \sigma$	-0.56	-0.52

5.9 Option Greeks Under Stable and Unstable Regimes

This section examines the five canonical option Greeks: Delta, Gamma, Theta, Vega, and Rho. Each Greek represents the sensitivity of the option price to a specific financial input. Delta measures the option's responsiveness to movements in the underlying asset price, while Gamma represents the curvature of that sensitivity, i.e. the rate of change of Delta. Theta represents the time decay of the option's value as it approaches maturity. Vega indicates the sensitivity of the option price to volatility.

Each Greek is computed using three distinct methods. The first method applies the classical Black–Scholes formula using observed market-implied volatility. The second method also uses the Black–Scholes formula, but with inputs derived from BPCM. This variant is referred to as `FitWellBLS`, where `FitWell` denotes the output of the final BPCM layer and `BLS` refers to the Black–Scholes model. The third method, denoted `FitWellHSTN`, computes the Greeks using the Heston pricing model, with parameters calibrated from BPCM. These three sets are denoted `BLS_i`, `FitWellBLS_i`, and `FitWellHSTN_i`, respectively, where $i = 1, \ldots, 5$ corresponds, respectively, to Delta, Gamma, Theta, Vega, and Rho.

Approximating Vega via Effective Volatility

In the Black–Scholes model, the volatility σ is constant, and the option price

$$C_{\mathrm{BS}} = C_{\mathrm{BS}}(S_0, K, r, T; \sigma)$$

has Vega defined in the usual way as

$$\mathrm{Vega}_{\mathrm{BS}} \coloneqq \frac{\partial C_{\mathrm{BS}}}{\partial \sigma}.$$

In the Heston model, however, there is no single constant volatility. Instead, the variance process

$$\mathrm{d}v_t \;=\; \kappa(\bar{v} - v_t)\,\mathrm{d}t \;+\; \sigma\sqrt{v_t}\,\mathrm{d}W_t,$$

is stochastic. To obtain a Black–Scholes–comparable quantity, we approximate the random path $(v_t)_{t\in[0,T]}$ by a *deterministic proxy*, namely its conditional mean profile

$$\phi(t) \;\coloneqq\; \mathbb{E}[\,v_t\,] \;=\; \bar{v} + (v_0 - \bar{v})e^{-\kappa t}, \qquad t \in [0, T].$$

In particular, we work with the expected *time-averaged variance* over the option's life,

$$\Phi_T \;\coloneqq\; \frac{1}{T}\,\mathbb{E}\!\left[\int_0^T v_t\,\mathrm{d}t\right] \;=\; \frac{1}{T}\int_0^T \phi(t)\,\mathrm{d}t \;=\; \bar{v} + (v_0 - \bar{v})\,A, \quad A \coloneqq \frac{1-e^{-\kappa T}}{\kappa T}.$$

The *effective volatility* is then defined as $\sigma_{\mathrm{EFF}} \coloneqq \sqrt{\Phi_T}$, which serves as the Black–Scholes volatility input in the approximation that follows.

Differentiating with respect to the initial variance v_0, we get

$$\frac{\partial \Phi_T}{\partial v_0} = A, \qquad \frac{\partial \sigma_{\mathrm{EFF}}}{\partial v_0} = \frac{1}{2\sigma_{\mathrm{EFF}}}\,\frac{\partial \Phi_T}{\partial v_0} = \frac{A}{2\sigma_{\mathrm{EFF}}}.$$

Approximating the Heston price by a Black–Scholes price with $\sigma = \sigma_{\mathrm{EFF}}$,

$$C_{\mathrm{Heston}} \;\approx\; C_{\mathrm{BS}}(S_0, K, r, T; \sigma_{\mathrm{EFF}}),$$

the chain rule gives

$$\frac{\partial C}{\partial v_0} \;\approx\; \mathrm{Vega}_{\mathrm{BS}}\,\frac{\partial \sigma_{\mathrm{EFF}}}{\partial v_0} = \mathrm{Vega}_{\mathrm{BS}}\,\frac{A}{2\sigma_{\mathrm{EFF}}}.$$

Rearranging yields the practical expression

$$\boxed{\mathrm{Vega}_{\mathrm{BS}} \;\approx\; \frac{\partial C}{\partial v_0}\,\frac{2\,\sigma_{\mathrm{EFF}}}{A}}$$

Remark. As $T \to 0$ with κ fixed, $A \to 1$ and as $\kappa \to \infty$ with $T > 0$ fixed, $A \to 0$. For small κT, use the stable series expansion $A \sim 1 - \frac{\kappa T}{2}$ to avoid numerical instability. Note also that v_0, $\bar{v}$, and Φ_T are variances, while σ_{EFF} is a volatility.

In practical implementation, all Greek sensitivities are computed analytically, following the derivations presented in Chapter 3. To assess the stability of the Greeks across models and market data, we begin by comparing three benchmarks: the Black–Scholes model with raw and fitted inputs, and the Heston model with fitted inputs. To make these comparisons more robust, we aggregate observations in six–day bins and plot the median values, with medians defined separately by moneyness ($|x| \le 0.05$ for ATM, and the usual OTM/ITM conventions for calls and puts). For Theta, Vega, and Gamma, results are stratified by log-moneyness and, in some cases, by maturity slices. The Rho panel presents median profiles for calls and puts separately. Delta is illustrated through empirical densities on $[-1, 1]$, with calls shown as solid lines and puts as dashed lines. This aggregation reduces distinctive noise and

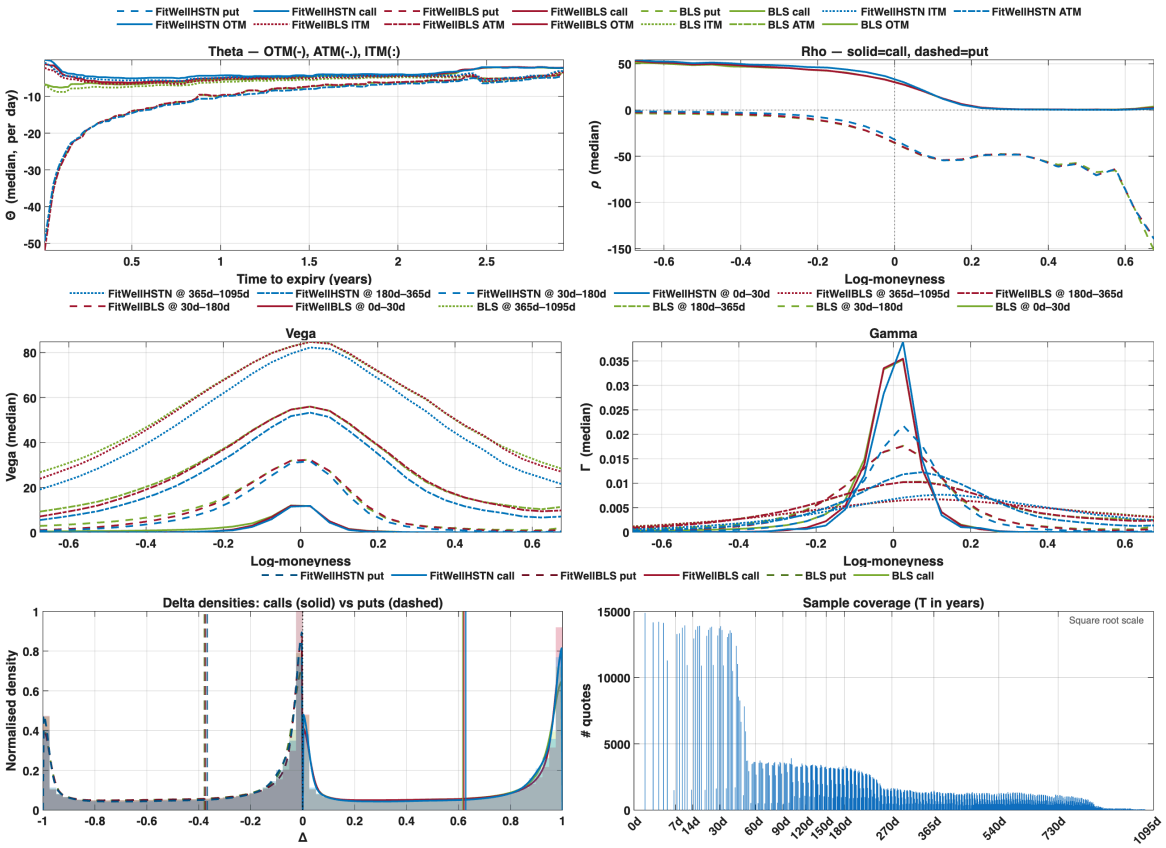


Figure 5.14: Median Greeks under stable conditions: comparison of market-implied values, Black–Scholes fits, and Heston fits.

highlights systematic differences across the three cases.

Figure 5.14 summarises the behaviour of Delta, Theta, Vega, Gamma, and Rho under stable market conditions. The fitted methods produce smooth curves consistent with theoretical expectations, while raw market Greeks often exhibit misshaping, such as negative ITM Theta or Rho tails that extend further into the wings than in the fitted surfaces, particularly for deep ITM puts where values drift below the levels implied by the fitted inputs. These deviations arise from unfiltered market data and are reduced once fitted inputs are applied.

The normalised distribution of Delta (Δ on the horizontal axis; density on the vertical). Solid lines denote calls and dashed lines denote puts. Calls concentrate near $\Delta \approx 1$ (high–sensitivity contracts), whereas puts concentrate near $\Delta \approx 0$. Both sides exhibit a sharp peak near $\Delta \approx 0$, with the concentration being more pronounced on the put side. The mean values (about 0.61 for calls and -0.39 for puts) highlight an asymmetry between the two sides. On the put side, many contracts cluster near $\Delta \approx 0$, reflecting their role as insurance instruments. These values close to zero contribute little to the overall average. On the call side, an influential mass lies near $\Delta \approx 1$, since

deep in-the-money calls are often used as stock alternates.

Both the Black–Scholes and Heston models predict sharply negative ATM Theta as $T \to 0$, while OTM and ITM Thetas converge toward zero for longer maturities. The median fits track these theoretical behaviours: ATM decay dominates at short maturities, whereas OTM and ITM lines remain close to zero. In contrast, the unfiltered market Black–Scholes Thetas sometimes display negative ITM values or excessively steep OTM decay, artefacts of microstructure noise and sparse data that are largely removed once filtering is applied.

All methods exhibit the expected ATM Vega peak and symmetric decay in the wings, consistent with the Black–Scholes scaling with $\sqrt{T}$. For longer maturities, the Heston surface lies below the Black–Scholes curves, with a peak that is slightly broader but of lower height, reflecting the dispersion introduced by stochastic variance. In the wings, the raw market curve departs from the fitted profiles, staying close on the right but rising above them on the left.

Gamma profiles display the expected bell shape, centred at the ATM region, with the sharpest curve for short maturities and flattening as maturity increases. At the money (ATM), all three methods align closely. In the wings, the market curve appears less smooth; the fitted Black–Scholes surface provides a stable intermediate profile, and the Heston surface lies slightly lower, particularly at longer maturities.

Rho follows the Black–Scholes prediction of a monotonic S-shape, increasing for calls and decreasing for puts. The fitted curves reproduce this behaviour, while the market curve shows irregularities in the tails caused by maturity mixing and the influence of deep OTM quotes. Heston Rho remains close to Black–Scholes, with only minor differences introduced through the impact of stochastic variance on the forward. The final panel reports the number of option contracts available in each maturity bucket. The distribution is heavily skewed toward short maturities, with most contracts concentrated between 0 and 60 days, and the number of observations steadily decreases as maturity increases.

In Figure 5.15, we employ the `FitWellBLS` surface as a fixed benchmark to systematically compare it with `BLS` and `FitWellHSTN`. This approach does not presuppose the correctness of any specific method but instead enables a transparent assessment of the degree and structure of disagreement among methodologies. The heatmaps reveal both areas where the methods agree and regions where they diverge.

For Delta, the differences are very small across the surface. Both the market-based `BLS` and the model-based `FitWellHSTN` remain close to the baseline, with only mild discrepancies appearing at longer maturities.

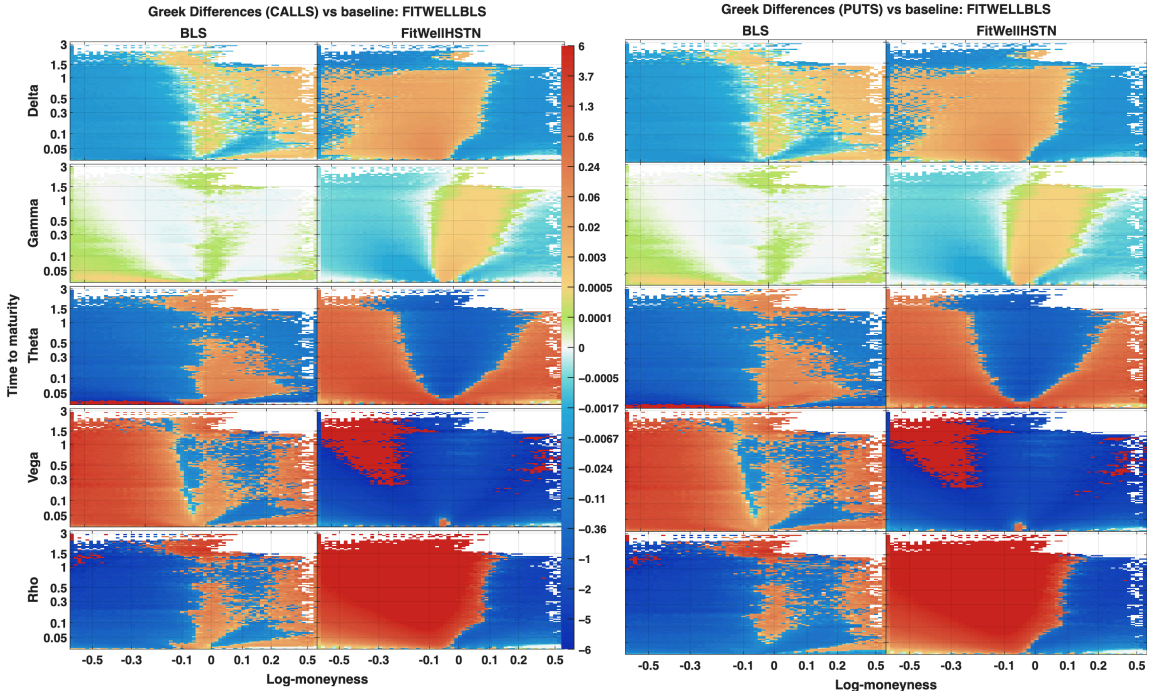


Figure 5.15: Heatmaps of differences in Greeks under alternative estimation methods.

For Gamma, the differences are concentrated at short maturities, especially near ATM, where the sensitivity of second-order derivatives amplifies even small discrepancies in the surfaces. At longer maturities, the differences diminish, and the Heston and Black–Scholes fits remain close to the baseline.

For Theta, the differences form a U-shaped profile with negative values near ATM and positive values in the deep ITM and OTM regions. The effect is strongest at short maturities and weakens as maturity increases, consistent with Theta's heightened sensitivity close to expiry. A similar profile appears in the Gamma panels, suggesting that the recurrence may reflect either the structural features of the option data or calibration effects of the model-fitting procedure.

For Vega, the differences are negative at short maturities across all strikes. As maturity increases, the sign of the negative log-moneyness reverses, but the sign of the positive log-moneyness remains unchanged. This arises from the leverage effect in Heston, where a negative correlation between price and variance shifts the probability mass into the lower tail, maintaining a higher Vega for strikes below the spot price. By contrast, for strikes above the spot, the tail remains lighter, so the adjustment there is minimal. The Black–Scholes framework, due to its inherent symmetry, is unable to capture this asymmetry in volatility sensitivities and consequently yields

systematically lower Vega estimates for strikes below the spot at longer maturities.

For Rho, clear asymmetry appears across moneyness. On the positive log-moneyness, the surface shows blue regions where `FitWellHSTN` produces smaller Rho, and red regions where it produces larger Rho. The mechanism follows from the Black–Scholes expressions

$$\rho_C = KTe^{-rT}N(d_2), \qquad \rho_P = -KTe^{-rT}N(-d_2),$$

with log-moneyness $x = \ln(K/F)$ and $d_2 = (\ln(F/K) - \frac{1}{2}\sigma^2 T)/(\sigma\sqrt{T})$. As with Delta ($\Delta_C = N(d_1)$, $\Delta_P = N(d_1)-1$), Rho exhibits sign asymmetry with $\rho_C \geq 0$ and $\rho_P \leq 0$. Across x, call Rho is largest on the left wing ($x < 0$, ITM) and decays to zero as x increases, while put Rho follows the mirror pattern on the right wing ($x > 0$, ITM). The Black–Scholes fit grows essentially linearly with T and thus tends to overstate tails, whereas the Heston fit, through stochastic volatility and negative spot–volatility correlation, distorts the skew nonlinearly and remains closer to the market, yielding alternating signs in the difference panels.

Collectively, these visualisations and quantitative surfaces localise the limitations of each model and clarify the domains in which market Greeks are unreliable, highlighting regions where refined modelling or smoothing may provide the most impact.

Appendix F.1 compares Black–Scholes and Heston Greeks across the stable and unstable datasets. In the stable case, Delta, Vega and Rho show near-perfect correlation, while Gamma and Theta show moderate alignment. Under unstable conditions, Delta, Vega, Rho, and Theta remain strong, but Gamma collapses, confirming that curvature-based sensitivities are the most fragile under stress. The correlations between Delta and Rho remain consistently strong across both Black–Scholes and Heston estimates (≈ 0.67 in stable periods and ≈ 0.78 in unstable). This indicates that sensitivities to the underlying asset and to the interest rate are structurally well-identified, even under market stress, consistent with the theoretical framework of Heston (1993).

5.10 Correlation Analysis Across Market Regimes

This section presents an empirical analysis of the correlation structures between model calibration outputs and key market observables, distinguishing between stable and unstable market regimes. The primary objective is to validate the effectiveness of BPCM in preserving economically coherent relationships and providing clear diagnostic signals during periods of market turbulence. The full correlation matrix is reported in

Appendix F.1, Table F.1.1, for stable and unstable data with $|\texttt{corr}| \geq 0.30$; weaker associations are omitted to avoid spurious inferences (Aït-Sahalia and Jacod, 2014; Hull, 2018).

The Layer 1 outputs include the implied spot price (`ImpliedUnderlyingPrice`), which is determined through parity constraints, and the arbitrage-free reconstructed option price (`ImpliedOptionPrice`), which is derived under parity conditions. These outputs exhibit near-perfect correlation ($\texttt{corr} \approx 1.00$) with observed market prices in both regimes. These results confirm that BPCM preserves arbitrage relationships even under severe market turbulence, consistent with theoretical expectations from Andersen and Piterbarg (2010).

The correlation patterns in the stable regime are considerably consistent with theoretical expectations. Days to Expiration shows a strong and stable link with Vega (around $\texttt{corr} \approx 0.60$), while the gradient $\partial C/\partial \bar{v}$ (`FitWellGradient_2`) also exhibits a strong correlation (around $\texttt{corr} \approx 0.58$). Extrinsic value correlates more moderately (around $\texttt{corr} \approx 0.48$), reflecting the classical role of maturity in driving option sensitivity. These results align well with option pricing theory, which attributes the time value of options to both volatility and maturity effects (Hull, 2018, Chapter 10). In addition, the gradient with respect to the long-run variance $\bar{v}$ shows an even stronger correlation with Vega, reaching $\texttt{corr} \approx 0.93$ across all methods, which further supports the modelling intuition that Vega is largely driven by the long-run component of volatility in the Heston framework. In contrast, when instability is introduced, theoretical consistency weakens. The correlation between Days to Expiration and Vega falls to more moderate levels, while the link between Days to Expiration and the gradient $\partial C/\partial \bar{v}$ drops sharply to $\texttt{corr} \approx 0.20$, and the correlation with extrinsic value declines to about $\texttt{corr} \approx 0.30$. Likewise, the correlation between extrinsic value and Vega weakens from $\texttt{corr} \approx 0.85$ to about 0.50–0.55, showing that noise disrupts this structural link. This breakdown illustrates why classical risk metrics become unreliable in stressed market conditions. By contrast, gradient-based diagnostics continue to capture systematic misalignment, thereby offering a more robust calibration tool, consistent with sensitivity analyses by Foglia et al. (2009) and robustness concerns raised by Hill and Tiedeman (2007).

By contrast, the second set of findings departs from theoretical intuition. Under instability, Vega–maturity correlations weaken, and the associations for $\partial C/\partial \bar{v}$ and extrinsic value drop sharply, contradicting their expected structural roles. At the same time, `FitWellGradient_3` exhibits strengthened correlations with `FitWellGreeksBS_4` and `GreeksBS_4`, rising from moderate values in the stable regime

(around corr ≈ 0.30) to strong levels in the unstable regime (about corr ≈ 0.77–0.79), while the link with FitWellGreekHSTN_4 increases from corr ≈ 0.35 to corr ≈ 0.44, highlighting the relative robustness of the Heston framework. The correlation with FitWellGradient_5 also deepens (from corr ≈ -0.51 to corr ≈ -0.72), and sensitivity to Days to Expiration grows significantly (from corr ≈ 0.12 to corr ≈ 0.52). Such shifts illustrate how instability distorts maturity-driven relationships and redistributes interactions among gradients and Greeks. This inconsistency with theoretical structures serves as a key marker of model misalignment (Hill and Tiedeman, 2007; Hull, 2018).

In the stable regime, residuals exhibit little to no correlation with option prices, parameters, gradients, or Greeks, indicating that errors are well-distributed and do not interfere with model sensitivities. In contrast, the unstable regime reveals that while residuals remain weakly linked to parameters directly, strong correlations emerge with gradients and Greeks, showing that residual noise is absorbed into the sensitivity layer of the model. This shift demonstrates that stable calibrations preserve the independence of residuals, whereas unstable calibrations redistribute errors into gradients and Greeks, undermining interpretability and robustness. Future work may require a more structured calibration framework that reduces boundary effects and systematically prevents residual noise from propagating into sensitivities, thereby preserving their interpretability.

Market observables, such as Strike, Volume, and OpenInterest, exhibit varying degrees of correlation with model outputs. Intrinsic value continues to correlate strongly with ImpliedOptionPrice (corr ≥ 0.99), underscoring the successful enforcement of no-arbitrage conditions even during market instability.

The comparison between regimes shows that significant correlations ($|\text{corr}| \geq 0.30$) are much fewer in the stable case (178) than in the unstable case (333). This simpler structure in the stable regime is preferable, as it highlights a smaller set of strong and economically meaningful links. By contrast, the unstable regime generates numerous additional connections, most of which are weak and driven by noise, making the results more difficult to interpret. The BPCM framework addresses this issue by using adaptive filtering to isolate the latent signal from noise, ensuring interpretability and aligning with the need for diagnostic refinement as emphasised by Foglia et al. (2009).

Under classical theoretical expectations (see Hull, 2018), Delta should be positively correlated with the underlying asset price. Empirically, the adjusted Black–Scholes Delta exhibits a positive correlation with market price (corr $\approx +0.35$) in the stable regime, aligning with theory. However, in the unstable regime, this turns into a strong

negative correlation ($\texttt{corr} \approx -0.83$), indicating a breakdown in price sensitivity and a divergence from classical behaviour. The Heston Delta similarly aligns with market price under stability ($\texttt{corr} \approx +0.33$) but not under stress ($\texttt{corr} \approx -0.83$), highlighting the insufficiency of stochastic volatility modelling without diagnostic layering. Intrinsic value and all Delta show the expected positive correlation ($\texttt{corr} \approx +0.31$), consistent with moneyness logic. Under instability, this relationship reverses sharply ($\texttt{corr} \approx -0.85$), indicating that pricing no longer reflects intrinsic characteristics and further affirming BPCM's role in identifying misalignment.

These findings collectively demonstrate that BPCM not only calibrates effectively under normal market conditions but also preserves economically coherent relationships in unstable or noisy environments. This robustness suggests its potential applicability in high-frequency contexts where microstructure noise, including bid–ask bounce, latency effects such as delays in quote dissemination or recording, and timing mismatches in quotes, can distort observed prices (Aït-Sahalia and Jacod, 2014). By constructing bounded pricing surfaces that enforce key financial constraints, such as put–call parity and bid–ask bounds, BPCM maintains internal consistency even when quotes are sparse or irregular. This structure suppresses spurious variance and autocorrelation patterns often observed at high sampling frequencies, providing a model-based filtering approach that aligns with the principles of structural econometrics. Its flexibility under adverse data conditions positions BPCM as a promising candidate for further investigation in high-frequency model calibration, a direction discussed in the conclusion.

Chapter 6

Conclusion and Future Research Directions

6.1 Summary and Contextualisation of Results

This thesis addressed the following research question: how can we robustly calibrate the Heston stochastic volatility model under realistic market conditions characterised by data imperfections? In addressing the question, this thesis aimed to bridge the gap between theoretical model calibration and practical financial markets and provide a sophisticated, practical calibration method. The BPCM was developed with a layered structure capable of robustly fitting the Heston stochastic volatility model to large option datasets. The method comprises three core layers whose outputs can be further refined by additional calibration layers. Layer 1 cleans the market data and extracts reliable pricing signals. This layer considers each option's quotes, strike price, time to maturity, risk-free interest rate, and dividend yield and aims to establish put–call parity. The estimated option data are then used to calibrate the Heston model in Layer 2, which outputs a set of calibrated parameters for each trading day. Layer 3 interprets the residuals from Layer 2 and redistributes them across identified noisy inputs to improve parameter calibration. In this layer, parameter adjustments are only allowed within small bounds to ensure the parameters remain economically reasonable. Layer 3 normally provides a robust, reliable calibration; in unusual market conditions or cases of exceptionally noisy data, BPCM initiates additional refinement iterations to reduce the residuals below a specified threshold.

This thesis contributes to the calibration literature by aligning the BPCM methodology with established findings in high-frequency financial econometrics. Rather than relying on raw, potentially noisy market quotes, BPCM incorporates structural filtering, put–call parity enforcement, and bounded optimisation to mitigate the effects of microstructure noise. These design choices directly address the concerns raised in

recent econometric studies, ensuring that the resulting model outputs remain both robust and economically interpretable.

BPCM also circumvents the limitations of assuming the existence of a singular, definitive market price. Rather than enforcing rigid adherence to the observed mid or last price, BPCM produces internally coherent pricing snapshots that satisfy the model's economic and structural constraints. Additionally, as BPCM classifies quotes based on structural coherence, it retains even unstable observations within controlled bounds. This allows the model to maximise information retention and deliver accurate calibration based on real-world data. This approach distinguishes BPCM from traditional heuristic or rule-based filtering techniques, which tend to inadvertently exclude valuable information during data cleaning (Aït-Sahalia and Jacod, 2014).

6.2 Novel Contributions

This research makes four key contributions to the literature and practices related to Heston model calibration. First, the balanced premium calibration method is a methodological innovation that improves Heston calibration stability, interpretability, and computational efficiency. This novel framework is based on the systematic integration of data quality assurance directly into the calibration process. Unlike traditional approaches, which separate data cleaning from parameter estimation, BPCM combines data filtering, preliminary calibration, and joint optimisation into a coherent three-layer structure. This approach increases model robustness, significantly reduces calibration errors, and offers improved adherence to arbitrage bounds, particularly during periods of market stress.

This thesis extends existing work on analytical calibration by deriving closed-form expressions for the second-order sensitivities (Hessian matrix) of the Heston model with respect to the model parameters. While first-order analytical gradients were introduced by Cui et al. (2017), the analytical structure developed here allows for the construction of an approximate Hessian using the Jacobian matrix of the model prices. This formulation supports second-order optimisation schemes, which are typically more expensive and susceptible to numerical artefacts. In practice, these artefacts arise because the Hessian includes products of small second-order sensitivities, making it more sensitive to numerical noise and round-off errors.

In addition, the thesis derives closed-form analytical expressions for both the first- and second-order derivatives of the Heston pricing function with respect to time to maturity τ and the risk-free interest rate r. These expressions offer a mathematically

rigorous way to compute the Greeks Theta and Rho, and provide theoretical tools for examining the model's sensitivity to maturity and rate inputs.

Third, this thesis repurposed calibration residuals from error metrics into powerful diagnostic tools that provide insight into market structures and dynamics. The systematic interpretation of residuals can reveal important issues and aspects such as model misspecifications, changes in market regimes, and periods of market stress. This approach fundamentally reframes calibration from a mere parameter fitting exercise into a sophisticated instrument for market analysis and informed decision-making.

Fourth, this thesis evaluated the BPCM framework on a real-world dataset of over 1.5 million Apple option contracts spanning the years 2018 to 2024. For stable market conditions, 91.58% of the model-generated option prices fell within the observed bid–ask spreads, as illustrated in Figure 5.1. Similarly, 95.77% of the underlying asset prices were contained within the corresponding daily high–low ranges, as shown in Figure 5.6. Both figures are based exclusively on stable data periods and confirm that BPCM produces outputs consistent with market-observed bounds, supporting its reliability for pricing and risk management applications under stable market conditions.

6.3 Limitations of BPCM

Despite its demonstrated strengths, BPCM has several limitations. First, the calibration depends significantly on estimated inputs such as the quote time, which may reflect modelling assumptions or data timestamp uncertainty rather than true market recording times. Thus, the resulting calibrated time series may occasionally reflect imposed assumptions rather than market behaviours.

Second, the BPCM boundaries for latent variables such as interest rates, dividend yields, and time adjustments are empirically defined. Without precise intra-day or tick-level data, these boundaries introduce subjective elements that might affect calibration consistency across different market regimes.

Third, BPCM assumes that the Heston model is structurally correct. However, significant market shifts, regime changes, or complexities arising from options with early-exercise features may undermine this core assumption. In such cases, BPCM may require adaptations to maintain its calibration effectiveness.

6.4 Future Research Directions

Future research may focus on improving the BPCM estimation of latent market inputs, particularly interest rates and dividend yields. Employing inter-market data, forward curve estimation, or leveraging cross-asset dynamics could substantially enhance the reliability and economic coherence of these variables. Moreover, as introduced in Chapter 3, first- and second-order sensitivities of option prices are derived with respect to interest rates and time to maturity, as well as second-order derivatives with respect to the Heston model parameters. These include cross-partial terms in the Hessian matrix, which capture local curvature in the calibration surface. Although this structure is not fully explored in the present analysis, it provides a promising direction for future work on higher-order diagnostics and the stability of calibration across regimes.

Moreover, BPCM could be extended to encompass multiple correlated assets or indices. Examining the relationships between individual stocks, exchange-traded funds, and indices could yield significant insights into market dynamics that could then be used to quantify systemic noise and improve calibration robustness.

A key avenue for future research may be the development of BPCM into a dynamic, time-aware calibration framework using state–space modelling techniques such as Kalman filtering. Incorporating the robust, noise-filtered dataset produced by BPCM into dynamic parameter estimation would enable the extraction of smooth and interpretable parameter trajectories. The revised model could rapidly identify and adapt to volatility regime shifts, evolving risk premiums, and market-level structural changes, significantly enhancing both parameter calibration and theoretical understanding.

Finally, extending the BPCM framework to include jumps or rough volatility components could address inherent model misspecifications in extreme market conditions. Another promising direction is to integrate BPCM with finite-difference solvers, such as ADI schemes, which are efficient and stable for multi-dimensional option-pricing PDEs. For example, in 't Hout and Foulon (2010) established ADI efficiency for the two-dimensional Heston PDE with correlation, while Haentjens and in 't Hout (2012) extended these methods to the three-dimensional Heston–Hull–White PDE with stochastic interest rates. This would broaden BPCM's applicability from vanilla options to a wider class of derivatives while maintaining robustness in calibration.

Bibliography

Aït-Sahalia, Y. and Jacod, J. (2014). *High-Frequency Financial Econometrics*. Princeton University Press.

Aït-Sahalia, Y. and Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53(2):499–547.

Aït-Sahalia, Y., Mykland, P. A., and Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies*, 18(2):351–416.

Albrecher, H., Mayer, P., Schoutens, W., and Tistaert, J. (2006). The little Heston trap. Working paper. Linz and Graz University of Technology.

Alòs, E., De Santiago, R., and Vives, J. (2015). Calibration of stochastic volatility models via second order approximation: The Heston model case. *International Journal of Theoretical and Applied Finance*, 18(6):1550036.

Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance*, 11(3):1–42.

Andersen, L. B. G. and Piterbarg, V. V. (2010). *Interest Rate Modeling, Volume III: Products and Risk Management*, volume 3. Atlantic Financial Press. See Chapter 16.

Andreasen, J. and Huge, B. N. (2010). Volatility interpolation: Arbitrage-free interpolation of option prices. Technical Report 1694972, SSRN Electronic Journal. Available at `https://ssrn.com/abstract=1694972`, revised Feb 2011.

Andres, H., Arrouy, P., Bonnefoy, P., Boumezoued, A., and Mehalla, S. (2020). Fast calibration of the LIBOR market model with stochastic volatility based on analytical gradient. arXiv preprint. arXiv:2006.13521.

Apple Inc. (2020a). Apple announces Mac transition to Apple Silicon. Available at: https://www.apple.com/newsroom/2020/06/apple-announces-mac-transition-to-apple-silicon/ (accessed 23 June 2025).

Apple Inc. (2020b). Apple reports record first quarter results. Available at: https://www.apple.com/uk/newsroom/2020/01/apple-reports-record-first-quarter-results/ (accessed 23 June 2025).

Apple Inc. (2020c). Apple reports second quarter results. Technical report, Apple Newsroom. Fiscal Q2 2020 results published April 30, 2020. Available at: https://www.apple.com/newsroom/2020/04/apple-reports-second-quarter-results/ (accessed June 23, 2025).

Apple Inc. (2020d). Apple unveils new iPad Pro with LiDAR and trackpad support. Technical report, Apple Newsroom. Announces 4th-generation iPad Pro with A12Z chip, LiDAR Scanner, and new Magic Keyboard support (accessed June 23, 2025).

Apple Inc. (2020e). iPhone SE: A powerful new smartphone in a popular design. Technical report, Apple Newsroom. Product launch press release published April 15, 2020. Available at: https://www.apple.com/newsroom/2020/04/iphone-se-a-powerful-new-smartphone-in-a-popular-design/ (accessed June 23, 2025).

Apple Inc. (2020f). New MacBook Air has more to love and is now just $999. Technical report, Apple Newsroom. Announces updated MacBook Air with Magic Keyboard, twice the storage, and improved price (accessed June 23, 2025).

Apple Inc. (2021a). Apple introduces new developer tools and technologies to create even better apps. Technical report, Apple Newsroom. Announcing developer betas for iOS 15, iPadOS 15, macOS Monterey, watchOS 8 and tvOS 15 (accessed June 23, 2025).

Apple Inc. (2021b). Apple reports second quarter results. Available at: https://www.apple.com/newsroom/2021/04/apple-reports-second-quarter-results/ (accessed 23 June 2025).

Apple Inc. (2021c). Apple unveils game-changing MacBook Pro powered by M1 Pro and M1 Max chips. Technical report, Apple Newsroom. Available at: https:

//www.apple.com/uk/newsroom/2021/10/apple-unveils-game-changing-macbook-pro/ (accessed June 23, 2025).

Apple Inc. (2022a). Apple reports first quarter results. Technical report, Apple Newsroom. Apple posted revenue of USD 123.9 B — a record quarter even under supply chain constraints (accessed June 23, 2025).

Apple Inc. (2022b). Apple unveils all-new MacBook Air supercharged by the M2 chip. Technical report, Apple Newsroom. Launch announcement of redesigned M2 MacBook Air (accessed June 23, 2025).

Apple Inc. (2022c). Apple's self service repair now available. Technical report, Apple Newsroom. Press release announcing Self Service Repair availability (accessed June 23, 2025).

Apple Inc. (2023a). Apple introduces Apple Pay Later. Available at: https://www.apple.com/newsroom/2023/03/apple-introduces-apple-pay-later/ (accessed 23 June 2025).

Apple Inc. (2023b). Apple unveils M3, M3 Pro, and M3 Max chips for next generation of Mac. Technical report, Apple Newsroom. Official press release announcing new M3 chip family and MacBook Pro/iMac updates (accessed June 23, 2025).

Apple Inc. (2023c). WWDC 2023 Highlights: Apple unveils Vision Pro and updates across platforms. Available at https://www.apple.com/newsroom/2023/06/wwdc23-highlights/ (accessed June 23, 2025).

Apple Inc. (2024a). Apple introduces M4 chip. Technical report, Apple Newsroom. Official press release unveiling the M4 chip in the new iPad Pro (accessed June 23, 2025).

Apple Inc. (2024b). Apple reports second quarter results. Available at: https://www.apple.com/uk/newsroom/2024/05/apple-reports-second-quarter-results/ (accessed 23 June 2025).

Apple Inc. (2025). Dividend history. Available at: https://investor.apple.com/dividend-history/default.aspx (accessed 23 June 2025).

Axios (2021). Apple's quarterly sales top $100 billion for first time. *Axios*. Apple posted $111B in revenue in Q1 2021, marking its first-ever $100B quarter (accessed June 23, 2025).

Bain, A., Crisan, D., and Laachir, I. (2019). Calibration of local stochastic volatility models with stochastic rates and credit spreads. *arXiv preprint arXiv:1911.00877*.

Balu, N. and Randewich, N. (2023). Apple's stock market value falls below $2 trillion. *Reuters*. Shares fell 3.7% following downgrade by BNP Paribas; market value dropped below $2 trillion for the first time since March 2021 (accessed June 23, 2025).

Bauer, J. (2017). Model calibration in thinly traded derivatives markets. *SSRN Electronic Journal*, 2017(10):1–31. Available at SSRN: `https://ssrn.com/abstract=2849497`.

BBC News (2020). Apple and Google release marks 'watershed moment' for contact tracing. Available at: `https://www.bbc.co.uk/news/technology-52740131` (accessed 23 June 2025).

Bilton, N. (2024). Why Tim Cook is going all in on the Apple Vision Pro. *Vanity Fair*. Includes first public photos of Tim Cook wearing the Vision Pro headset; publication date aligns with early February 2024 launch (accessed June 23, 2025).

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Boness, A. J. (1964). Elements of a theory of stock-option value. *Journal of Political Economy*, 72(2):163–175.

Boyarchenko, S. and Levendorskii, S. (2024). Correct implied volatility shapes and reliable pricing in the rough Heston model. SSRN. Available at: `https://ssrn.com/abstract=5060018`.

Canon, G. and Jolly, J. (2021). Apple sales rise to $90 bn amid COVID buying surge. Available at: `https://www.theguardian.com/technology/2021/apr/28/apple-quarterly-earnings-covid-sales` (accessed 23 June 2025).

Carr, P. and Madan, D. (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2(4):61–73.

Chan, K., Chung, Y. P., and Johnson, H. (1993). Why option prices lag stock prices: A trading-based explanation. *Journal of Finance*, 48(5):1957–1967.

Clevenhaus, A., Totzeck, C., and Ehrhardt, M. (2024). A gradient-based calibration method for the Heston model. *International Journal of Computer Mathematics*, 101(9–10):1094–1112.

CNBC (2020). Wall Street's fear gauge hits highest level ever. Available at: `https://www.cnbc.com/2020/03/16/wall-streets-fear-gauge-hits-highest-level-ever.html` (accessed 23 June 2025).

Coleman, T. F., Li, Y., and Verma, A. (1999). Reconstructing the term structure of volatilities from option prices. *Journal of Computational Finance*, 2(3):77–116.

Conn, A. R., Gould, N. I. M., and Toint, P. L. (2000). *Trust Region Methods*. SIAM.

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236.

Cont, R. and Tankov, P. (2004). *Financial Modelling with Jump Processes*. Chapman and Hall, London.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.

Cui, Y., del Baño Rollin, S., and Germano, G. (2017). Full and fast calibration of the Heston stochastic volatility model. *European Journal of Operational Research*, 263(2):625–638.

del Baño Rollin, S., Ferreiro-Castilla, A., and Utzet, F. (2010). On the density of log-spot in the Heston volatility model. *Stochastic Processes and their Applications*, 120(10):2037–2063.

Elliott, D. and Johnston, P. R. (2008). Gauss–Legendre quadrature for the evaluation of integrals involving the Hankel function. *Journal of Computational and Applied Mathematics*, 211(1):23–35.

Etherington, D. (2020). Apple and Google launch exposure notification API. Available at: `https://techcrunch.com/2020/05/20/apple-and-google-launch-exposure-notification-api-enabling-public-health-authorities-to-release-apps/` (accessed 23 June 2025).

Evensen, G., Vossepoel, F., and van Leeuwen, P. J. (2022). *Data Assimilation Fundamentals: A Unified Formulation of State and Parameter Estimation*. Springer, Cham.

Fengler, M. R. (2009). Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance*, 9(4):417–428.

Foglia, L., Hill, M. C., Mehl, S. W., Burlando, P., Werner, A. D., and Stauffer, F. (2009). Sensitivity analysis, calibration and testing of a distributed hydrological model using error-based weighting and multiple performance measures. *Water Resources Research*, 45(6).

Gaetano, C. (2020). Coronavirus fear spooking markets. Available at: `https://www.nysscpa.org/news/publications/the-trusted-professional/article/coronavirus-fear-spooking-markets-012720` (accessed 23 June 2025).

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. John Wiley & Sons, Inc.

Gibbs, S. (2021). Apple launches new AirPods and revamped MacBook Pro laptops. Available at: `https://www.theguardian.com/technology/2021/oct/18/apple-launches-new-airpods-and-revamped-macbook-pro-laptops` (accessed 23 June 2025).

Gilli, M. and Schumann, E. (2013). Calibrating option pricing models with heuristics. In Brabazon, A., O'Neill, M., and Maringer, D., editors, *Natural Computing in Computational Finance: Volume 4*, pages 9–37. Springer Berlin Heidelberg.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Applications of Mathematics. Springer.

Gradojevic, N. and Kukolj, D. (2024). Unlocking the black box: Non-parametric option pricing before and during COVID-19. *Annals of Operations Research*, 334:59–82.

Gudmundsson, H. and Vyncke, D. (2019). On the calibration of the 3/2 model. *European Journal of Operational Research*, 276(3):1178–1192.

Gupta, R. (2025). Computational insights into optimal household portfolio decisions: A stochastic approach with Heston model and finite difference scheme. *International Journal of Information Technology*, 17:677–688.

Guterding, D. and Boenkost, W. (2018). The Heston stochastic volatility model with piecewise constant parameters – efficient calibration and pricing of window barrier options. *Journal of Computational and Applied Mathematics*, 343:353–362.

Haentjens, T. and in 't Hout, K. J. (2012). Alternating direction implicit finite difference schemes for the Heston–Hull–White PDE. *The Journal of Computational Finance*, 16(1):83–110. Efficient ADI schemes for 3D Heston–Hull–White pricing.

Hagströmer, B. (2019). Bias in the effective bid–ask spread. Working paper.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343.

Hill, M. C. and Tiedeman, C. R. (2007). *Effective Groundwater Model Calibration: With Analysis of Data, Sensitivities, Predictions, and Uncertainty*. John Wiley & Sons.

Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2):281–300.

Hull, J. C. (2018). *Options, Futures, and Other Derivatives*. Pearson Education Limited, Harlow, England, 9th edition.

in 't Hout, K. J. and Foulon, S. (2010). ADI finite difference schemes for option pricing in the Heston model with correlation. *International Journal of Numerical Analysis and Modeling*, 7(2):303–320.

Ivanovas, A. (2015). *Option Data, Missing Tails, and the Intraday Variation of Implied Moments*. PhD thesis, University of St. Gallen.

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2):433–451.

Jacquier, A. and Shi, F. (2019). The randomized Heston model. *SIAM Journal on Financial Mathematics*, 10(1):89–129.

Kahl, C. and Jäckel, P. (2006). Not-so-complex logarithms in the Heston model. *Wilmott Magazine*, pages 74–103.

Krauskopf, L. and Davies, M. (2022). Massive turnaround for stocks puts traders on alert for more volatility. *Reuters*. Discusses heightened volatility and hedging activity following rapid market moves (accessed June 23, 2025).

Lagnado, R. and Osher, S. (1997). A technique for calibrating derivative security pricing models to market data. In *Computational Finance*.

Lee, R. W. (2005). Option pricing by transform methods: Extensions, unification, and error control. *Journal of Computational Finance*, 7(3):51–86.

Levenberg, K. (1944). A method for the solution of certain problems in least squares. *Quarterly Applied Mathematics*, 2:164–168.

Lord, R., Koekkoek, R., and van Dijk, D. (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2):177–194.

Macrotrends (2025). Apple – 45 year stock price history | AAPL. Available at: `https://www.macrotrends.net/stocks/charts/AAPL/apple/stock-price-history` (accessed 23 June 2025).

Manaster, S. and Rendleman, R. J. (1982). Option prices as predictors of equilibrium stock prices. *Journal of Finance*, 37(4):1043–1057.

Marcato, G. and Sebehela, T. (2022). The paradoxical prices of options. *Review of Pacific Basin Financial Markets and Policies*, 25(2):2250009.

Marquardt, D. (1963). An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11:431–441.

Meier, P. (2015). *Essays on Pricing Kernel Estimation, Option Data Filtering and Risk-Neutral Density Tail Estimation*. PhD thesis, University of St. Gallen.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183.

Mo, Y. and Goh, B. (2024). Apple iPhone sales in China fall 30% in first week of 2024 – Jefferies. *Reuters*. Reports 30% year-on-year drop in iPhone sales during first week of January 2024 (accessed June 23, 2025).

Moré, J. J. (1978). The Levenberg–Marquardt algorithm: Implementation and theory. In Watson, G. A., editor, *Numerical Analysis*, volume 630 of *Lecture Notes in Mathematics*, pages 105–116. Springer.

Murray, P. (2020). The 5 most exciting things Apple announced at WWDC 2020. Available at: `https://time.com/5857500/apple-wwdc-2020/` (accessed 23 June 2025).

Neely, A. (2023). Apple to hold first in-person summit in years, focused on AI. *AppleInsider*. Reports Apple's internal AI summit held at the Steve Jobs Theater, published February 6, 2023 (accessed June 23, 2025).

Nelis, J. (2022). Detecting and repairing arbitrage: From European to American options. MSc thesis, Imperial College London.

Nellis, S. (2022). Apple results top estimates as iPhone demand escapes economic slump. *Reuters*. Reports Apple stock rising 3.5% as CFO notes easing supply shortages and strong iPhone demand (accessed June 23, 2025).

Nellis, S. (2023a). Apple forecasts another drop in revenue, proclaims iPhone production problems over. *Reuters*. Reports Apple's Q1 FY23 revenue drop and cites iPhone supply chain recovery in China following COVID disruptions (accessed June 23, 2025).

Nellis, S. (2023b). Apple profit tops Wall Street targets, strong services counter weaker iPhone. *Reuters*. Accessed: 23 June 2025. Reports $81.8 billion revenue for fiscal Q3.

Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer, 2nd edition.

Reuters (2022). Apple rolls out self-service repair program for iPhones, with Mac support planned. *Reuters*. Reporting date matching the April 2022 launch announcement (accessed June 23, 2025).

Reuters (2023). Apple says its new high-end handset models are facing supply constraints. *Reuters*. Cook confirmed that iPhone 15 Pro and Pro Max models were still supply-constrained more than a month after launch (accessed June 23, 2025).

Ribeiro, A. and Poulsen, R. (2013). Approximation behoves calibration. *Quantitative Finance Letters*, 1(1):36–40.

Rossignol, J. (2022). Apple announces new iPad Pro with M2 chip and other new features. *MacRumors*. Announces iPad Pro (M2) with Apple Pencil hover, Wi-Fi 6E, and ProRes video support (accessed June 23, 2025).

Rouah, F. D. (2013). *The Heston Model and its Extensions in MATLAB and C#*. Wiley, Hoboken, NJ.

Ruf, J. and Wang, W. (2020). Neural networks for option pricing and hedging: A literature review. *Journal of Computational Finance*, 24(1):1–46.

Scott, L. O. (1987). Option pricing when the variance changes randomly: Theory, estimation, and an application. *The Journal of Financial and Quantitative Analysis*, 22(4):419–438.

Sepp, A. (2008). Pricing options on realized variance in the Heston model with jumps in returns and volatility. *Journal of Computational Finance*, 11(4):33–70.

Sepp, A. (2012). Pricing options on realized variance in the Heston model with jumps in returns and volatility – Part II: An approximate distribution of discrete variance. *Journal of Computational Finance*, 16(2):3–32.

Shen, J. (2020). *Direct Inversion for the Heston Model*. PhD thesis, Heriot-Watt University, Edinburgh, United Kingdom.

Slivka, E. (2021). Apple reports record-setting Q1 2021 results: $28.8b profit on $111.4b revenue. *MacRumors*. Covers Apple's record Q1 FY21 earnings driven by iPhone 12 demand (accessed June 23, 2025).

Sridi, A. and Bilokon, P. (2023). Applying Deep Learning to calibrate Stochastic Volatility models. arXiv preprint: `https://arxiv.org/abs/2309.07843v2`.

Stephan, J. A. and Whaley, R. E. (1990). Intraday price change and trading volume relations in the stock and stock option markets. *Journal of Finance*, 45(1):191–220.

Su, T. and Guo, W. (2006). Option put–call parity relations when the underlying security pays dividends. *International Journal of Business and Economics*, 5(3):225–230.

The Verge (2022). Apple made more money than ever, even with the supply chain crunch. *The Verge*. Highlights Apple estimated losing USD 6 B in revenue due to supply constraints, signaling inflation and infrastructure pressure (accessed June 23, 2025).

The Verge (2023). Apple 'scary fast' Mac launch event: The 4 biggest announcements. *The Verge*. Covers introduction of M3-powered MacBook Pro and iMac at Apple's Oct 30 event (accessed June 23, 2025).

Wacker, K. (2021). Apple news roundup: July 2021. Available at: `https://eshop.macsales.com/blog/76893-apple-news-roundup-july-2021/` (accessed 23 June 2025).

Wang, Y. and Bai, Y. (2025). European option pricing under the approximate fractional Heston jump–diffusion model with a stochastic long-term mean. *Alexandria Engineering Journal*, 123(3):145–156.

Wiggins, J. B. (1987). Option values under stochastic volatility: Theory and empirical estimates. *Journal of Financial Economics*, 19(2):351–372.

Wired Staff (2023). Everything Apple announced at WWDC 2023. *Wired*. Coverage includes Vision Pro launch, macOS Sonoma, iOS 17 features.

World Economic Forum (2020). Mad march: How the stock market is being hit by COVID-19. Available at: `https://www.weforum.org/stories/2020/03/stock-market-volatility-coronavirus/` (accessed 23 June 2025).

World Health Organization (2020). Statement on the second meeting of the International Health Regulations (2005) emergency committee regarding the outbreak of novel coronavirus (2019-nCoV). Accessed: 1 August 2025.

Zhang, C. (2023). Calibrating the Heston model with deep differential networks. *Preprint*. Available on request or institutional repository.

Appendix A

Derivation of the Heston Option Pricing PDE

This appendix systematically derives the fundamental partial differential equation (PDE) of the Heston stochastic volatility model, an essential cornerstone for pricing options within a stochastic volatility framework. Beginning with the specification of the model under the risk-neutral measure, we rigorously apply Itô's Lemma to derive the PDE governing option prices. We then simplify the resulting PDE through a log-price transformation, providing a tractable form crucial for subsequent analytic and numerical analyses. The derivation is clearly articulated, conforming to academic standards for mathematical rigor and precision.

A.1 Application of Itô's Lemma to the Heston Model

The Heston model describes the dynamics of the asset price S_t and the variance V_t using two coupled stochastic differential Equations (2.1). Under this model, the Brownian motions $\mathrm{d}W_t^s$ and $\mathrm{d}W_t^v$ are correlated with correlation ρ, so:

$$\mathrm{d}W_t^s \cdot \mathrm{d}W_t^v = \rho \, \mathrm{d}t.$$

Apply Itô's Lemma to a twice differentiable function $u(S_t, V_t, t)$, where u depends on the time t, the stock price S_t, and the variance V_t:

$$\mathrm{d}u = \frac{\partial u}{\partial t}\mathrm{d}t + \frac{\partial u}{\partial S}\,\mathrm{d}S_t + \frac{\partial u}{\partial V}\,\mathrm{d}V_t + \frac{1}{2}\frac{\partial^2 u}{\partial S^2}\,(\mathrm{d}S_t)^2 + \frac{1}{2}\frac{\partial^2 u}{\partial V^2}\,(\mathrm{d}V_t)^2 + \frac{\partial^2 u}{\partial S \partial V}\,\mathrm{d}S_t\,\mathrm{d}V_t. \tag{A.1.1}$$

Substituting the dynamics of S_t and V_t from Equations (2.1) into the first-order terms, we obtain:

$$\frac{\partial u}{\partial S}\,\mathrm{d}S_t = \frac{\partial u}{\partial S}\left(\mu S_t \mathrm{d}t + \sqrt{V_t} S_t\,\mathrm{d}W_t^s\right), \tag{A.1.2}$$

$$\frac{\partial u}{\partial V}\,\mathrm{d}V_t = \frac{\partial u}{\partial V}\left(\kappa(\bar{v} - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_t^v\right). \tag{A.1.3}$$

Next, compute the quadratic terms $(\mathrm{d}S_t)^2$, $(\mathrm{d}V_t)^2$, and $\mathrm{d}S_t\,\mathrm{d}V_t$.

- For $(\mathrm{d}S_t)^2$:

$$\begin{aligned}(\mathrm{d}S_t)^2 &= \left(\mu S_t \mathrm{d}t + \sqrt{V_t} S_t\,\mathrm{d}W_t^s\right)^2,\\ (\mathrm{d}S_t)^2 &= (\mu S_t \mathrm{d}t)^2 + 2(\mu S_t \mathrm{d}t)(\sqrt{V_t} S_t\,\mathrm{d}W_t^s) + (\sqrt{V_t} S_t\,\mathrm{d}W_t^s)^2.\end{aligned}$$

- For $(\mathrm{d}V_t)^2$:

$$\begin{aligned}(\mathrm{d}V_t)^2 &= \left(\kappa(\bar{v} - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_t^v\right)^2,\\ (\mathrm{d}V_t)^2 &= (\kappa(\bar{v} - V_t)\mathrm{d}t)^2 + 2(\kappa(\bar{v} - V_t)\mathrm{d}t)(\sigma\sqrt{V_t}\,\mathrm{d}W_t^v) + (\sigma\sqrt{V_t}\,\mathrm{d}W_t^v)^2.\end{aligned}$$

- For $\mathrm{d}S_t\,\mathrm{d}V_t$:

$$\begin{aligned}\mathrm{d}S_t\,\mathrm{d}V_t &= \left(\mu S_t \mathrm{d}t + \sqrt{V_t} S_t\,\mathrm{d}W_t^s\right)\left(\kappa(\bar{v} - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_t^v\right),\\ \mathrm{d}S_t\,\mathrm{d}V_t &= \mu S_t\kappa(\bar{v} - V_t)(\mathrm{d}t)^2 + \mu S_t\sigma\sqrt{V_t}\mathrm{d}t\,\mathrm{d}W_t^v\\ &\quad + \sqrt{V_t} S_t\kappa(\bar{v} - V_t)\,\mathrm{d}W_t^s\mathrm{d}t + \sqrt{V_t} S_t\sigma\sqrt{V_t}\,\mathrm{d}W_t^s\,\mathrm{d}W_t^v.\end{aligned}$$

We classify the stochastic differential terms by their asymptotic order to justify which terms are kept and which are eliminated when applying Itô's Lemma.

- $(\mathrm{d}t)^2 = o(\mathrm{d}t)$, so its contribution is negligible as $\mathrm{d}t \to 0$.
- $\mathrm{d}t \cdot \mathrm{d}W_t^s = \mathcal{O}(\mathrm{d}t^{3/2})$, which also vanishes relative to dt.
- $(\mathrm{d}W_t^s)^2 = \mathrm{d}t$ for a Wiener process.
- $\mathrm{d}W_t^s\,\mathrm{d}W_t^v = \rho\mathrm{d}t$ for correlated Wiener processes ρ.

Thus,

$$\begin{aligned}(\mathrm{d}S_t)^2 &= V_t S_t^2 \mathrm{d}t,\\ (\mathrm{d}V_t)^2 &= \sigma^2 V_t \mathrm{d}t,\\ \mathrm{d}S_t\,\mathrm{d}V_t &= \rho\sigma V_t S_t \mathrm{d}t.\end{aligned}$$

Next, substitute these terms into Equation (A.1.1):

$$\begin{aligned}\mathrm{d}u = \frac{\partial u}{\partial t}\mathrm{d}t + \frac{\partial u}{\partial S}\left(\mu S_t \mathrm{d}t + \sqrt{V_t}S_t\,\mathrm{d}W_t^s\right) + \frac{\partial u}{\partial V}\left(\kappa(\bar{v}-V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_t^v\right)\\ + \frac{1}{2}\frac{\partial^2 u}{\partial S^2}V_t S_t^2 \mathrm{d}t + \frac{1}{2}\frac{\partial^2 u}{\partial V^2}\sigma^2 V_t \mathrm{d}t + \frac{\partial^2 u}{\partial S \partial V}\rho\sigma V_t S_t \mathrm{d}t. \qquad \text{(A.1.4)}\end{aligned}$$

Finally, group the terms proportional to $\mathrm{d}t$, $\mathrm{d}W_t^s$, and $\mathrm{d}W_t^v$. **Drift terms ($\mathrm{d}t$):**

$$\begin{aligned}&\frac{\partial u}{\partial t} \;+\; \mu S_t \frac{\partial u}{\partial S} \;+\; \kappa(\bar{v}-V_t)\frac{\partial u}{\partial V}\\ &\quad+\; \tfrac{1}{2}V_t S_t^2 \frac{\partial^2 u}{\partial S^2} \;+\; \tfrac{1}{2}\sigma^2 V_t \frac{\partial^2 u}{\partial V^2} \;+\; \rho\sigma V_t S_t \frac{\partial^2 u}{\partial S \partial V}.\end{aligned}$$

Using the results above, we now substitute everything back into Itô's formula and collect terms by powers of $\mathrm{d}t$, $\mathrm{d}W_t^s$, and $\mathrm{d}W_t^v$. This yields:

$$\begin{aligned}\mathrm{d}u = \left[\frac{\partial u}{\partial t} + \mu S_t \frac{\partial u}{\partial S} + \kappa(\bar{v}-V_t)\frac{\partial u}{\partial V} + \frac{1}{2}\frac{\partial^2 u}{\partial S^2}V_t S_t^2 + \frac{1}{2}\frac{\partial^2 u}{\partial V^2}\sigma^2 V_t + \frac{\partial^2 u}{\partial S \partial V}\rho\sigma V_t S_t\right]\mathrm{d}t\\ + \sqrt{V_t}S_t\frac{\partial u}{\partial S}\,\mathrm{d}W_t^s + \sigma\sqrt{V_t}\frac{\partial u}{\partial V}\,\mathrm{d}W_t^v.\end{aligned}$$

Transition to the Risk-Neutral Measure

To price options under the no-arbitrage principle, we adopt the **risk-neutral measure** and replace the drift term μ in the asset price dynamics with the risk-free rate r. This reflects the fact that, under the risk-neutral measure, the expected return of the asset matches the risk-free rate, ensuring no arbitrage opportunities. Explicit references to the risk-neutral measure (e.g. $\mathbb{Q}$) are omitted to simplify notation.

Incorporation of the Discounting Term $-ru$

To account for the time value of money under the risk-neutral measure, the option price $u(S_t, V_t, t)$ is expressed as the discounted expected value of its payoff. Thus, the process $e^{-rt}u(S_t, V_t, t)$ must be a martingale under $\mathbb{Q}$, meaning that its expected

change must be zero:

$$\mathbb{E}\left[\mathrm{d}\left(e^{-rt}u(S_t, V_t, t)\right)\right] = 0.$$

Applying Itô's Lemma to $e^{-rt}u(S_t, V_t, t)$, we obtain:

$$\mathrm{d}\left(e^{-rt}u(S_t, V_t, t)\right) = e^{-rt}\left(\mathrm{d}u - ru\mathrm{d}t\right).$$

Substitute the expression for $\mathrm{d}u$ (now under the risk-neutral measure, so μ is replaced with r):

$$\begin{aligned}\mathrm{d}\left(e^{-rt}u(S_t, V_t, t)\right) = e^{-rt}\Bigg[&\frac{\partial u}{\partial t} + rS_t\frac{\partial u}{\partial S} + \kappa(\bar{v} - V_t)\frac{\partial u}{\partial V} \\ &+ \frac{1}{2}\frac{\partial^2 u}{\partial S^2}V_t S_t^2 + \frac{1}{2}\frac{\partial^2 u}{\partial V^2}\sigma^2 V_t \\ &+ \frac{\partial^2 u}{\partial S \partial V}\rho\sigma V_t S_t - ru\Bigg]\mathrm{d}t \\ &+ e^{-rt}\sqrt{V_t}S_t\frac{\partial u}{\partial S}\,\mathrm{d}W_t^s + e^{-rt}\sigma\sqrt{V_t}\frac{\partial u}{\partial V}\,\mathrm{d}W_t^v.\end{aligned}$$

The last two terms are stochastic integrals with respect to Brownian motions $\mathrm{d}W_t^s$ and $\mathrm{d}W_t^v$, both of which are martingales with zero mean. Hence, when taking expectations, these terms vanish. Setting the expectation of the full expression to zero retains only the drift component:

$$\mathbb{E}\left[\mathrm{d}\left(e^{-rt}u(S_t, V_t, t)\right)\right] = 0.$$

This gives the governing partial differential equation for the Heston model:

$$\frac{\partial u}{\partial t} + rS_t\frac{\partial u}{\partial S} + \kappa(\bar{v} - V_t)\frac{\partial u}{\partial V} + \frac{1}{2}\frac{\partial^2 u}{\partial S^2}V_t S_t^2 + \frac{1}{2}\frac{\partial^2 u}{\partial V^2}\sigma^2 V_t + \frac{\partial^2 u}{\partial S \partial V}\rho\sigma V_t S_t - ru = 0. \quad \text{(A.1.5)}$$

This is the Heston option pricing PDE under the risk-neutral measure. In the next section, we simplify this PDE by transforming it to log-price coordinates.

A.2 Log-Transformation of the Heston PDE

To simplify the PDE and eliminate the explicit dependence on S_t^2, we introduce a change of variables by defining $s = \ln S_t$. We also simplify the notation by writing $V_t = v$. This change of variables simplifies the second-derivative term in the PDE by eliminating the explicit appearance of S_t, making the equation more suitable for

Fourier analysis in the next section. Using Equation (2.1) for S_t under the risk-neutral measure, we obtain:

$$\mathrm{d}S_t = rS_t\mathrm{d}t + \sqrt{V_t}S_t\,\mathrm{d}W^s.$$

Next, apply Itô's Lemma to $s = \ln S_t$. The derivative terms are:

$$\mathrm{d}s = \frac{1}{S_t}\mathrm{d}S_t - \frac{1}{2}\frac{1}{S_t^2}(\mathrm{d}S_t)^2.$$

Substitute $\mathrm{d}S_t = rS_t\mathrm{d}t + \sqrt{v}S_t\,\mathrm{d}W^s$ and $(\mathrm{d}S_t)^2 = vS_t^2\mathrm{d}t$:

$$\begin{aligned}\mathrm{d}s &= \frac{1}{S_t}\left(rS_t\mathrm{d}t + \sqrt{v}S_t\,\mathrm{d}W^s\right) - \frac{1}{2}\frac{1}{S_t^2}\left(vS_t^2\mathrm{d}t\right),\\ \mathrm{d}s &= \left(r - \frac{1}{2}v\right)\mathrm{d}t + \sqrt{v}\,\mathrm{d}W^s.\end{aligned}$$

Thus, the dynamics of $x = \ln S$ are:

$$\mathrm{d}s = \left(r - \frac{1}{2}v\right)\mathrm{d}t + \sqrt{v}\,\mathrm{d}W^s.$$

The variance $v(t)$ follows a square-root process:

$$\mathrm{d}v = \kappa(\bar{v} - v)\mathrm{d}t + \sigma\sqrt{v}\,\mathrm{d}W^v.$$

Let $f(x, v, t) = u(S_t, V_t, t)$ denote the transformed option price expressed in terms of log-price $s = \ln S_t$, variance $v = V_t$, and time t.

$$\mathrm{d}f = \frac{\partial f}{\partial t}\mathrm{d}t + \frac{\partial f}{\partial s}\,\mathrm{d}s + \frac{\partial f}{\partial v}\,\mathrm{d}v + \frac{1}{2}\frac{\partial^2 f}{\partial s^2}\,(\mathrm{d}s)^2 + \frac{1}{2}\frac{\partial^2 f}{\partial v^2}\,(\mathrm{d}v)^2 + \frac{\partial^2 f}{\partial s\partial v}\,\mathrm{d}s\,\mathrm{d}v.$$

From the dynamics, we get:

$$\begin{aligned}&\mathrm{d}s = \left(r - \frac{1}{2}v\right)\mathrm{d}t + \sqrt{v}\,\mathrm{d}W^s, \quad (\mathrm{d}s)^2 = v\mathrm{d}t,\\ &\mathrm{d}v = \kappa(\bar{v} - v)\mathrm{d}t + \sigma\sqrt{v}\,\mathrm{d}W^v, \quad (\mathrm{d}v)^2 = \sigma^2 v\mathrm{d}t,\\ &\mathrm{d}s\,\mathrm{d}v = \rho\sigma v\mathrm{d}t.\end{aligned}$$

Substitute these into Itô's Lemma for $f(s, v, t)$:

$$\mathrm{d}f = \left[\frac{\partial f}{\partial t} + \left(r - \frac{1}{2}v\right)\frac{\partial f}{\partial s} + \kappa(\bar{v} - v)\frac{\partial f}{\partial v} + \frac{1}{2}v\frac{\partial^2 f}{\partial s^2} + \frac{1}{2}\sigma^2 v\frac{\partial^2 f}{\partial v^2} + \rho\sigma v\frac{\partial^2 f}{\partial s \partial v}\right]\mathrm{d}t$$
$$+ \sqrt{v}\frac{\partial f}{\partial s}\,\mathrm{d}W^s + \sigma\sqrt{v}\frac{\partial f}{\partial v}\,\mathrm{d}W^v.$$

Under the risk-neutral measure, $f(s, v, t)$ must be a martingale, meaning $\mathbb{E}[\mathrm{d}f] = 0$. This eliminates the stochastic terms involving $\mathrm{d}W^s$ and $\mathrm{d}W^v$, leaving the following partial deferentialEquation:

$$\frac{\partial f}{\partial t} + \left(r - \frac{1}{2}v\right)\frac{\partial f}{\partial s} + \kappa(\bar{v} - v)\frac{\partial f}{\partial v} + \frac{1}{2}v\frac{\partial^2 f}{\partial s^2} + \frac{1}{2}\sigma^2 v\frac{\partial^2 f}{\partial v^2} + \rho\sigma v\frac{\partial^2 f}{\partial s \partial v} = 0. \quad \text{(A.2.1)}$$

This PDE is now expressed entirely in terms of $s = \ln S_t$ and v, with no explicit dependence on the asset price S_t. This prepares the equation for Fourier-based solution techniques in the next section.

Appendix B

Derivation of the Heston Model's Characteristic Function

This appendix derives the characteristic function of the Heston stochastic volatility model, a crucial analytical result enabling efficient option pricing via Fourier inversion techniques. Starting from the simplified PDE obtained in Appendix A, we employ Fourier transform methods and systematically solve the resulting Riccati equations for the functions $D(\tau)$ and $C(\tau)$. Each derivation step is clearly delineated to ensure transparency and reproducibility, aligning closely with academic rigor and facilitating further computational applications in calibration and pricing.

B.1 Derivation of the Characteristic Function for the Heston Model

We now derive the closed-form characteristic function of the Heston model by solving the PDE derived above using a Fourier transform method Heston (1993). In particular, we will obtain differential equations for two functions, $C(\tau)$ and $D(\tau)$, which determine the characteristic function. The derivation is based on Heston (1993), with additional explanations included for clarity and completeness. We begin with the log-space PDE presented in Equation (A.2.1), which forms the basis for deriving the Heston model's characteristic function. To align the PDE with the standard form in Heston (1993), we set $u = -\frac{1}{2}$ (a specific complex shift for the exponent), $a = \kappa\bar{v}$, and $b = \kappa$. These substitutions simplify the structure of the solution and ensure the moment-generating function is well-behaved. With these definitions, the PDE becomes:

$$-\frac{\partial f}{\partial t} = \frac{1}{2}v\frac{\partial^2 f}{\partial s^2} + \rho\sigma v\frac{\partial^2 f}{\partial s \partial v} + \frac{1}{2}\sigma^2 v\frac{\partial^2 f}{\partial v^2} + (r + uv)\frac{\partial f}{\partial s} + (a - bv)\frac{\partial f}{\partial v}. \tag{B.1.1}$$

Assume a solution of the form

$$f(s, v, t) = \exp[C(\tau) + D(\tau)v + i\phi s],$$

which corresponds to a Fourier transform in the s-direction. This ansatz reflects the characteristic function $\varphi(\phi; \tau) = \mathbb{E}[e^{i\phi s_T}]$. Substituting this expression into Equation (B.1.1), we proceed to compute all required derivatives. Differentiating f with respect to each variable gives:

$$\frac{\partial f}{\partial s} = i\phi f, \quad \frac{\partial^2 f}{\partial s^2} = -\phi^2 f, \quad \frac{\partial f}{\partial v} = Df, \quad \frac{\partial^2 f}{\partial v^2} = D^2 f, \quad \frac{\partial^2 f}{\partial s \partial v} = i\phi D f.$$

Since f depends on time through $C(\tau)$ and $D(\tau)$, and $\tau = T - t$, we apply the chain rule:

$$\frac{\partial f}{\partial t} = -\left(\frac{dC}{d\tau} + v\frac{dD}{d\tau}\right) f.$$

Substitute all derivatives into the PDE Equation (B.1.1) and divide through by $f \neq 0$:

$$\begin{aligned}\left(-\frac{\partial C}{\partial t} - v\frac{\partial D}{\partial t}\right) f = \frac{1}{2}v\left(-\phi^2 f\right) + \rho\sigma v \cdot i\phi D f + \frac{1}{2}\sigma^2 v D^2 f \\ +(r + uv) \cdot i\phi f + (a - bv)Df.\end{aligned}$$

Simplify the terms:

$$\begin{aligned}-\frac{\partial C}{\partial t} - v\frac{\partial D}{\partial t} = -\frac{1}{2}v\phi^2 + \rho\sigma v \cdot i\phi D + \frac{1}{2}\sigma^2 v D^2 \\ + i\phi r + i\phi uv + aD - bvD.\end{aligned}$$

Next, divide through by f and group the terms based on their dependence on v. The terms independent of v are:

$$\frac{\partial C}{\partial t} = -i\phi r - aD.$$

The terms proportional to v are:

$$-v\frac{\partial D}{\partial t} = -\frac{1}{2}v\phi^2 + \rho\sigma v\mathrm{i}\phi D + \frac{1}{2}\sigma^2 v D^2 + i\phi uv - bvD.$$

Dividing through by v, we get:

$$\frac{\partial D}{\partial t} = \frac{1}{2}\phi^2 - \rho\sigma i\phi D - \frac{1}{2}\sigma^2 D^2 - i\phi u + bD.$$

In the Heston model, it is convenient to rewrite this equation in terms of $\tau = T - t$ (time to maturity), which implies that:

$$\frac{\partial \tau}{\partial t} = -1,$$

$$\partial \tau = -\partial t.$$

Thus, any time derivative with respect to t can be rewritten in terms of τ as:

$$\frac{\partial D}{\partial t} = -\frac{\partial D}{\partial \tau}.$$

Then,

$$\begin{aligned} \frac{\partial D}{\partial \tau} &= -\frac{1}{2}\phi^2 + \rho\sigma \mathrm{i}\phi D + \frac{1}{2}\sigma^2 D^2 + i\phi u - bD, \\ \frac{\partial C}{\partial \tau} &= +i\phi r + aD. \end{aligned}$$

These two equations determine the time evolution of $C(\tau)$ and $D(\tau)$. Note that the equation for D is nonlinear (Riccati-type), while the equation for C is linear once $D(\tau)$ is known.

The next section presents the explicit solution to the Riccati equation for $D(\tau)$, and the corresponding integration to obtain $C(\tau)$.

B.2 Solutions for the Riccati Equations $D(\tau)$ and $C(\tau)$

This section presents a full, step-by-step solution to the ordinary differential equations for $D(\tau)$ and $C(\tau)$, obtained previously in Section B.1. We begin by solving the nonlinear Riccati equation for $D(\tau)$, then use this solution to integrate the linear equation for $C(\tau)$. These expressions define the closed-form characteristic function of the Heston model.

Solving the Riccati Equation for $D(\tau)$

The Riccati equation for $D(\tau)$ is:

$$\frac{\partial D}{\partial \tau} = CD^2 + BD + A, \tag{B.2.1}$$

where:

- $C = \frac{1}{2}\sigma^2$
- $B = \rho\sigma i\phi - b$
- $A = -\frac{1}{2}\phi^2 + i\phi u$.

This is a nonlinear Riccati equation due to the quadratic term in $D(\tau)$. We now apply the standard Riccati-to-linear transformation. Define a new function $y(\tau)$ via the substitution:

$$D(\tau) = -\frac{1}{C}\frac{\partial \ln y(\tau)}{\partial \tau} = -\frac{1}{C}\frac{y'(\tau)}{y(\tau)}, \tag{B.2.2}$$

where $y(\tau)$ is a new function. Taking the derivative of $D(\tau)$, we have:

$$\frac{\partial D}{\partial \tau} = -\frac{1}{C}\frac{\partial}{\partial \tau}\left(\frac{y'(\tau)}{y(\tau)}\right).$$

Using the quotient rule for derivatives, we get:

$$\frac{\partial}{\partial \tau}\left(\frac{y'(\tau)}{y(\tau)}\right) = \frac{y''(\tau)y(\tau) - (y'(\tau))^2}{(y(\tau))^2}.$$

Substitute this into Equation (B.2.1):

$$-\frac{1}{C}\left[\frac{y''(\tau)}{y(\tau)} - \left(\frac{y'(\tau)}{y(\tau)}\right)^2\right] = C\left(-\frac{1}{C}\frac{y'(\tau)}{y(\tau)}\right)^2 + B\left(-\frac{1}{C}\frac{y'(\tau)}{y(\tau)}\right) + A,$$
$$-\frac{1}{C}\frac{y''(\tau)}{y(\tau)} + \frac{1}{C}\left(\frac{y'(\tau)}{y(\tau)}\right)^2 = \frac{1}{C}\left(\frac{y'(\tau)}{y(\tau)}\right)^2 - \frac{B}{C}\frac{y'(\tau)}{y(\tau)} + A.$$

Cancel $\frac{1}{C}\left(\frac{y'(\tau)}{y(\tau)}\right)^2$ on both sides and multiply through by $-C$:

$$\frac{y''(\tau)}{y(\tau)} = B\frac{y'(\tau)}{y(\tau)} - AC.$$

This yields the second-order linear ODE:

$$y''(\tau) - By'(\tau) + ACy(\tau) = 0.$$

The general solution for a second-order linear differential equation is:

$$y(\tau) = c_1 e^{\lambda_1 \tau} + c_2 e^{\lambda_2 \tau},$$

where λ_1 and λ_2 are the roots of the characteristic Equation:

$$\lambda^2 - B\lambda + AC = 0.$$

The roots are:

$$\lambda_{1,2} = \frac{B \pm \sqrt{B^2 - 4AC}}{2}.$$

Let us define the discriminant:

$$\nu \coloneqq \sqrt{B^2 - 4AC} = \sqrt{(\rho\sigma i\phi - b)^2 - 2\sigma^2(i\phi u - \tfrac{1}{2}\phi^2)}.$$

Then, the roots are:

$$\lambda_1 = \frac{-b + \rho\sigma i\phi + \nu}{2}, \quad \lambda_2 = \frac{-b + \rho\sigma i\phi - \nu}{2}.$$

Substitute $y(\tau)$ into $D(\tau)$ using the Riccati substitution:

$$D(\tau) = -\frac{1}{C}\frac{\partial \ln y(\tau)}{\partial \tau},$$

where:

$$y(\tau) = c_1 e^{\lambda_1 \tau} + c_2 e^{\lambda_2 \tau}.$$

This is the general solution to the second-order linear ODE, where $\lambda_{1,2} = \frac{1}{2}(B \pm \nu)$. Compute $\ln y(\tau)$ by taking the natural logarithm of $y(\tau)$:

$$\ln y(\tau) = \ln\left(c_1 e^{\lambda_1 \tau} + c_2 e^{\lambda_2 \tau}\right).$$

Differentiate $\ln y(\tau)$ with respect to τ:

$$\frac{\partial \ln y(\tau)}{\partial \tau} = \frac{\frac{\partial y(\tau)}{\partial \tau}}{y(\tau)}.$$

Compute $\frac{\partial y(\tau)}{\partial \tau}$ by using the above form of $y(\tau)$:

$$\frac{\partial y(\tau)}{\partial \tau} = c_1\lambda_1 e^{\lambda_1 \tau} + c_2\lambda_2 e^{\lambda_2 \tau}.$$

Next, substitute this into equation (B.2.2):

$$D(\tau) = -\frac{1}{C}\left[\frac{\frac{\partial y(\tau)}{\partial \tau}}{y(\tau)}\right].$$

Substitute $y(\tau)$ and $\frac{\partial y(\tau)}{\partial \tau}$:

$$D(\tau) = -\frac{1}{C}\left[\frac{c_1\lambda_1 e^{\lambda_1\tau} + c_2\lambda_2 e^{\lambda_2\tau}}{c_1 e^{\lambda_1\tau} + c_2 e^{\lambda_2\tau}}\right].$$

To simplify, divide the numerator and denominator by $c_1 e^{\lambda_1\tau}$:

$$D(\tau) = -\frac{1}{C}\left[\frac{\lambda_1 + c\lambda_2 e^{(\lambda_2-\lambda_1)\tau}}{1 + ce^{(\lambda_2-\lambda_1)\tau}}\right], \quad \text{where} \quad c = \frac{c_2}{c_1}.$$

From the boundary condition $D(0) = 0$, we get:

$$\lambda_1 + c\lambda_2 = 0 \quad \Longrightarrow \quad c = -\frac{\lambda_1}{\lambda_2}.$$

Substitute c into the simplified form:

$$D(\tau) = -\frac{1}{C}\left[\frac{\lambda_1 - \frac{\lambda_1}{\lambda_2}\lambda_2 e^{(\lambda_2-\lambda_1)\tau}}{1 - \frac{\lambda_1}{\lambda_2}e^{(\lambda_2-\lambda_1)\tau}}\right].$$

Simplify the expression:

$$D(\tau) = -\frac{1}{C}\left[\frac{\lambda_1\left(1 - e^{(\lambda_2-\lambda_1)\tau}\right)}{1 - \frac{\lambda_1}{\lambda_2}e^{(\lambda_2-\lambda_1)\tau}}\right].$$

Now, recall the Heston-specific parameters:

$$-\nu = \lambda_2 - \lambda_1, \quad h = \frac{\lambda_1}{\lambda_2}.$$

Substitute these parameters into the expression:

$$D(\tau) = -\frac{\lambda_1}{C}\frac{1 - e^{-\nu\tau}}{1 - he^{-\nu\tau}}.$$

Substitute $\lambda_1 = \frac{-b+\rho\sigma i\phi+\nu}{2}$ and $\lambda_2 = \frac{-b+\rho\sigma i\phi-\nu}{2}$. Using $C = \frac{1}{2}\sigma^2$, we get:

$$D(\tau) = \frac{b - \rho\sigma i\phi - \nu}{\sigma^2}\left[\frac{1 - e^{-\nu\tau}}{1 - he^{-\nu\tau}}\right],$$

where:

$$\nu = \sqrt{(\rho\sigma i\phi - b)^2 - \sigma^2(2i\phi u - \phi^2)},$$
$$h = \frac{b - \rho\sigma i\phi - \nu}{b - \rho\sigma i\phi + \nu}.$$

If we substitute $\lambda_1 = \frac{-b+\rho\sigma i\phi-\nu}{2}$ and $\lambda_2 = \frac{-b+\rho\sigma i\phi+\nu}{2}$ and use $C = \frac{1}{2}\sigma^2$, we obtain:

$$D(\tau) = \frac{b - \rho\sigma i\phi + \nu}{\sigma^2}\left[\frac{1 - e^{\nu\tau}}{1 - he^{\nu\tau}}\right],$$

where:

$$\nu = \sqrt{(\rho\sigma i\phi - b)^2 - \sigma^2(2i\phi u - \phi^2)},$$
$$h = \frac{b - \rho\sigma i\phi + \nu}{b - \rho\sigma i\phi - \nu}.$$

Solving for $C(\tau)$

The equation for $C(\tau)$ is:

$$\frac{\partial C}{\partial \tau} = i\phi r + aD(\tau).$$

Substitute $D(\tau)$ into the differentialEquation:

$$\frac{\partial C}{\partial \tau} = +i\phi r + a\frac{(b - \rho\sigma i\phi - \nu)}{\sigma^2}\left[\frac{1 - e^{-\nu\tau}}{1 - he^{-\nu\tau}}\right].$$

Integrate both sides with respect to τ:

$$C(\tau) = \int_0^\tau \left[i\phi r + \frac{a(b - \rho\sigma i\phi - \nu)}{\sigma^2}\left(\frac{1 - e^{-\nu s}}{1 - he^{-\nu s}}\right)\right] \mathrm{d}s.$$

Next, split the equation into two terms. The integral of $-i\phi r$ is straightforward:

$$\int_0^\tau i\phi r \,\mathrm{d}s = i\phi r\tau.$$

The second term is:

$$\int_0^\tau \frac{1 - e^{-\nu s}}{1 - he^{-\nu s}} \,\mathrm{d}s.$$

To solve the integral, split the fraction into two terms:

$$\int_0^\tau \frac{1-e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s = \int_0^\tau \frac{1}{1-he^{-\nu s}}\,\mathrm{d}s - \int_0^\tau \frac{e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s.$$

The first term is:

$$\int_0^\tau \frac{1}{1-he^{-\nu s}}\,\mathrm{d}s.$$

Let $u = 1 - he^{-\nu s}$ so that $e^{-\nu s} = \frac{1-u}{h}$. Differentiating u, we obtain:

$$\mathrm{d}u = h\nu e^{-\nu s}\,\mathrm{d}s = \nu(1-u)\,\mathrm{d}s.$$

Solve for $\mathrm{d}s$:

$$\mathrm{d}s = \frac{\mathrm{d}u}{\nu(1-u)}.$$

Rewrite the integral:

$$\int_0^\tau \frac{1}{1-he^{-\nu s}}\,\mathrm{d}s = \int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{u}\cdot\frac{1}{\nu(1-u)}\,\mathrm{d}u.$$

Next, decompose $\frac{1}{u(1-u)}$ as

$$\frac{1}{u(1-u)} = \frac{A}{u} + \frac{B}{1-u},$$

where A and B are constants. Multiply through by $u(1-u)$ and simplify:

$$\begin{aligned} 1 &= A(1-u) + Bu, \\ 1 &= A - Au + Bu, \\ 1 &= A + (B-A)u, \end{aligned}$$

where:

$$A = 1, \quad B - A = 0 \implies B = 1.$$

Thus,

$$\frac{1}{u(1-u)} = \frac{1}{u} + \frac{1}{1-u}.$$

Use the partial fraction decomposition to obtain:

$$\int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{u(1-u)}\,\mathrm{d}u = \int_{1-h}^{1-he^{-\nu\tau}} \left(\frac{1}{u} + \frac{1}{1-u}\right)\,\mathrm{d}u.$$

Split the integral:

$$\int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{u(1-u)}\,\mathrm{d}u = \int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{u}\,\mathrm{d}u + \int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{1-u}\,\mathrm{d}u.$$

The first integral ($\frac{1}{u}$) is:

$$\int \frac{1}{u}\,\mathrm{d}u = \ln|u|.$$

Evaluate the definite integral:

$$\int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{u}\,\mathrm{d}u = \ln\left(1-he^{-\nu\tau}\right) - \ln(1-h) = \ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right).$$

The second integral ($\frac{1}{1-u}$) is:

$$\int \frac{1}{1-u}\,\mathrm{d}u = -\ln|1-u|.$$

Evaluate the definite integral:

$$\int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{1-u}\,\mathrm{d}u = -\ln\left(he^{-\nu\tau}\right) + \ln(h) = \ln\frac{h}{he^{-\nu\tau}} = \nu\tau.$$

Then,

$$\begin{aligned}\int_0^\tau \frac{1}{1-he^{-\nu s}}\,\mathrm{d}s &= \frac{1}{\nu}\left[\nu\tau + \ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right)\right] \\ &= \tau + \frac{1}{\nu}\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right).\end{aligned}$$

The second term is:

$$\int_0^\tau \frac{e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s.$$

Use the same substitution to obtain:

$$u = 1-he^{-\nu s}, \quad e^{-\nu s} = \frac{1-u}{h}.$$

Differentiating u, we get:

$$\mathrm{d}u = h\nu e^{-\nu s}\,\mathrm{d}s = \nu(1-u)\,\mathrm{d}s.$$

Solve for $\mathrm{d}s$:

$$\mathrm{d}s = \frac{\mathrm{d}u}{d(1-u)}.$$

Rewrite and simplify the integral:

$$\int_0^\tau \frac{e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s = -\int_{1-h}^{1-he^{-\nu\tau}} \frac{\frac{1-u}{h}}{u}\cdot\frac{1}{\nu(1-u)}\,\mathrm{d}u,$$
$$\int_0^\tau \frac{e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s = \frac{1}{\nu h}\int_{1-h}^{1-he^{-\nu\tau}} \frac{1}{u}\,\mathrm{d}u.$$

The integral of $\frac{1}{u}$ is:

$$\int \frac{1}{u}\,\mathrm{d}u = \ln|u|.$$

Then,

$$\int_0^\tau \frac{e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s = \frac{1}{\nu h}\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right).$$

Combine the results of the two terms:

$$\begin{aligned}\int_0^\tau \frac{1-e^{-\nu s}}{1-he^{-\nu s}}\,\mathrm{d}s &= \tau + \frac{1}{\nu}\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right) - \frac{1}{\nu h}\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right)\\ &= \tau + \left(\frac{1}{\nu}-\frac{1}{\nu h}\right)\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right)\\ &= \tau + \left(\frac{h-1}{\nu h}\right)\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right).\end{aligned}$$

Next, combine the all terms:

$$C(\tau) = -i\phi r\tau - \frac{a(b-\rho\sigma i\phi-\nu)}{\sigma^2}\cdot\left[\tau + \left(\frac{h-1}{\nu h}\right)\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right)\right].$$

We know that $h(b-\rho\sigma i\phi+\nu) = (b-\rho\sigma i\phi-\nu)$. Then,

$$C(\tau) = -i\phi r\tau - \frac{a}{\sigma^2}\cdot\left[(b-\rho\sigma i\phi-\nu)\tau + h(b-\rho\sigma i\phi+\nu)\left(\frac{h-1}{\nu h}\right)\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right)\right],$$

$$\begin{aligned}C(\tau) = i\,\phi\,r\,\tau + \frac{a}{\sigma^2}\Big[&(b-\rho\sigma i\phi-\nu)\,\tau\\ &+ (b-\rho\sigma i\phi+\nu)\left(\frac{(b-\rho\sigma i\phi-\nu)-(b-\rho\sigma i\phi+\nu)}{\nu\,(b-\rho\sigma i\phi+\nu)}\right)\ln\left(\frac{1-he^{-\nu\tau}}{1-h}\right)\Big].\end{aligned}$$

Simplify the term:

$$C(\tau) = i\phi r\tau + \frac{a}{\sigma^2} \cdot \left[(b - \rho\sigma i\phi - \nu)\tau + \left(\frac{-2\nu}{\nu} \right) \ln \left(\frac{1 - he^{-\nu\tau}}{1 - h} \right) \right],$$
$$C(\tau) = i\phi r\tau + \frac{a}{\sigma^2} \cdot \left[(b - \rho\sigma i\phi - \nu)\tau + -2 \ln \left(\frac{1 - he^{-\nu\tau}}{1 - h} \right) \right].$$

Finally, substitute $\lambda_1 = \frac{-b+\rho\sigma i\phi-\nu}{2}$ and $\lambda_2 = \frac{-b+\rho\sigma i\phi+\nu}{2}$ to obtain the following expression:

$$C(\tau) = i\phi r\tau + \frac{a}{\sigma^2} \cdot \left[(b - \rho\sigma i\phi + \nu)\tau + -2 \ln \left(\frac{1 - he^{\nu\tau}}{1 - h} \right) \right]. \tag{B.2.3}$$

These explicit solutions for $D(\tau)$ and $C(\tau)$ define the characteristic function of the Heston model as:

$$\varphi(\phi; \tau) = \exp\left(C(\tau) + D(\tau)v_0 + i\phi \ln S_0 \right),$$

which will be used in calibration and numerical pricing throughout the remainder of the thesis.

Appendix C

Second Derivatives of the Characteristic Function

This appendix presents detailed derivations of the second-order partial derivatives of the Heston characteristic function with respect to key model parameters.

C.1 Second Derivatives of φ with Respect to v_0

Let us define the intermediate function

$$f(v_0) = -\frac{A(v_0)}{v_0},$$

where both A and v_0 are treated as differentiable functions. The first derivative of φ with respect to v_0 is:

$$\frac{\partial \varphi}{\partial v_0} = f(v_0) \cdot \varphi.$$

To obtain the second derivative, differentiate again using the product rule:

$$\frac{\partial^2 \varphi}{\partial v_0^2} = \left(f'(v_0) + f(v_0)^2\right) \cdot \varphi.$$

Now, compute the exact form of $f'(v_0)$ using the quotient rule:

$$f'(v_0) = \frac{d}{\mathrm{d}v_0}\left(-\frac{A(v_0)}{v_0}\right) = -\left(\frac{v_0 \cdot \frac{\partial A}{\partial v_0} - A}{v_0^2}\right) = -\left(\frac{A - A}{v_0^2}\right) = 0.$$

This captures the dependencies on both v_0 and $A(v_0)$. Substituting into the second derivative, we get:

$$\frac{\partial^2 \varphi}{\partial v_0^2} = \left(\frac{A^2}{v_0^2}\right) \cdot \varphi.$$

This expression is exact and accounts for all v_0-dependence, both directly through

the denominator and indirectly via the function $A(v_0)$. This result is particularly relevant when assembling second-order derivatives for use in Newton-type calibration algorithms.

C.2 Second Derivatives of φ with Respect to ρ

The first derivative is:

$$\frac{\partial \varphi}{\partial \rho} = \varphi \cdot h_\rho, \quad \text{where} \quad h_\rho = -\frac{\partial A}{\partial \rho} + \frac{2\kappa\bar{v}}{\sigma^2 d}\left(\frac{\partial d}{\partial \rho} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \rho}\right) - \frac{\kappa\bar{v}\tau i u}{\sigma}.$$

The second derivative is:

$$\frac{\partial^2 \varphi}{\partial \rho^2} = \frac{\partial}{\partial \rho}\left(\varphi \cdot h_\rho\right) = \varphi \cdot h_\rho^2 + \varphi \cdot \frac{\partial h_\rho}{\partial \rho},$$

where $A, A_2, d,$ and ξ depend on ρ and the rest are constants. First, differentiate h_ρ with respect to ρ:

$$\frac{\partial h_\rho}{\partial \rho} = -\frac{\partial^2 A}{\partial \rho^2} + \frac{2\kappa\bar{v}}{\sigma^2}\cdot\frac{\partial}{\partial \rho}\left[\frac{1}{d}\left(\frac{\partial d}{\partial \rho} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \rho}\right)\right].$$

Let us denote

$$X \coloneqq \frac{\partial d}{\partial \rho} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \rho}.$$

Then, using the product rule:

$$\frac{\partial}{\partial \rho}\left(\frac{1}{d}X\right) = -\frac{1}{d^2}\frac{\partial d}{\partial \rho}X + \frac{1}{d}\cdot\frac{\partial X}{\partial \rho}.$$

Expand $\frac{\partial X}{\partial \rho}$:

$$\frac{\partial X}{\partial \rho} = \frac{\partial^2 d}{\partial \rho^2} - \frac{\partial}{\partial \rho}\left(\frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \rho}\right).$$

The second term expands as:

$$\frac{\partial}{\partial \rho}\left(\frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \rho}\right) = \left(\frac{A_2\frac{\partial d}{\partial \rho} - d\frac{\partial A_2}{\partial \rho}}{A_2^2}\right)\cdot\frac{\partial A_2}{\partial \rho} + \frac{d}{A_2}\cdot\frac{\partial^2 A_2}{\partial \rho^2}.$$

Substitute this back into $\frac{\partial h_\rho}{\partial \rho}$:

$$\frac{\partial h_\rho}{\partial \rho} = -\frac{\partial^2 A}{\partial \rho^2} + \frac{2\kappa\bar{v}}{\sigma^2}\left[-\frac{1}{d^2}\frac{\partial d}{\partial \rho}\left(\frac{\partial d}{\partial \rho} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \rho}\right) + \frac{1}{d}\left(\frac{\partial^2 d}{\partial \rho^2} - \left(\frac{A_2\frac{\partial d}{\partial \rho} - d\frac{\partial A_2}{\partial \rho}}{A_2^2}\right)\cdot\frac{\partial A_2}{\partial \rho} - \frac{d}{A_2}\cdot\frac{\partial^2 A_2}{\partial \rho^2}\right)\right].$$

First and Second Derivatives of A with Respect to ρ

The first derivative of $A = \frac{A_1}{A_2}$ with respect to ρ is:

$$\frac{\partial A}{\partial \rho} = \frac{1}{A_2}\frac{\partial A_1}{\partial \rho} - \frac{A}{A_2}\frac{\partial A_2}{\partial \rho}.$$

The second derivative is:

$$\frac{\partial^2 A}{\partial \rho^2} = \left(-\frac{1}{A_2^2}\frac{\partial A_2}{\partial \rho}\frac{\partial A_1}{\partial \rho} + \frac{1}{A_2}\frac{\partial^2 A_1}{\partial \rho^2}\right) - \left(\frac{1}{A_2}\frac{\partial A}{\partial \rho}\frac{\partial A_2}{\partial \rho} - \frac{A}{A_2^2}\left(\frac{\partial A_2}{\partial \rho}\right)^2 + \frac{A}{A_2}\frac{\partial^2 A_2}{\partial \rho^2}\right). \tag{C.2.1}$$

Simplified, this becomes:

$$\boxed{\frac{\partial^2 A}{\partial \rho^2} = \frac{1}{A_2}\frac{\partial^2 A_1}{\partial \rho^2} - \frac{1}{A_2^2}\frac{\partial A_1}{\partial \rho}\frac{\partial A_2}{\partial \rho} - \frac{1}{A_2}\frac{\partial A}{\partial \rho}\frac{\partial A_2}{\partial \rho} + \frac{A}{A_2^2}\left(\frac{\partial A_2}{\partial \rho}\right)^2 - \frac{A}{A_2}\frac{\partial^2 A_2}{\partial \rho^2}.}$$

Second Derivatives of A_1 and A_2 under $\xi(\rho)$

Let $\xi = \kappa - \rho\sigma i u$, where $d = \sqrt{(\kappa - \rho\sigma i u)^2 + \sigma^2(u^2 + iu)}$. Then,

$$\frac{\partial \xi}{\partial \rho} = -\sigma i u, \qquad \frac{\partial d}{\partial \rho} = \frac{1}{2d}\cdot\frac{\partial}{\partial \rho}\left[(\kappa - \rho\sigma i u)^2 + \sigma^2(u^2 + iu)\right] = -\frac{\sigma i u(\kappa - \rho\sigma i u)}{d}.$$

First, obtain the second derivative of A_2. Given

$$\frac{\partial A_2}{\partial \rho} = -\frac{\sigma i u(2 + \xi\tau)}{2dv_0}\left(\xi\cosh\frac{d\tau}{2} + d\sinh\frac{d\tau}{2}\right),$$

differentiate

$$\frac{\partial^2 A_2}{\partial \rho^2} = -\frac{\sigma i u}{2v_0}\left[\frac{\tau\frac{\partial \xi}{\partial \rho}\left(\xi\cosh\frac{d\tau}{2} + d\sinh\frac{d\tau}{2}\right)}{d} + \frac{(2 + \xi\tau)}{d^2}\cdot\left(d\frac{\partial g}{\partial \rho} - g\frac{\partial d}{\partial \rho}\right)\right],$$

where

$$g = \xi \cosh \frac{d\tau}{2} + d \sinh \frac{d\tau}{2}$$

and

$$\frac{\partial g}{\partial \rho} = \frac{\partial \xi}{\partial \rho} \cosh \frac{d\tau}{2} + \left(\xi \frac{\tau}{2} + 1\right) \sinh \frac{d\tau}{2} \frac{\partial d}{\partial \rho} + \frac{\tau d}{2} \cosh \frac{d\tau}{2} \frac{\partial d}{\partial \rho}.$$

Next, obtain the second derivative of A_1. Given

$$\frac{\partial A_1}{\partial \rho} = -\frac{iu(u^2 + iu)\tau\xi\sigma}{2d} \cosh \frac{d\tau}{2},$$

differentiate

$$\frac{\partial^2 A_1}{\partial \rho^2} = -\frac{iu(u^2 + iu)\tau\sigma}{2d^2} \left[\left(d\frac{\partial \xi}{\partial \rho} - \xi \frac{\partial d}{\partial \rho}\right) \cosh \frac{d\tau}{2} + \frac{\xi\tau d}{2} \sinh \frac{d\tau}{2} \frac{\partial d}{\partial \rho}\right].$$

Derivative of d with Respect to ρ

Given $d = \sqrt{\xi^2 + \sigma^2(u^2 + iu)}$, where $\xi = \kappa - \sigma\rho iu$, the first derivative of d is:

$$\frac{\partial d}{\partial \rho} = \frac{\xi}{d} \cdot \frac{\partial \xi}{\partial \rho} = \frac{\xi}{d}(-\sigma iu) = -\frac{\sigma iu\xi}{d}.$$

The second derivative of d is obtained as follows. Start with

$$\frac{\partial^2 d}{\partial \rho^2} = -\sigma iu \cdot \frac{\partial}{\partial \rho}\left(\frac{\xi}{d}\right)$$

and apply the quotient rule:

$$\frac{\partial}{\partial \rho}\left(\frac{\xi}{d}\right) = \frac{d \cdot \frac{\partial \xi}{\partial \rho} - \xi \cdot \frac{\partial d}{\partial \rho}}{d^2}.$$

Substitute

$$\frac{\partial \xi}{\partial \rho} = -\sigma iu, \quad \frac{\partial d}{\partial \rho} = -\frac{\sigma iu\xi}{d}. \tag{C.2.2}$$

Then,

$$\frac{\partial}{\partial \rho}\left(\frac{\xi}{d}\right) = \frac{-\sigma iud + \frac{\sigma iu\xi^2}{d}}{d^2} = \frac{\sigma iu(\xi^2 - d^2)}{d^3}.$$

Since

$$d^2 = \xi^2 + \sigma^2(u^2 + iu) \Rightarrow \xi^2 - d^2 = -\sigma^2(u^2 + iu),$$

we obtain:

$$\frac{\partial}{\partial \rho}\left(\frac{\xi}{d}\right) = -\frac{\sigma^3 iu(u^2+iu)}{d^3}.$$

The second derivative is then:

$$\frac{\partial^2 d}{\partial \rho^2} = -\sigma iu \cdot \left(-\frac{\sigma^3 iu(u^2+iu)}{d^3}\right) = \frac{\sigma^4 i^2 u^2(u^2+iu)}{d^3}.$$

Since $i^2 = -1$,

$$\boxed{\frac{\partial^2 d}{\partial \rho^2} = -\frac{\sigma^4 u^2(u^2+iu)}{d^3}.}$$

C.3 Second Derivatives of φ with Respect to κ

Given

$$\frac{\partial \varphi}{\partial \kappa} = \varphi \cdot h_\kappa,$$

$$h_\kappa = \frac{1}{\sigma iu} \cdot \frac{\partial A}{\partial \rho} + \frac{2\bar{v}}{\sigma^2} D + \frac{2\kappa \bar{v}}{\sigma^2 B} \cdot \frac{\partial B}{\partial \kappa} - \frac{\bar{v}\rho\tau iu}{\sigma},$$

the second derivative is:

$$\frac{\partial^2 \varphi}{\partial \kappa^2} = \frac{\partial}{\partial \kappa}(\varphi \cdot h_\kappa) = \varphi \cdot h_\kappa^2 + \varphi \cdot \frac{\partial h_\kappa}{\partial \kappa}.$$

To compute $\frac{\partial h_\kappa}{\partial \kappa}$, note that A, B, and D are functions of κ, while all other parameters are constant with respect to κ. Differentiating each term:

The first term is:

$$\frac{\partial}{\partial \kappa}\left(\frac{1}{\sigma iu} \cdot \frac{\partial A}{\partial \rho}\right) = \frac{1}{\sigma iu} \cdot \frac{\partial^2 A}{\partial \rho \partial \kappa}.$$

The second term is:

$$\frac{\partial}{\partial \kappa}\left(\frac{2\bar{v}}{\sigma^2} D\right) = \frac{2\bar{v}}{\sigma^2} \cdot \frac{\partial D}{\partial \kappa}.$$

The third term is (applying the product rule):

$$\frac{\partial}{\partial \kappa}\left(\frac{2\kappa\bar{v}}{\sigma^2 B} \cdot \frac{\partial B}{\partial \kappa}\right) = \frac{2\bar{v}}{\sigma^2}\left(\frac{1}{B} \cdot \frac{\partial B}{\partial \kappa} + \kappa \cdot \frac{\partial}{\partial \kappa}\left(\frac{1}{B} \cdot \frac{\partial B}{\partial \kappa}\right)\right).$$

Now compute the inner derivative:

$$\frac{\partial}{\partial \kappa}\left(\frac{1}{B} \cdot \frac{\partial B}{\partial \kappa}\right) = -\frac{1}{B^2}\left(\frac{\partial B}{\partial \kappa}\right)^2 + \frac{1}{B} \cdot \frac{\partial^2 B}{\partial \kappa^2}.$$

Thus, the third term becomes:

$$\frac{2\bar{v}}{\sigma^2}\left(\frac{1}{B}\cdot\frac{\partial B}{\partial\kappa}+\kappa\left(-\frac{1}{B^2}\left(\frac{\partial B}{\partial\kappa}\right)^2+\frac{1}{B}\cdot\frac{\partial^2 B}{\partial\kappa^2}\right)\right).$$

The fourth term is constant with respect to κ, so its derivative is zero:

$$\frac{\partial}{\partial\kappa}\left(-\frac{\bar{v}\rho\tau iu}{\sigma}\right)=0.$$

Combining all these, we obtain:

$$\boxed{\frac{\partial h_\kappa}{\partial\kappa}=\frac{1}{\sigma iu}\cdot\frac{\partial^2 A}{\partial\rho\partial\kappa}+\frac{2\bar{v}}{\sigma^2}\left(\frac{\partial D}{\partial\kappa}+\frac{1}{B}\cdot\frac{\partial B}{\partial\kappa}-\frac{\kappa}{B^2}\left(\frac{\partial B}{\partial\kappa}\right)^2+\frac{\kappa}{B}\cdot\frac{\partial^2 B}{\partial\kappa^2}\right).}$$

Second Derivative of A with Respect to κ

The first derivative is:

$$\frac{\partial A}{\partial\kappa}=\frac{i}{\sigma u}\cdot\frac{\partial A}{\partial\rho},$$
$$\frac{\partial A}{\partial\rho}=\frac{1}{A_2}\frac{\partial A_1}{\partial\rho}-\frac{A_1}{A_2^2}\frac{\partial A_2}{\partial\rho}.$$

The second derivative is:

$$\frac{\partial^2 A}{\partial\kappa^2}=\frac{i}{\sigma u}\cdot\frac{\partial}{\partial\kappa}\left(\frac{\partial A}{\partial\rho}\right).$$

Now, compute each term:

$$\frac{\partial}{\partial\kappa}\left(\frac{1}{A_2}\cdot\frac{\partial A_1}{\partial\rho}\right)=-\frac{1}{A_2^2}\cdot\frac{\partial A_2}{\partial\kappa}\cdot\frac{\partial A_1}{\partial\rho}+\frac{1}{A_2}\cdot\frac{\partial^2 A_1}{\partial\rho\partial\kappa},$$
$$\frac{\partial}{\partial\kappa}\left(\frac{A_1}{A_2^2}\cdot\frac{\partial A_2}{\partial\rho}\right)=\left(\frac{1}{A_2^2}\cdot\frac{\partial A_1}{\partial\kappa}-\frac{2A_1}{A_2^3}\cdot\frac{\partial A_2}{\partial\kappa}\right)\cdot\frac{\partial A_2}{\partial\rho}+\frac{A_1}{A_2^2}\cdot\frac{\partial^2 A_2}{\partial\rho\partial\kappa}.$$

The final expression is:

$$\boxed{\begin{aligned}\frac{\partial^2 A}{\partial\kappa^2}=\frac{i}{\sigma u}\Bigg[&\frac{1}{A_2}\cdot\frac{\partial^2 A_1}{\partial\rho\partial\kappa}-\frac{A_1}{A_2^2}\cdot\frac{\partial^2 A_2}{\partial\rho\partial\kappa}\\&-\frac{1}{A_2^2}\left(\frac{\partial A_2}{\partial\kappa}\cdot\frac{\partial A_1}{\partial\rho}+\frac{\partial A_1}{\partial\kappa}\cdot\frac{\partial A_2}{\partial\rho}\right)+\frac{2A_1}{A_2^3}\cdot\frac{\partial A_2}{\partial\kappa}\cdot\frac{\partial A_2}{\partial\rho}\Bigg].\end{aligned}}$$

Second Derivative of B with Respect to κ

Let us consider

$$\frac{\partial B}{\partial \kappa} = \frac{i}{\sigma u} \cdot \frac{\partial B}{\partial \rho} + \frac{B\tau}{2} \quad \text{and} \quad \frac{\partial B}{\partial \rho} = \frac{e^{k\tau/2}}{v_0} \left(\frac{1}{A_2} \cdot \frac{\partial d}{\partial \rho} - \frac{d}{A_2^2} \cdot \frac{\partial A_2}{\partial \rho} \right).$$

First, differentiate:

$$\frac{\partial^2 B}{\partial \kappa^2} = \frac{i}{\sigma u} \cdot \frac{\partial}{\partial \kappa} \left(\frac{\partial B}{\partial \rho} \right) + \frac{\tau}{2} \cdot \frac{\partial B}{\partial \kappa}.$$

Second, expand the derivatives:

$$\frac{\partial^2 B}{\partial \kappa^2} = \frac{i}{\sigma u} \cdot \frac{e^{k\tau/2}}{v_0} \cdot \frac{\partial}{\partial \kappa} \left(\frac{1}{A_2} \cdot \frac{\partial d}{\partial \rho} - \frac{d}{A_2^2} \cdot \frac{\partial A_2}{\partial \rho} \right) + \frac{\tau}{2} \left(\frac{i}{\sigma u} \cdot \frac{\partial B}{\partial \rho} + \frac{B\tau}{2} \right).$$

Third, apply the product rule:

$$\frac{\partial}{\partial \kappa} \left(\frac{1}{A_2} \cdot \frac{\partial d}{\partial \rho} \right) = -\frac{1}{A_2^2} \cdot \frac{\partial A_2}{\partial \kappa} \cdot \frac{\partial d}{\partial \rho} + \frac{1}{A_2} \cdot \frac{\partial^2 d}{\partial \rho \partial \kappa},$$

$$\frac{\partial}{\partial \kappa} \left(\frac{d}{A_2^2} \cdot \frac{\partial A_2}{\partial \rho} \right) = \left(\frac{1}{A_2^2} \cdot \frac{\partial d}{\partial \kappa} - \frac{2d}{A_2^3} \cdot \frac{\partial A_2}{\partial \kappa} \right) \cdot \frac{\partial A_2}{\partial \rho} + \frac{d}{A_2^2} \cdot \frac{\partial^2 A_2}{\partial \rho \partial \kappa}.$$

The final expression is then:

$$\boxed{\begin{aligned} \frac{\partial^2 B}{\partial \kappa^2} = \frac{i}{\sigma u} \cdot \frac{e^{k\tau/2}}{v_0} \Bigg[& \frac{1}{A_2} \cdot \frac{\partial^2 d}{\partial \rho \partial \kappa} - \frac{d}{A_2^2} \cdot \frac{\partial^2 A_2}{\partial \rho \partial \kappa} - \frac{1}{A_2^2} \left(\frac{\partial A_2}{\partial \kappa} \cdot \frac{\partial d}{\partial \rho} + \frac{\partial d}{\partial \kappa} \cdot \frac{\partial A_2}{\partial \rho} \right) \\ & + \frac{2d}{A_2^3} \cdot \frac{\partial A_2}{\partial \kappa} \cdot \frac{\partial A_2}{\partial \rho} \Bigg] + \frac{\tau}{2} \left(\frac{i}{\sigma u} \cdot \frac{\partial B}{\partial \rho} + \frac{B\tau}{2} \right). \end{aligned}}$$

Mixed Second Partial Derivative $\frac{\partial^2 d}{\partial \rho \partial \kappa}$

Start by obtaining the first partial derivative:

$$\frac{\partial d}{\partial \rho} = -\frac{\sigma i u \xi}{d}.$$

Differentiate with respect to κ:

$$\frac{\partial^2 d}{\partial \rho \partial \kappa} = -\sigma i u \cdot \frac{\partial}{\partial \kappa} \left(\frac{\xi}{d} \right).$$

Apply the quotient rule:

$$\frac{\partial}{\partial \kappa}\left(\frac{\xi}{d}\right) = \frac{d \cdot \frac{\partial \xi}{\partial \kappa} - \xi \cdot \frac{\partial d}{\partial \kappa}}{d^2}.$$

Then, substitute the derivatives:

$$\frac{\partial \xi}{\partial \kappa} = 1, \quad \frac{\partial d}{\partial \kappa} = \frac{\xi}{d} \Rightarrow \frac{\partial}{\partial \kappa}\left(\frac{\xi}{d}\right) = \frac{d^2 - \xi^2}{d^3}.$$

Simplify using identities:

$$d^2 = \xi^2 + \sigma^2(u^2 + iu) \Rightarrow d^2 - \xi^2 = \sigma^2(u^2 + iu).$$

The derivative is then:

$$\frac{\partial}{\partial \kappa}\left(\frac{\xi}{d}\right) = \frac{\sigma^2(u^2 + iu)}{d^3} \Rightarrow \frac{\partial^2 d}{\partial \rho \partial \kappa} = -\sigma i u \cdot \frac{\sigma^2(u^2 + iu)}{d^3} = -\frac{\sigma^3 iu(u^2 + iu)}{d^3},$$

$$\boxed{\frac{\partial^2 d}{\partial \rho \partial \kappa} = -\frac{\sigma^3 iu(u^2 + iu)}{d^3}.}$$

Mixed Second Partial Derivative of A_1

Compute the mixed second partial derivative $\frac{\partial^2 A_1}{\partial \rho \partial \kappa}$ using the intermediate expression

$$\begin{aligned}\frac{\partial A_1}{\partial \rho} &= -\frac{iu(u^2 + iu)\tau\xi\sigma}{2d}\cosh\left(\frac{d\tau}{2}\right) \\ &= C \cdot \frac{\xi}{d}\cosh\left(\frac{d\tau}{2}\right),\end{aligned}$$

where

$$C = -\frac{iu(u^2 + iu)\tau\sigma}{2}.$$

Differentiate with respect to κ to obtain:

$$\begin{aligned}\frac{\partial^2 A_1}{\partial \rho \partial \kappa} &= C \cdot \frac{\partial}{\partial \kappa}\left(\frac{\xi}{d}\cosh\left(\frac{d\tau}{2}\right)\right) \\ &= C\left(\frac{\partial}{\partial \kappa}\left(\frac{\xi}{d}\right)\cosh\left(\frac{d\tau}{2}\right) + \frac{\xi}{d} \cdot \frac{\partial}{\partial \kappa}\cosh\left(\frac{d\tau}{2}\right)\right).\end{aligned}$$

Using the quotient and chain rules, we have:

$$\frac{\partial}{\partial\kappa}\left(\frac{\xi}{d}\right) = \frac{\sigma^2(u^2+iu)}{d^3},$$
$$\frac{\partial}{\partial\kappa}\cosh\left(\frac{d\tau}{2}\right) = \sinh\left(\frac{d\tau}{2}\right)\cdot\frac{\tau}{2}\cdot\frac{\xi}{d}.$$

Combine the terms:

$$\frac{\partial^2 A_1}{\partial\rho\partial\kappa} = C\left(\frac{\sigma^2(u^2+iu)}{d^3}\cosh\left(\frac{d\tau}{2}\right) + \frac{\tau\xi^2}{2d^2}\sinh\left(\frac{d\tau}{2}\right)\right).$$

Lastly, substitute C:

$$\boxed{\frac{\partial^2 A_1}{\partial\rho\partial\kappa} = -\frac{iu(u^2+iu)\tau\sigma^3(u^2+iu)}{2d^3}\cosh\left(\frac{d\tau}{2}\right) - \frac{iu(u^2+iu)\tau^2\sigma\xi^2}{4d^2}\sinh\left(\frac{d\tau}{2}\right).}$$

Mixed Partial Derivative $\frac{\partial^2 A_2}{\partial\rho\partial\kappa}$

Given

$$\frac{\partial A_2}{\partial\rho} = -\frac{\sigma iu(2+\xi\tau)}{2dv_0}\left(\xi\cosh\left(\frac{d\tau}{2}\right) + d\sinh\left(\frac{d\tau}{2}\right)\right),$$

let

$$f(\xi,d) = -\frac{\sigma iu(2+\xi\tau)}{2dv_0}, \quad g(\xi,d) = \xi\cosh\left(\frac{d\tau}{2}\right) + d\sinh\left(\frac{d\tau}{2}\right).$$

First, differentiate the product:

$$\frac{\partial^2 A_2}{\partial\rho\partial\kappa} = \frac{\partial f}{\partial\kappa}\cdot g + f\cdot\frac{\partial g}{\partial\kappa}.$$

Second, compute $\frac{\partial f}{\partial\kappa}$:

$$\frac{\partial f}{\partial\kappa} = -\frac{\sigma iu\tau}{2dv_0} + \frac{\sigma iu(2+\xi\tau)\xi}{2d^3v_0}.$$

Third, compute $\frac{\partial g}{\partial\kappa}$:

$$\frac{\partial g}{\partial\kappa} = \cosh\left(\frac{d\tau}{2}\right) + \frac{\xi}{d}\left[\left(\frac{\xi\tau}{2}+1\right)\sinh\left(\frac{d\tau}{2}\right) + \frac{d\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\right].$$

The final expression is then:

$$\boxed{\begin{aligned}\frac{\partial^2 A_2}{\partial\rho\partial\kappa} = &\left(-\frac{\sigma i u\tau}{2dv_0} + \frac{\sigma i u(2+\xi\tau)\xi}{2d^3 v_0}\right)\left(\xi\cosh\left(\frac{d\tau}{2}\right) + d\sinh\left(\frac{d\tau}{2}\right)\right)\\ &- \frac{\sigma i u(2+\xi\tau)}{2dv_0}\left(\cosh\left(\frac{d\tau}{2}\right) + \frac{\xi}{d}\Big[\left(\frac{\xi\tau}{2}+1\right)\sinh\left(\frac{d\tau}{2}\right) + \frac{d\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\Big]\right).\end{aligned}} \tag{C.3.1}$$

Let $L = \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}e^{-d\tau}$. Then,

$$D = \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa-d)\tau}{2} - \ln L.$$

First Derivative of D with Respect to κ

Let us consider

$$\begin{aligned}\frac{\partial\xi}{\partial\kappa} &= 1,\\ \frac{\partial d}{\partial\kappa} &= \frac{\xi}{d},\\ \frac{\partial}{\partial\kappa}\ln\left(\frac{d}{v_0}\right) &= \frac{\xi}{d^2},\\ \frac{\partial}{\partial\kappa}\left(\frac{(\kappa-d)\tau}{2}\right) &= \frac{\tau}{2}\left(1-\frac{\xi}{d}\right).\end{aligned}$$

The derivative of A with respect to κ is:

$$\frac{\partial L}{\partial\kappa} = \frac{1}{2v_0}\left[\left(\frac{\xi}{d}+1\right) + \left(\frac{\xi}{d}-1\right)e^{-d\tau} - \tau(d-\xi)\cdot\frac{\xi}{d}e^{-d\tau}\right].$$

Hence,

$$\frac{\partial D}{\partial\kappa} = \frac{\xi}{d^2} + \frac{\tau}{2}\left(1-\frac{\xi}{d}\right) - \frac{1}{L}\cdot\frac{\partial L}{\partial\kappa}.$$

Derivative of A_1 with Respect to κ

Using the chain rule, we obtain:

$$\begin{aligned}\frac{\partial A_1}{\partial \kappa} &= (u^2 + iu) \cdot \frac{\partial}{\partial \kappa}\left(\sinh\left(\frac{d\tau}{2}\right)\right) \\ &= (u^2 + iu) \cdot \cosh\left(\frac{d\tau}{2}\right) \cdot \frac{\tau}{2} \cdot \frac{\partial d}{\partial \kappa}.\end{aligned}$$

Since $\frac{\partial d}{\partial \kappa} = \frac{\xi}{d}$, we get:

$$\boxed{\frac{\partial A_1}{\partial \kappa} = \frac{(u^2 + iu)\tau\xi}{2d} \cosh\left(\frac{d\tau}{2}\right).}$$

Derivative of A_2 with Respect to κ

Applying the product and chain rules, we get:

$$\begin{aligned}\frac{\partial A_2}{\partial \kappa} = &\frac{1}{v_0}\frac{\partial d}{\partial \kappa}\cosh\left(\frac{d\tau}{2}\right) + \frac{d}{v_0} \cdot \frac{\partial}{\partial \kappa}\left(\cosh\left(\frac{d\tau}{2}\right)\right) \\ &+ \frac{1}{v_0}\frac{\partial \xi}{\partial \kappa}\sinh\left(\frac{d\tau}{2}\right) + \frac{\xi}{v_0} \cdot \frac{\partial}{\partial \kappa}\left(\sinh\left(\frac{d\tau}{2}\right)\right).\end{aligned}$$

Simplify each term:

$$\begin{aligned}&\text{(i)} \quad \frac{1}{v_0} \cdot \frac{\xi}{d}\cosh\left(\frac{d\tau}{2}\right), \\ &\text{(ii)} \quad \frac{\tau\xi}{2v_0}\sinh\left(\frac{d\tau}{2}\right), \\ &\text{(iii)} \quad \frac{1}{v_0}\sinh\left(\frac{d\tau}{2}\right), \\ &\text{(iv)} \quad \frac{\tau\xi^2}{2v_0 d}\cosh\left(\frac{d\tau}{2}\right).\end{aligned}$$

Next, combine all the terms:

$$\begin{aligned}\frac{\partial A_2}{\partial \kappa} &= \frac{\xi}{v_0 d}\cosh\left(\frac{d\tau}{2}\right) + \frac{\tau\xi}{2v_0}\sinh\left(\frac{d\tau}{2}\right) + \frac{1}{v_0}\sinh\left(\frac{d\tau}{2}\right) + \frac{\tau\xi^2}{2v_0 d}\cosh\left(\frac{d\tau}{2}\right) \\ &= \frac{1}{v_0}\left[\left(\frac{\xi}{d} + \frac{\tau\xi^2}{2d}\right)\cosh\left(\frac{d\tau}{2}\right) + \left(1 + \frac{\tau\xi}{2}\right)\sinh\left(\frac{d\tau}{2}\right)\right].\end{aligned}$$

Thus,

$$\boxed{\frac{\partial A_2}{\partial \kappa} = \frac{1}{v_0}\left[\frac{\xi}{d}\left(1+\frac{\tau\xi}{2}\right)\cosh\left(\frac{d\tau}{2}\right)+\left(1+\frac{\tau\xi}{2}\right)\sinh\left(\frac{d\tau}{2}\right)\right].}$$

C.4 Second Derivatives of φ with Respect to σ

Given

$$\frac{\partial \varphi}{\partial \sigma} = \varphi \cdot h_\sigma,$$

$$h_\sigma = -\frac{\partial A}{\partial \sigma} - \frac{4\kappa\bar{v}}{\sigma^3}D + \frac{2\kappa\bar{v}}{\sigma^2 d}\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \sigma}\right) + \frac{\kappa\bar{v}\rho\tau i u}{\sigma^2},$$

the second derivative is derived as follows:

$$\frac{\partial^2 \varphi}{\partial \sigma^2} = \frac{\partial}{\partial \sigma}\left(\varphi \cdot h_\sigma\right) = \varphi \cdot h_\sigma^2 + \varphi \cdot \frac{\partial h_\sigma}{\partial \sigma}.$$

First Derivative of h_σ with Respect to σ

Here, we compute $\frac{\partial h_\sigma}{\partial \sigma}$. The derivative of $-\frac{\partial A}{\partial \sigma}$ is:

$$\frac{\partial}{\partial \sigma}\left(-\frac{\partial A}{\partial \sigma}\right) = -\frac{\partial^2 A}{\partial \sigma^2}.$$

The derivative of $-\frac{4\kappa\bar{v}}{\sigma^3}D$ is:

$$\frac{\partial}{\partial \sigma}\left(-\frac{4\kappa\bar{v}}{\sigma^3}D\right) = \frac{12\kappa\bar{v}}{\sigma^4}D - \frac{4\kappa\bar{v}}{\sigma^3}\cdot\frac{\partial D}{\partial \sigma}.$$

Next, obtain the derivative of $\frac{2\kappa\bar{v}}{\sigma^2 d}\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \sigma}\right)$. Let us define:

$$T = \frac{2\kappa\bar{v}}{\sigma^2 d}\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \sigma}\right).$$

Apply the product rule:

$$\frac{\partial T}{\partial \sigma} = \left(\frac{\partial}{\partial \sigma}\frac{2\kappa\bar{v}}{\sigma^2 d}\right)\cdot\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \sigma}\right) + \frac{2\kappa\bar{v}}{\sigma^2 d}\cdot\frac{\partial}{\partial \sigma}\left(\frac{\partial d}{\partial \sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial \sigma}\right).$$

The subcomponents are as follows. The coefficient derivative is:

$$\frac{\partial}{\partial\sigma}\left(\frac{2\kappa\bar{v}}{\sigma^2 d}\right) = -\frac{4\kappa\bar{v}}{\sigma^3 d} - \frac{2\kappa\bar{v}}{\sigma^2 d^2}\cdot\frac{\partial d}{\partial\sigma}.$$

The inner derivative is:

$$\frac{\partial}{\partial\sigma}\left(\frac{\partial d}{\partial\sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial\sigma}\right) = \frac{\partial^2 d}{\partial\sigma^2} - \left(\frac{A_2\cdot\frac{\partial d}{\partial\sigma} - d\cdot\frac{\partial A_2}{\partial\sigma}}{A_2^2}\right)\cdot\frac{\partial A_2}{\partial\sigma} - \frac{d}{A_2}\cdot\frac{\partial^2 A_2}{\partial\sigma^2}.$$

Next, obtain the derivative of $\frac{\kappa\bar{v}\rho\tau iu}{\sigma^2}$:

$$\frac{\partial}{\partial\sigma}\left(\frac{\kappa\bar{v}\rho\tau iu}{\sigma^2}\right) = -\frac{2\kappa\bar{v}\rho\tau iu}{\sigma^3}.$$

Finally, the combined expression is:

$$\begin{aligned}\frac{\partial h_\sigma}{\partial\sigma} = &-\frac{\partial^2 A}{\partial\sigma^2} + \frac{12\kappa\bar{v}}{\sigma^4}D - \frac{4\kappa\bar{v}}{\sigma^3}\cdot\frac{\partial D}{\partial\sigma} \\ &- \left(\frac{4\kappa\bar{v}}{\sigma^3 d} + \frac{2\kappa\bar{v}}{\sigma^2 d^2}\cdot\frac{\partial d}{\partial\sigma}\right)\cdot\left(\frac{\partial d}{\partial\sigma} - \frac{d}{A_2}\cdot\frac{\partial A_2}{\partial\sigma}\right) \\ &+ \frac{2\kappa\bar{v}}{\sigma^2 d}\left[\frac{\partial^2 d}{\partial\sigma^2} - \frac{A_2\left(\frac{\partial d}{\partial\sigma}\right)^2 - d\cdot\frac{\partial d}{\partial\sigma}\cdot\frac{\partial A_2}{\partial\sigma}}{A_2^2} - \frac{d}{A_2}\cdot\frac{\partial^2 A_2}{\partial\sigma^2}\right] \\ &- \frac{2\kappa\bar{v}\rho\tau iu}{\sigma^3}.\end{aligned}$$

Second Derivative of A with Respect to σ

We compute $\frac{\partial^2 A}{\partial\sigma^2}$ using the expression for $\frac{\partial A}{\partial\sigma}$ fromEquation (30d) of Cui et al. (2017):

$$\frac{\partial A}{\partial\sigma} = \frac{1}{A_2}\cdot\frac{\partial A_1}{\partial\sigma} - \frac{A}{A_2}\cdot\frac{\partial A_2}{\partial\sigma}.$$

Differentiating with respect to σ:

$$\begin{aligned}\frac{\partial^2 A}{\partial \sigma^2} &= \frac{\partial}{\partial \sigma}\left(\frac{1}{A_2}\cdot\frac{\partial A_1}{\partial \sigma}\right) - \frac{\partial}{\partial \sigma}\left(\frac{A}{A_2}\cdot\frac{\partial A_2}{\partial \sigma}\right) \\ &= \left[-\frac{1}{A_2^2}\cdot\frac{\partial A_2}{\partial \sigma}\cdot\frac{\partial A_1}{\partial \sigma} + \frac{1}{A_2}\cdot\frac{\partial^2 A_1}{\partial \sigma^2}\right] \\ &\quad - \left[\left(\frac{1}{A_2}\frac{\partial A}{\partial \sigma} - \frac{A}{A_2^2}\frac{\partial A_2}{\partial \sigma}\right)\frac{\partial A_2}{\partial \sigma} + \frac{A}{A_2}\cdot\frac{\partial^2 A_2}{\partial \sigma^2}\right].\end{aligned}$$

Substitute $\frac{\partial A}{\partial \sigma}$ and simplify:

$$\begin{aligned}&= -\frac{1}{A_2^2}\frac{\partial A_2}{\partial \sigma}\frac{\partial A_1}{\partial \sigma} + \frac{1}{A_2}\frac{\partial^2 A_1}{\partial \sigma^2} \\ &\quad - \left[\frac{1}{A_2}\left(\frac{1}{A_2}\frac{\partial A_1}{\partial \sigma} - \frac{A}{A_2}\frac{\partial A_2}{\partial \sigma}\right)\frac{\partial A_2}{\partial \sigma} - \frac{A}{A_2^2}\left(\frac{\partial A_2}{\partial \sigma}\right)^2 + \frac{A}{A_2}\frac{\partial^2 A_2}{\partial \sigma^2}\right] \\ &= \frac{1}{A_2}\frac{\partial^2 A_1}{\partial \sigma^2} - \frac{A}{A_2}\frac{\partial^2 A_2}{\partial \sigma^2} - \frac{2}{A_2^2}\frac{\partial A_1}{\partial \sigma}\frac{\partial A_2}{\partial \sigma} + \frac{2A}{A_2^2}\left(\frac{\partial A_2}{\partial \sigma}\right)^2.\end{aligned}$$

The final expression is:

$$\boxed{\frac{\partial^2 A}{\partial \sigma^2} = \frac{1}{A_2}\frac{\partial^2 A_1}{\partial \sigma^2} - \frac{A}{A_2}\frac{\partial^2 A_2}{\partial \sigma^2} - \frac{2}{A_2^2}\frac{\partial A_1}{\partial \sigma}\frac{\partial A_2}{\partial \sigma} + \frac{2A}{A_2^2}\left(\frac{\partial A_2}{\partial \sigma}\right)^2.}$$

Second Derivatives $\frac{\partial^2 A_1}{\partial \sigma^2}$ and $\frac{\partial^2 A_2}{\partial \sigma^2}$

First, for A_1:

$$\frac{\partial A_1}{\partial \sigma} = \frac{(u^2+iu)\tau}{2}\cdot\frac{\partial d}{\partial \sigma}\cosh\left(\frac{d\tau}{2}\right),$$

$$\frac{\partial^2 A_1}{\partial \sigma^2} = \frac{(u^2+iu)\tau}{2}\left[\frac{\partial^2 d}{\partial \sigma^2}\cosh\left(\frac{d\tau}{2}\right) + \frac{\tau}{2}\left(\frac{\partial d}{\partial \sigma}\right)^2\sinh\left(\frac{d\tau}{2}\right)\right].$$

With:

$$\frac{\partial d}{\partial \sigma} = \frac{\sigma(u^2+iu) - \xi\rho i u}{d}, \quad \frac{\partial^2 d}{\partial \sigma^2} = \frac{u^2+iu+\rho^2u^2}{d} - \frac{\left(\sigma(u^2+iu) - \xi\rho i u\right)^2}{d^3}.$$

Thus:

$$\boxed{\begin{aligned}\frac{\partial^2 A_1}{\partial\sigma^2} = \frac{(u^2+iu)\tau}{2}\Bigg[&\frac{u^2+iu+\rho^2u^2}{d}\cosh\left(\frac{d\tau}{2}\right)\\ &+\frac{\tau}{2}\cdot\frac{\left(\sigma(u^2+iu)-\xi\rho iu\right)^2}{d^2}\sinh\left(\frac{d\tau}{2}\right)\\ &-\frac{\left(\sigma(u^2+iu)-\xi\rho iu\right)^2}{d^3}\cosh\left(\frac{d\tau}{2}\right)\Bigg].\end{aligned}}$$

For A_2:

$$\frac{\partial A_2}{\partial\sigma} = \frac{\rho}{\sigma}\frac{\partial A_2}{\partial\rho} - \frac{2+\tau\xi}{v_0\tau\xi}\frac{\partial A_1}{\partial\rho} + \frac{\sigma\tau A_1}{2v_0},$$

$$\boxed{\begin{aligned}\frac{\partial^2 A_2}{\partial\sigma^2} = &-\frac{\rho}{\sigma^2}\frac{\partial A_2}{\partial\rho} + \frac{\rho}{\sigma}\frac{\partial^2 A_2}{\partial\rho\partial\sigma} + \frac{\rho iu(2+\tau\xi)}{v_0\sigma\xi^2}\frac{\partial A_1}{\partial\rho}\\ &-\frac{2+\tau\xi}{v_0\tau\xi}\frac{\partial^2 A_1}{\partial\rho\partial\sigma} + \frac{\tau A_1}{2v_0} + \frac{\sigma\tau}{2v_0}\frac{\partial A_1}{\partial\sigma}.\end{aligned}}$$

Derivatives of d and D with Respect to σ

For d:

$$\frac{\partial d}{\partial\sigma} = \frac{1}{d}\left(\sigma(u^2+iu)-\xi\rho iu\right),$$

$$\boxed{\frac{\partial^2 d}{\partial\sigma^2} = \frac{1}{d^3}\left[d^2(u^2+iu+\rho^2u^2) - \left(\sigma(u^2+iu)-\xi\rho iu\right)^2\right].}$$

For D:

$$D = \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa-d)\tau}{2} - \ln L,$$

$$L = \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}e^{-d\tau},$$

$$\boxed{\begin{aligned}\frac{\partial D}{\partial\sigma} = &\frac{-\xi\rho iu+\sigma(u^2+iu)}{d^2} + \frac{\tau(\xi\rho iu-\sigma(u^2+iu))}{2d}\\ &-\frac{1}{L}\cdot\frac{1}{2v_0}\Bigg[\left(\frac{\partial d}{\partial\sigma}-\rho iu\right)+\left(\frac{\partial d}{\partial\sigma}+\rho iu\right)e^{-d\tau}\\ &-\tau(d-\xi)\frac{\partial d}{\partial\sigma}e^{-d\tau}\Bigg].\end{aligned}}$$

Appendix D

Derivatives with Respect to Time-to-Maturity

Derivatives of $D(u,\tau)$ and the Heston Characteristic Function with Respect to Maturity τ

This appendix presents analytical derivatives of the Heston characteristic function with respect to time to maturity τ. We derive first and second-order derivatives for both the auxiliary function $D(u,\tau)$ and the characteristic function $\varphi(u;\tau)$, using the stabilized formulation from Cui et al. (2017).

Characteristic Function Specification

The characteristic function is defined as:

$$\varphi(u;\tau) = \exp\left(iu\ln\frac{F}{S_0} - \frac{\kappa\bar{v}\rho\tau iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right),$$

with components:

$$\begin{aligned}
F &= S_0 e^{(r-q)\tau} \quad \text{(forward price)},\\
d &= \sqrt{\xi^2 + \sigma^2(iu+u^2)},\\
\xi &= \kappa - \rho\sigma iu,\\
A &= \frac{A_1}{A_2},\\
A_1 &= (u^2+iu)\sinh\frac{d\tau}{2},\\
A_2 &= \frac{d}{v_0}\cosh\frac{d\tau}{2} + \frac{\xi}{v_0}\sinh\frac{d\tau}{2}.
\end{aligned}$$

The function $D(u,\tau)$ is defined through:

$$D(u,\tau) = \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L(\tau),$$

where $L(\tau)$ is:

$$L(\tau) \coloneqq \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}e^{-d\tau}.$$

first Derivative of $D(u,\tau)$ define:

$$\begin{aligned}\frac{\partial D}{\partial \tau} &= \frac{\kappa - d}{2} - \frac{1}{L}\frac{\partial L}{\partial \tau} \\ &= \frac{\kappa - d}{2} + \frac{d(d-\xi)e^{-d\tau}}{(d+\xi) + (d-\xi)e^{-d\tau}} \qquad \text{(D.0.1)} \\ &= \frac{\kappa - d}{2} + \frac{d(d-\xi)e^{-d\tau}}{2v_0 L(\tau)} \qquad \text{(D.0.2)}\end{aligned}$$

where:

$$\frac{\partial L}{\partial \tau} = -\frac{d(d-\xi)}{2v_0}e^{-d\tau}$$

where we note that $(d+\xi) + (d-\xi)e^{-d\tau} = 2v_0 L(\tau)$.

Second Derivative of $D(u,\tau)$ define:

$$f(\tau) \equiv \frac{\partial D}{\partial \tau} - \frac{\kappa - d}{2} = \frac{d(d-\xi)e^{-d\tau}}{2v_0 L(\tau)}$$

Differentiate $f(\tau)$:

$$\begin{aligned}\frac{\partial f}{\partial \tau} &= \frac{1}{2v_0}\frac{\partial}{\partial \tau}\left(\frac{d(d-\xi)e^{-d\tau}}{L(\tau)}\right) \\ &= \frac{1}{2v_0}\cdot\frac{\left[-d^2(d-\xi)e^{-d\tau}\right]L(\tau) - \left[d(d-\xi)e^{-d\tau}\right]\frac{\partial L}{\partial \tau}}{L(\tau)^2} \\ &= \frac{1}{2v_0}\cdot\frac{-d^2(d-\xi)e^{-d\tau}L(\tau) + \frac{d^2(d-\xi)^2}{2v_0}e^{-2d\tau}}{L(\tau)^2}\end{aligned}$$

Thus:

$$\boxed{\frac{\partial^2 D}{\partial \tau^2} = \frac{-d^2(d-\xi)e^{-d\tau}}{2v_0 L(\tau)} + \frac{d^2(d-\xi)^2 e^{-2d\tau}}{4v_0^2 L(\tau)^2}}$$

First Derivative of the Characteristic Function

Let $X = iu \ln \frac{F}{S_0} - \frac{\kappa\bar{v}\rho\tau iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2} D$, so $\varphi = e^X$ and:

$$\frac{\partial \varphi}{\partial \tau} = \frac{\partial X}{\partial \tau} \varphi.$$

The component derivatives are:

$$\begin{aligned} \frac{\partial}{\partial \tau}\left(iu \ln \frac{F}{S_0}\right) &= iu(r-q), \\ \frac{\partial}{\partial \tau}\left(-\frac{\kappa\bar{v}\rho\tau iu}{\sigma}\right) &= -\frac{\kappa\bar{v}\rho iu}{\sigma}, \\ \frac{\partial A}{\partial \tau} &= \frac{\frac{\partial A_1}{\partial \tau} A_2 - A_1 \frac{\partial A_2}{\partial \tau}}{A_2^2}, \end{aligned}$$

where:

$$\begin{aligned} \frac{\partial A_1}{\partial \tau} &= (u^2 + iu) \cdot \frac{d}{2} \cosh \frac{d\tau}{2}, \\ \frac{\partial A_2}{\partial \tau} &= \frac{d^2}{2v_0} \sinh \frac{d\tau}{2} + \frac{d\xi}{2v_0} \cosh \frac{d\tau}{2}. \end{aligned}$$

Combining terms:

$$\boxed{\frac{\partial \varphi}{\partial \tau} = \left(iu(r-q) - \frac{\kappa\bar{v}\rho iu}{\sigma} - \frac{\partial A}{\partial \tau} + \frac{2\kappa\bar{v}}{\sigma^2} \cdot \frac{\partial D}{\partial \tau}\right) \varphi(u;\tau).}$$

Second Derivative of the Characteristic Function

The second derivative is:

$$\frac{\partial^2 \varphi}{\partial \tau^2} = \left(\frac{\partial^2 X}{\partial \tau^2} + \left(\frac{\partial X}{\partial \tau}\right)^2\right) \varphi,$$

where:

$$\frac{\partial^2 X}{\partial \tau^2} = -\frac{\partial^2 A}{\partial \tau^2} + \frac{2\kappa\bar{v}}{\sigma^2} \cdot \frac{\partial^2 D}{\partial \tau^2}.$$

The second derivative of A is:

$$\frac{\partial^2 A}{\partial \tau^2} = \frac{1}{A_2^3}\left[(A_1'' A_2 - A_1 A_2'') A_2 - 2(A_1' A_2 - A_1 A_2') A_2'\right],$$

with:

$$A_1'' = \frac{\partial^2 A_1}{\partial \tau^2} = (u^2 + iu) \cdot \frac{d^2}{4} \sinh \frac{d\tau}{2},$$
$$A_2'' = \frac{\partial^2 A_2}{\partial \tau^2} = \frac{d^3}{4v_0} \cosh \frac{d\tau}{2} + \frac{d^2 \xi}{4v_0} \sinh \frac{d\tau}{2}.$$

The complete expression is:

$$\boxed{\frac{\partial^2 \varphi}{\partial \tau^2} = \left(\frac{\partial^2 X}{\partial \tau^2} + \left(iu(r-q) - \frac{\kappa \bar{v} \rho i u}{\sigma} - \frac{\partial A}{\partial \tau} + \frac{2\kappa \bar{v}}{\sigma^2} \cdot \frac{\partial D}{\partial \tau} \right)^2 \right) \varphi.}$$

Appendix E

Cross Derivatives of the Heston Characteristic Function

This appendix derives the second-order cross-partial derivatives of the Heston characteristic function φ with respect to all distinct pairs of model parameters. These mixed derivatives correspond to the off-diagonal entries of the Hessian matrix used in calibration. We focus on the parameter set $\{v_0, \bar{v}, \rho, \kappa, \sigma\}$, and for each pair (θ_i, θ_j) with $i < j$, we compute $\partial^2\varphi/\partial\theta_i\partial\theta_j$.

Each Section (E.1–E.10) presents a self-contained derivation, following a unified methodology: we write each first derivative as $\partial\varphi/\partial\theta = h_\theta\varphi$, then apply the product rule for the second. Where h_θ itself depends on the second parameter, its derivative is computed explicitly. We note that all mixed derivatives commute under standard smoothness assumptions (Clairaut's theorem), i.e. $\partial^2\varphi/\partial\theta_i\partial\theta_j = \partial^2\varphi/\partial\theta_j\partial\theta_i$.

E.1 Derivation of the Cross Derivative $\frac{\partial^2\varphi}{\partial v_0\,\partial\bar{v}}$

In this section, we derive the second-order cross derivative of the Heston characteristic function φ with respect to the initial variance v_0 and the long-term variance $\bar{v}$.

First, we obtain the first derivative with respect to v_0. The Heston characteristic function φ depends on v_0 through the term $A(v_0, \tau)$ in its exponent. The first derivative is:

$$\frac{\partial\varphi}{\partial v_0} = \left(\frac{\partial A}{\partial v_0}\right)\varphi = -\frac{A}{v_0}\varphi =: h(v_0)\varphi,$$

where $h(v_0) = -\frac{A}{v_0}$ and A is a function of v_0 and other parameters.

Second, we obtain the first derivative with respect to $\bar{v}$. The long-term variance $\bar{v}$ appears in the drift term of the Heston SDE. Differentiating φ with respect to $\bar{v}$ yields

$$\frac{\partial\varphi}{\partial\bar{v}} = \left(-\frac{\kappa\rho\tau iu}{\sigma} + \frac{2\kappa}{\sigma^2}D\right)\varphi =: h(\bar{v})\varphi,$$

where D is a function arising from the Heston PDE solution.

Third, we obtain the second-order cross derivative. Applying the product rule to $\frac{\partial}{\partial \bar{v}}\left(\frac{\partial \varphi}{\partial v_0}\right)$, we get:

$$\frac{\partial^2 \varphi}{\partial v_0 \partial \bar{v}} = \frac{\partial}{\partial \bar{v}}\left(h(v_0)\varphi\right) = \underbrace{\frac{\partial h(v_0)}{\partial \bar{v}}}_{=0} \cdot \varphi + h(v_0) \cdot \frac{\partial \varphi}{\partial \bar{v}}.$$

Since $h(v_0) = -\frac{A}{v_0}$ does not depend on $\bar{v}$, $\frac{\partial h(v_0)}{\partial \bar{v}} = 0$. Thus,

$$\frac{\partial^2 \varphi}{\partial v_0 \partial \bar{v}} = h(v_0) \cdot h(\bar{v}) \cdot \varphi.$$

The final expression is obtained by substituting $h(v_0)$ and $h(\bar{v})$:

$$\frac{\partial^2 \varphi}{\partial v_0 \partial \bar{v}} = \left(-\frac{A}{v_0}\right)\left(-\frac{\kappa \rho \tau i u}{\sigma} + \frac{2\kappa}{\sigma^2} D\right)\varphi.$$

Rearranging the terms produces the following simplified result:

$$\boxed{\frac{\partial^2 \varphi}{\partial v_0 \partial \bar{v}} = \frac{A}{v_0}\left(\frac{\kappa \rho \tau i u}{\sigma} - \frac{2\kappa}{\sigma^2} D\right)\varphi.}$$

E.2 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial v_0 \, \partial \rho}$

First, obtain the derivative of $h(v_0)$ with respect to ρ:

$$\frac{\partial h(v_0)}{\partial \rho} = -\frac{1}{v_0}\frac{\partial A}{\partial \rho}.$$

Second, apply the quotient rule to $A = A_1/A_2$:

$$\frac{\partial A}{\partial \rho} = \frac{\frac{\partial A_1}{\partial \rho} A_2 - A_1 \frac{\partial A_2}{\partial \rho}}{A_2^2}.$$

Third, obtain the component derivatives. The derivative of A_1 is:

$$\begin{aligned} A_1 &= (u^2 + iu)v_0 \sinh\left(\frac{d\tau}{2}\right), \\ \frac{\partial A_1}{\partial \rho} &= (u^2 + iu)v_0 \cdot \frac{\tau}{2} \cosh\left(\frac{d\tau}{2}\right) \frac{\partial d}{\partial \rho}, \\ \frac{\partial d}{\partial \rho} &= -\frac{\xi \sigma i u}{d}, \quad \xi = \kappa - \sigma\rho i u. \end{aligned}$$

The derivative of A_2 is:

$$\begin{aligned} A_2 &= d \cosh\left(\frac{d\tau}{2}\right) + \xi \sinh\left(\frac{d\tau}{2}\right), \\ \frac{\partial A_2}{\partial \rho} &= \left[\cosh\left(\frac{d\tau}{2}\right) + \frac{d\tau}{2}\sinh\left(\frac{d\tau}{2}\right) + \frac{\xi\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\right]\left(-\frac{\xi\sigma i u}{d}\right) \\ &\quad - \sigma i u \sinh\left(\frac{d\tau}{2}\right). \end{aligned}$$

The next step is obtaining the first derivative with respect to ρ:

$$\frac{\partial \varphi}{\partial \rho} = \left(-\frac{\kappa\bar{v}\tau i u}{\sigma} - \frac{\partial A}{\partial \rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \rho}\right)\varphi =: h(\rho)\varphi.$$

Then, obtain the derivative of D:

$$\begin{aligned} D &= \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L, \\ \frac{\partial D}{\partial \rho} &= \left(\frac{1}{d} - \frac{\tau}{2}\right)\left(-\frac{\xi\sigma i u}{d}\right) \\ &\quad - \frac{1}{L}\left[\frac{1}{2v_0}\left(\frac{\partial d}{\partial \rho} - \sigma i u\right) + \frac{1}{2v_0}\left(\frac{\partial d}{\partial \rho} + \sigma i u\right)e^{-d\tau}\right. \\ &\quad \left. - \frac{d - \xi}{2v_0}\tau e^{-d\tau}\frac{\partial d}{\partial \rho}\right]. \end{aligned}$$

The cross derivative is obtained by applying the product rule:

$$\frac{\partial^2 \varphi}{\partial v_0 \partial \rho} = \underbrace{\frac{\partial h(v_0)}{\partial \rho}}_{-\frac{1}{v_0}\frac{\partial A}{\partial \rho}}\varphi + h(v_0) \cdot h(\rho)\varphi.$$

The final result is:

$$\boxed{\frac{\partial^2\varphi}{\partial v_0 \partial\rho} = \left[-\frac{1}{v_0}\frac{\partial A}{\partial\rho} + \left(-\frac{A}{v_0}\right)\left(-\frac{\kappa\bar{v}\tau iu}{\sigma} - \frac{\partial A}{\partial\rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\rho}\right)\right]\varphi.}$$

E.3 Derivation of the Cross Derivative $\frac{\partial^2\varphi}{\partial v_0\,\partial\kappa}$

First, differentiate $h(v_0)$ with respect to κ:

$$\frac{\partial h(v_0)}{\partial\kappa} = -\frac{1}{v_0}\frac{\partial A}{\partial\kappa}.$$

Apply the quotient rule to $\frac{\partial A}{\partial\kappa}$:

$$\frac{\partial A}{\partial\kappa} = \frac{\frac{\partial A_1}{\partial\kappa}A_2 - A_1\frac{\partial A_2}{\partial\kappa}}{A_2^2}.$$

Next, obtain the component derivatives. The derivative of A_1 is:

$$\begin{aligned} A_1 &= (u^2+iu)v_0\sinh\left(\frac{d\tau}{2}\right), \\ \frac{\partial A_1}{\partial\kappa} &= (u^2+iu)v_0\cdot\frac{\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\frac{\partial d}{\partial\kappa}, \\ \frac{\partial d}{\partial\kappa} &= \frac{\xi}{d}, \quad \xi = \kappa - \sigma\rho iu. \end{aligned}$$

The derivative of A_2 is:

$$\begin{aligned} A_2 &= d\cosh\left(\frac{d\tau}{2}\right) + \xi\sinh\left(\frac{d\tau}{2}\right), \\ \frac{\partial A_2}{\partial\kappa} &= \left[\cosh\left(\frac{d\tau}{2}\right) + \frac{\xi\tau}{2}\cosh\left(\frac{d\tau}{2}\right) + \frac{d\tau}{2}\sinh\left(\frac{d\tau}{2}\right)\right]\frac{\xi}{d} \\ &\quad + \sinh\left(\frac{d\tau}{2}\right). \end{aligned}$$

The next step is obtaining the first derivative with respect to κ:

$$\frac{\partial\varphi}{\partial\kappa} = \left(-\frac{\bar{v}\rho\tau iu}{\sigma} - \frac{\partial A}{\partial\kappa} + \frac{2\bar{v}}{\sigma^2}D + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\kappa}\right)\varphi =: h(\kappa)\varphi.$$

Then, the cross derivative via the product rule is:

$$\frac{\partial^2 \varphi}{\partial v_0 \partial \kappa} = \underbrace{\frac{\partial h(v_0)}{\partial \kappa}}_{-\frac{1}{v_0}\frac{\partial A}{\partial \kappa}} \varphi + h(v_0) \cdot h(\kappa)\varphi.$$

Finally, the result is:

$$\boxed{\frac{\partial^2 \varphi}{\partial v_0 \partial \kappa} = \left[-\frac{1}{v_0}\frac{\partial A}{\partial \kappa} + \left(-\frac{A}{v_0}\right)\left(-\frac{\bar{v}\rho\tau i u}{\sigma} - \frac{\partial A}{\partial \kappa} + \frac{2\bar{v}}{\sigma^2}D + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \kappa}\right)\right]\varphi.}$$

E.4 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial v_0 \, \partial \sigma}$

First, obtain the first derivative with respect to v_0. From Itô's lemma, we have:

$$\frac{\partial \varphi}{\partial v_0} = h(v_0) \cdot \varphi, \quad \text{where} \quad h(v_0) = -\frac{A}{v_0}.$$

Second, differentiate $h(v_0)$ with respect to σ:

$$\frac{\partial h(v_0)}{\partial \sigma} = -\frac{1}{v_0}\frac{\partial A}{\partial \sigma}.$$

Applying the quotient rule to $A = \frac{A_1}{A_2}$, we obtain:

$$\frac{\partial A}{\partial \sigma} = \frac{\frac{\partial A_1}{\partial \sigma}A_2 - A_1\frac{\partial A_2}{\partial \sigma}}{A_2^2}.$$

Next, obtain the component derivatives. The derivative of A_1 is:

$$\begin{aligned} A_1 &= (u^2 + iu)v_0 \sinh\left(\frac{d\tau}{2}\right), \\ \frac{\partial A_1}{\partial \sigma} &= (u^2 + iu)v_0 \cdot \frac{\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\frac{\partial d}{\partial \sigma}, \\ \frac{\partial d}{\partial \sigma} &= \frac{\sigma(u^2 + iu) - \rho\xi i u}{d}. \end{aligned}$$

The derivative of A_2 is:

$$\begin{aligned} A_2 &= d\cosh\left(\frac{d\tau}{2}\right) + \xi\sinh\left(\frac{d\tau}{2}\right), \\ \frac{\partial A_2}{\partial\sigma} &= \frac{\partial d}{\partial\sigma}\cosh\left(\frac{d\tau}{2}\right) + d\cdot\frac{\tau}{2}\sinh\left(\frac{d\tau}{2}\right)\frac{\partial d}{\partial\sigma} \\ &\quad + \frac{\partial\xi}{\partial\sigma}\sinh\left(\frac{d\tau}{2}\right) + \xi\cdot\frac{\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\frac{\partial d}{\partial\sigma}, \\ \frac{\partial\xi}{\partial\sigma} &= -\rho i u. \end{aligned}$$

Now, obtain the first derivative with respect to σ:

$$\frac{\partial\varphi}{\partial\sigma} = \left(\frac{\kappa\bar{v}\rho\tau i u}{\sigma^2} - \frac{\partial A}{\partial\sigma} - \frac{4\kappa\bar{v}}{\sigma^3}D + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\sigma}\right)\varphi.$$

The cross derivative via the product rule is:

$$\frac{\partial^2\varphi}{\partial v_0\partial\sigma} = \underbrace{\frac{\partial h(v_0)}{\partial\sigma}}_{-\frac{1}{v_0}\frac{\partial A}{\partial\sigma}}\varphi + h(v_0)\cdot\frac{\partial\varphi}{\partial\sigma}.$$

Finally, the result is:

$$\boxed{\frac{\partial^2\varphi}{\partial v_0\partial\sigma} = \left[-\frac{1}{v_0}\frac{\partial A}{\partial\sigma} - \frac{A}{v_0}\left(\frac{\kappa\bar{v}\rho\tau i u}{\sigma^2} - \frac{\partial A}{\partial\sigma} - \frac{4\kappa\bar{v}}{\sigma^3}D + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\sigma}\right)\right]\varphi.}$$

E.5 Derivation of the Cross Derivative $\frac{\partial^2\varphi}{\partial\bar{v}\,\partial\rho}$

First, obtain the first derivative with respect to $\bar{v}$. From the characteristic function

$$\varphi = \exp\left(iu\ln\frac{F}{S_t} - \frac{\kappa\bar{v}\rho\tau i u}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right),$$

the first derivative is:

$$\frac{\partial\varphi}{\partial\bar{v}} = \left(-\frac{\kappa\rho\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}D\right)\varphi =: h(\bar{v})\varphi.$$

Second, differentiate $h(\bar{v})$ with respect to ρ:

$$\frac{\partial h(\bar{v})}{\partial\rho} = -\frac{\kappa\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}\frac{\partial D}{\partial\rho}.$$

Third, compute $\frac{\partial D}{\partial \rho}$:

$$\begin{aligned}
D &= \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L, \\
L &= \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}e^{-d\tau}, \\
\frac{\partial D}{\partial \rho} &= \left(\frac{1}{d} - \frac{\tau}{2}\right)\left(-\frac{\xi\sigma i u}{d}\right) \\
&\quad - \frac{1}{L}\left[\frac{1}{2v_0}\left(-\frac{\xi\sigma i u}{d} - \sigma i u\right)\right. \\
&\quad + \frac{1}{2v_0}\left(-\frac{\xi\sigma i u}{d} + \sigma i u\right)e^{-d\tau} \\
&\quad \left. - \frac{d-\xi}{2v_0}\tau e^{-d\tau}\left(-\frac{\xi\sigma i u}{d}\right)\right].
\end{aligned}$$

Next, obtain the first derivative with respect to ρ:

$$\frac{\partial \varphi}{\partial \rho} = \left(-\frac{\kappa\bar{v}\tau i u}{\sigma} - \frac{\partial A}{\partial \rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \rho}\right)\varphi.$$

The cross derivative via the product rule is:

$$\frac{\partial^2 \varphi}{\partial \bar{v}\partial \rho} = \frac{\partial h(\bar{v})}{\partial \rho}\varphi + h(\bar{v})\frac{\partial \varphi}{\partial \rho}.$$

Substituting these results, we get:

$$\begin{aligned}
\frac{\partial^2 \varphi}{\partial \bar{v}\partial \rho} &= \left(-\frac{\kappa\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}\frac{\partial D}{\partial \rho}\right)\varphi \\
&\quad + \left(-\frac{\kappa\rho\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}D\right) \\
&\quad \times \left(-\frac{\kappa\bar{v}\tau i u}{\sigma} - \frac{\partial A}{\partial \rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \rho}\right)\varphi.
\end{aligned}$$

Finally, the result is:

$$\boxed{\begin{aligned}\frac{\partial^2\varphi}{\partial\bar{v}\partial\rho} = \Bigg[&-\frac{\kappa\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}\frac{\partial D}{\partial\rho} \\ &+ \left(-\frac{\kappa\rho\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}D\right) \\ &\times \left(-\frac{\kappa\bar{v}\tau i u}{\sigma} - \frac{\partial A}{\partial\rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\rho}\right)\Bigg]\varphi.\end{aligned}} \tag{E.5.1}$$

E.6 Derivation of the Cross Derivative $\frac{\partial^2\varphi}{\partial\bar{v}\,\partial\kappa}$

First, obtain the first derivative with respect to $\bar{v}$. The characteristic function

$$\varphi = \exp\left(iu\ln\frac{F}{S_t} - \frac{\kappa\bar{v}\rho\tau i u}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right)$$

yields

$$\frac{\partial\varphi}{\partial\bar{v}} = \left(-\frac{\kappa\rho\tau i u}{\sigma} + \frac{2\kappa}{\sigma^2}D\right)\varphi =: h(\bar{v})\varphi.$$

Second, differentiate $h(\bar{v})$ with respect to κ:

$$\frac{\partial h(\bar{v})}{\partial\kappa} = -\frac{\rho\tau i u}{\sigma} + \frac{2}{\sigma^2}D + \frac{2\kappa}{\sigma^2}\frac{\partial D}{\partial\kappa}.$$

Third, compute $\frac{\partial D}{\partial\kappa}$:

$$\begin{aligned}D &= \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa-d)\tau}{2} - \ln L, \\ L &= \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0}e^{-d\tau}, \\ \frac{\partial D}{\partial\kappa} &= \frac{\xi}{d^2} + \frac{\tau}{2}\left(1-\frac{\xi}{d}\right) - \frac{1}{L}\frac{\partial L}{\partial\kappa},\end{aligned}$$

where

$$\begin{aligned}\frac{\partial L}{\partial\kappa} = &\frac{1}{2v_0}\left(\frac{\xi}{d}+1\right) + \frac{1}{2v_0}\left(\frac{\xi}{d}-1\right)e^{-d\tau} \\ &- \frac{d-\xi}{2v_0}\tau e^{-d\tau}\frac{\xi}{d}.\end{aligned}$$

Next, obtain the first derivative with respect to κ:

$$\frac{\partial \varphi}{\partial \kappa} = \left(-\frac{\bar{v}\rho\tau i u}{\sigma} - \frac{\partial A}{\partial \kappa} + \frac{2\bar{v}}{\sigma^2} D + \frac{2\kappa\bar{v}}{\sigma^2} \frac{\partial D}{\partial \kappa} \right) \varphi.$$

Apply the quotient rule for A as follows:

$$\frac{\partial A}{\partial \kappa} = \frac{\frac{\partial A_1}{\partial \kappa} A_2 - A_1 \frac{\partial A_2}{\partial \kappa}}{A_2^2},$$

where

$$\begin{aligned}
\frac{\partial A_1}{\partial \kappa} &= (u^2 + iu) v_0 \frac{\tau}{2} \cosh\left(\frac{d\tau}{2}\right) \frac{\xi}{d}, \\
\frac{\partial A_2}{\partial \kappa} &= \left(1 + \frac{\xi\tau}{2}\right) \cosh\left(\frac{d\tau}{2}\right) \frac{\xi}{d} \\
&\quad + \frac{d\tau}{2} \sinh\left(\frac{d\tau}{2}\right) \frac{\xi}{d} + \sinh\left(\frac{d\tau}{2}\right).
\end{aligned}$$

Now, obtain the final cross derivative:

$$\frac{\partial^2 \varphi}{\partial \bar{v} \partial \kappa} = \frac{\partial h(\bar{v})}{\partial \kappa} \varphi + h(\bar{v}) \frac{\partial \varphi}{\partial \kappa}.$$

This yields the complete expression:

$$\boxed{\begin{aligned}
\frac{\partial^2 \varphi}{\partial \bar{v} \partial \kappa} = \Bigg[& -\frac{\rho\tau i u}{\sigma} + \frac{2}{\sigma^2} D + \frac{2\kappa}{\sigma^2} \frac{\partial D}{\partial \kappa} \\
& + \left(-\frac{\kappa\rho\tau i u}{\sigma} + \frac{2\kappa D}{\sigma^2} \right) \\
& \times \left(-\frac{\bar{v}\rho\tau i u}{\sigma} - \frac{\partial A}{\partial \kappa} + \frac{2\bar{v} D}{\sigma^2} + \frac{2\kappa\bar{v}}{\sigma^2} \frac{\partial D}{\partial \kappa} \right) \Bigg] \varphi.
\end{aligned}} \tag{E.6.1}$$

E.7 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \bar{v} \partial \sigma}$

First, obtain the first derivative with respect to $\bar{v}$. The characteristic function

$$\varphi = \exp\left(iu \ln \frac{F}{S_t} - \frac{\kappa\bar{v}\rho\tau i u}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2} D \right)$$

yields

$$\frac{\partial \varphi}{\partial \bar{v}} = \left(-\frac{\kappa \rho \tau i u}{\sigma} + \frac{2\kappa}{\sigma^2} D \right) \varphi =: h(\bar{v}) \varphi.$$

Second, consider the derivative of $h(\bar{v})$ to volatility:

$$\frac{\partial h(\bar{v})}{\partial \sigma} = \frac{\kappa \rho \tau i u}{\sigma^2} - \frac{4\kappa D}{\sigma^3} + \frac{2\kappa}{\sigma^2} \frac{\partial D}{\partial \sigma}. \tag{E.7.1}$$

Third, consider the derivative of the auxiliary term D, which decomposes as

$$\begin{aligned} D &= \ln\left(\frac{d}{v_0}\right) + \frac{(\kappa - d)\tau}{2} - \ln L, \\ L &= \frac{d+\xi}{2v_0} + \frac{d-\xi}{2v_0} e^{-d\tau}, \end{aligned}$$

with the following derivatives:

$$\frac{\partial D}{\partial \sigma} = \left(\frac{1}{d} - \frac{\tau}{2} \right) \frac{\partial d}{\partial \sigma} - \frac{1}{L} \frac{\partial L}{\partial \sigma}.$$

Here,

$$\begin{aligned} \frac{\partial d}{\partial \sigma} &= \frac{\sigma(u^2 + iu) - \rho \xi i u}{d}, \\ \frac{\partial L}{\partial \sigma} &= \frac{1}{2v_0} \left(\frac{\partial d}{\partial \sigma} - \rho i u \right) + \frac{1}{2v_0} \left(\frac{\partial d}{\partial \sigma} + \rho i u \right) e^{-d\tau} \\ &\quad - \frac{d - \xi}{2v_0} \tau e^{-d\tau} \frac{\partial d}{\partial \sigma}. \end{aligned}$$

Next, obtain the first derivative with respect to σ:

$$\frac{\partial \varphi}{\partial \sigma} = \left(\frac{\kappa \bar{v} \rho \tau i u}{\sigma^2} - \frac{\partial A}{\partial \sigma} - \frac{4\kappa \bar{v} D}{\sigma^3} + \frac{2\kappa \bar{v}}{\sigma^2} \frac{\partial D}{\partial \sigma} \right) \varphi. \tag{E.7.2}$$

Now, obtain the complete cross derivative by applying the product rule:

$$\frac{\partial^2 \varphi}{\partial \bar{v} \partial \sigma} = \frac{\partial h(\bar{v})}{\partial \sigma} \varphi + h(\bar{v}) \frac{\partial \varphi}{\partial \sigma}.$$

Substituting Equations (E.7.1) and (E.7.2), we get:

$$\boxed{\begin{aligned}\frac{\partial^2\varphi}{\partial\bar{v}\partial\sigma} = \Bigg[&\frac{\kappa\rho\tau iu}{\sigma^2} - \frac{4\kappa D}{\sigma^3} + \frac{2\kappa}{\sigma^2}\frac{\partial D}{\partial\sigma} \\ &+ \left(-\frac{\kappa\rho\tau iu}{\sigma} + \frac{2\kappa D}{\sigma^2}\right) \\ &\times \left(\frac{\kappa\bar{v}\rho\tau iu}{\sigma^2} - \frac{\partial A}{\partial\sigma} - \frac{4\kappa\bar{v}D}{\sigma^3} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\sigma}\right)\Bigg]\varphi.\end{aligned}}$$

Note the auxiliary calculation of $\frac{\partial A}{\partial\sigma}$. Using the quotient rule $A = A_1/A_2$, we get

$$\frac{\partial A}{\partial\sigma} = \frac{\frac{\partial A_1}{\partial\sigma}A_2 - A_1\frac{\partial A_2}{\partial\sigma}}{A_2^2},$$

where

$$\begin{aligned}\frac{\partial A_1}{\partial\sigma} &= (u^2 + iu)v_0\frac{\tau}{2}\cosh\left(\frac{d\tau}{2}\right)\frac{\partial d}{\partial\sigma}, \\ \frac{\partial A_2}{\partial\sigma} &= \left(1 + \frac{\xi\tau}{2} + \frac{d\tau}{2}\tanh\left(\frac{d\tau}{2}\right)\right)\frac{\partial d}{\partial\sigma}\cosh\left(\frac{d\tau}{2}\right) \\ &\quad - \rho iu\sinh\left(\frac{d\tau}{2}\right).\end{aligned}$$

E.8 Derivation of the Cross Derivative $\frac{\partial^2\varphi}{\partial\rho\,\partial\kappa}$

First, obtain the first derivative with respect to ρ. The characteristic function

$$\varphi = \exp\left(iu\ln\frac{F}{S_t} - \frac{\kappa\bar{v}\rho\tau iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right)$$

yields

$$\frac{\partial\varphi}{\partial\rho} = \left(-\frac{\kappa\bar{v}\tau iu}{\sigma} - \frac{\partial A}{\partial\rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\rho}\right)\varphi =: h(\rho)\varphi.$$

Second, obtain the mixed derivative of $h(\rho)$:

$$\frac{\partial h(\rho)}{\partial\kappa} = -\frac{\bar{v}\tau iu}{\sigma} - \frac{\partial^2 A}{\partial\rho\partial\kappa} + \frac{2\bar{v}}{\sigma^2}\left(\frac{\partial D}{\partial\rho} + \kappa\frac{\partial^2 D}{\partial\rho\partial\kappa}\right). \tag{E.8.1}$$

Third, obtain the second mixed derivative of A. For $A = A_1/A_2$:

$$\frac{\partial^2 A}{\partial\rho\partial\kappa} = \frac{1}{A_2^3}\left[\begin{array}{l}\left(\dfrac{\partial^2 A_1}{\partial\rho\partial\kappa}A_2 + \dfrac{\partial A_1}{\partial\rho}\dfrac{\partial A_2}{\partial\kappa} - \dfrac{\partial A_1}{\partial\kappa}\dfrac{\partial A_2}{\partial\rho} - A_1\dfrac{\partial^2 A_2}{\partial\rho\partial\kappa}\right)A_2 \\ -2\left(\dfrac{\partial A_1}{\partial\rho}A_2 - A_1\dfrac{\partial A_2}{\partial\rho}\right)\dfrac{\partial A_2}{\partial\kappa}\end{array}\right].$$

Next, obtain the second mixed derivative of D:

$$\frac{\partial^2 D}{\partial\rho\partial\kappa} = \frac{\partial}{\partial\kappa}\left[\left(\frac{1}{d} - \frac{\tau}{2}\right)\frac{\partial d}{\partial\rho} - \frac{1}{L}\frac{\partial L}{\partial\rho}\right].$$

This requires the computation of:

$$\begin{aligned}\frac{\partial^2 d}{\partial\rho\partial\kappa} &= \frac{\partial}{\partial\kappa}\left(-\frac{\xi\sigma i u}{d}\right), \\ \frac{\partial^2 L}{\partial\rho\partial\kappa} &= \frac{\partial}{\partial\kappa}\left[\frac{1}{2v_0}\left(\frac{\partial d}{\partial\rho} + \frac{\partial\xi}{\partial\rho}\right) + \cdots\right].\end{aligned}$$

The next step is obtaining the first derivative with respect to κ:

$$\frac{\partial\varphi}{\partial\kappa} = \left(-\frac{\bar{v}\rho\tau i u}{\sigma} - \frac{\partial A}{\partial\kappa} + \frac{2\bar{v}}{\sigma^2}D + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\kappa}\right)\varphi. \tag{E.8.2}$$

Now, obtain the complete cross derivative:

$$\frac{\partial^2\varphi}{\partial\rho\partial\kappa} = \frac{\partial h(\rho)}{\partial\kappa}\varphi + h(\rho)\frac{\partial\varphi}{\partial\kappa}.$$

Substituting Equations (E.8.1) and (E.8.2), we get:

$$\boxed{\begin{aligned}\frac{\partial^2\varphi}{\partial\rho\partial\kappa} = \Bigg[&-\frac{\bar{v}\tau i u}{\sigma} - \frac{\partial^2 A}{\partial\rho\partial\kappa} + \frac{2\bar{v}}{\sigma^2}\left(\frac{\partial D}{\partial\rho} + \kappa\frac{\partial^2 D}{\partial\rho\partial\kappa}\right) \\ &+ \left(-\frac{\kappa\bar{v}\tau i u}{\sigma} - \frac{\partial A}{\partial\rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\rho}\right) \\ &\times\left(-\frac{\bar{v}\rho\tau i u}{\sigma} - \frac{\partial A}{\partial\kappa} + \frac{2\bar{v}D}{\sigma^2} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\kappa}\right)\Bigg]\varphi.\end{aligned}}$$

E.9 Derivation of the Cross Derivative $\frac{\partial^2\varphi}{\partial\rho\,\partial\sigma}$

First, obtain the first derivative with respect to ρ. The characteristic function

$$\varphi = \exp\left(iu\ln\frac{F}{S_t} - \frac{\kappa\bar{v}\rho\tau iu}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right)$$

yields

$$\frac{\partial\varphi}{\partial\rho} = \left(-\frac{\kappa\bar{v}\tau iu}{\sigma} - \frac{\partial A}{\partial\rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\rho}\right)\varphi =: h(\rho)\varphi.$$

Second, obtain the mixed derivative of $h(\rho)$ with respect to σ:

$$\frac{\partial h(\rho)}{\partial\sigma} = \frac{\kappa\bar{v}\tau iu}{\sigma^2} - \frac{\partial^2 A}{\partial\rho\partial\sigma} - \frac{4\kappa\bar{v}}{\sigma^3}\frac{\partial D}{\partial\rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial^2 D}{\partial\rho\partial\sigma}. \tag{E.9.1}$$

Third, obtain the second mixed derivative of A. For $A = A_1/A_2$:

$$\frac{\partial^2 A}{\partial\rho\partial\sigma} = \frac{1}{A_2^3}\left[\begin{array}{l}\left(\frac{\partial^2 A_1}{\partial\rho\partial\sigma}A_2 + \frac{\partial A_1}{\partial\rho}\frac{\partial A_2}{\partial\sigma} - \frac{\partial A_1}{\partial\sigma}\frac{\partial A_2}{\partial\rho} - A_1\frac{\partial^2 A_2}{\partial\rho\partial\sigma}\right)A_2 \\ -2\left(\frac{\partial A_1}{\partial\rho}A_2 - A_1\frac{\partial A_2}{\partial\rho}\right)\frac{\partial A_2}{\partial\sigma}\end{array}\right].$$

Next, obtain the second mixed derivative of D:

$$\frac{\partial^2 D}{\partial\rho\partial\sigma} = \frac{\partial}{\partial\sigma}\left[\left(\frac{1}{d} - \frac{\tau}{2}\right)\frac{\partial d}{\partial\rho} - \frac{1}{L}\frac{\partial L}{\partial\rho}\right].$$

we compute $\frac{\partial L}{\partial\rho}$ as follows:

$$\begin{aligned}\frac{\partial L}{\partial\rho} &= \frac{1}{2v_0}\left[\frac{\partial}{\partial\rho}(d+\xi) + \frac{\partial}{\partial\rho}\left((d-\xi)e^{-d\tau}\right)\right] \\ &= \frac{1}{2v_0}\left[\left(\frac{\partial d}{\partial\rho} + \frac{\partial\xi}{\partial\rho}\right) + \left(\frac{\partial d}{\partial\rho} - \frac{\partial\xi}{\partial\rho}\right)e^{-d\tau} + (d-\xi)\frac{\partial}{\partial\rho}\left(e^{-d\tau}\right)\right] \\ &= \frac{1}{2v_0}\left[\left(\frac{\partial d}{\partial\rho} + \frac{\partial\xi}{\partial\rho}\right) + \left(\frac{\partial d}{\partial\rho} - \frac{\partial\xi}{\partial\rho}\right)e^{-d\tau} - \tau(d-\xi)\frac{\partial d}{\partial\rho}e^{-d\tau}\right]\end{aligned}$$

Substituting Equations (C.2.2) into the expression:

$$\frac{\partial L}{\partial \rho} = \frac{1}{2v_0}\left[\left(-\frac{\sigma i u \xi}{d} - \sigma i u\right) + \left(-\frac{\sigma i u \xi}{d} + \sigma i u\right) e^{-d\tau} - \tau(d-\xi)\left(-\frac{\sigma i u \xi}{d}\right) e^{-d\tau}\right] \tag{E.9.2}$$

$$= \frac{-\sigma i u}{2v_0}\left[\left(\frac{\xi}{d}+1\right) + \left(\frac{\xi}{d}-1\right) e^{-d\tau} + \tau(d-\xi)\frac{\xi}{d} e^{-d\tau}\right] \tag{E.9.3}$$

This requires the computation of:

$$\frac{\partial^2 d}{\partial \rho \partial \sigma} = \frac{\partial}{\partial \sigma}\left(-\frac{\xi \sigma i u}{d}\right),$$
$$\frac{\partial^2 L}{\partial \rho \partial \sigma} = \frac{\partial}{\partial \sigma}\left[\frac{1}{2v_0}\left(\frac{\partial d}{\partial \rho} - \rho i u\right) + \cdots\right].$$

Next, we obtain the first derivative with respect to σ, as previously defined in Equation (E.7.2). Now, obtain the complete cross derivative:

$$\frac{\partial^2 \varphi}{\partial \rho \partial \sigma} = \frac{\partial h(\rho)}{\partial \sigma}\varphi + h(\rho)\frac{\partial \varphi}{\partial \sigma}.$$

Substituting Equations (E.9.1) and (E.7.2), we get:

$$\boxed{\begin{aligned}\frac{\partial^2 \varphi}{\partial \rho \partial \sigma} = \Bigg[&\frac{\kappa \bar{v} \tau i u}{\sigma^2} - \frac{\partial^2 A}{\partial \rho \partial \sigma} - \frac{4\kappa\bar{v}}{\sigma^3}\frac{\partial D}{\partial \rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial^2 D}{\partial \rho \partial \sigma} \\ &+ \left(-\frac{\kappa\bar{v}\tau i u}{\sigma} - \frac{\partial A}{\partial \rho} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \rho}\right) \\ &\times \left(\frac{\kappa\bar{v}\rho\tau i u}{\sigma^2} - \frac{\partial A}{\partial \sigma} - \frac{4\kappa\bar{v}D}{\sigma^3} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \sigma}\right)\Bigg]\varphi.\end{aligned}}$$

E.10 Derivation of the Cross Derivative $\frac{\partial^2 \varphi}{\partial \kappa \, \partial \sigma}$

First, obtain the first derivative with respect to κ. The characteristic function

$$\varphi = \exp\left(i u \ln\frac{F}{S_t} - \frac{\kappa\bar{v}\rho\tau i u}{\sigma} - A + \frac{2\kappa\bar{v}}{\sigma^2}D\right)$$

yields

$$\frac{\partial \varphi}{\partial \kappa} = \left(-\frac{\bar{v}\rho\tau i u}{\sigma} - \frac{\partial A}{\partial \kappa} + \frac{2\bar{v}}{\sigma^2}D + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial \kappa}\right)\varphi =: h(\kappa)\varphi.$$

Second, obtain the mixed derivative of $h(\kappa)$:

$$\frac{\partial h(\kappa)}{\partial \sigma} = \frac{\bar{v}\rho\tau iu}{\sigma^2} - \frac{\partial^2 A}{\partial\kappa\partial\sigma} - \frac{4\bar{v}D}{\sigma^3} + \frac{2\bar{v}(1-2\kappa/\sigma)}{\sigma^2}\frac{\partial D}{\partial\kappa} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial^2 D}{\partial\kappa\partial\sigma}. \tag{E.10.1}$$

Third, obtain the second mixed derivative of A. For $A = A_1/A_2$:

$$\frac{\partial^2 A}{\partial\kappa\partial\sigma} = \frac{1}{A_2^3}\left[\begin{array}{l}\left(\dfrac{\partial^2 A_1}{\partial\kappa\partial\sigma}A_2 + \dfrac{\partial A_1}{\partial\kappa}\dfrac{\partial A_2}{\partial\sigma} - \dfrac{\partial A_1}{\partial\sigma}\dfrac{\partial A_2}{\partial\kappa} - A_1\dfrac{\partial^2 A_2}{\partial\kappa\partial\sigma}\right)A_2 \\ -2\left(\dfrac{\partial A_1}{\partial\kappa}A_2 - A_1\dfrac{\partial A_2}{\partial\kappa}\right)\dfrac{\partial A_2}{\partial\sigma}\end{array}\right].$$

Obtain the second mixed derivative of D:

$$\frac{\partial^2 D}{\partial\kappa\partial\sigma} = \frac{\partial}{\partial\sigma}\left[\left(\frac{1}{d} - \frac{\tau}{2}\right)\frac{\xi}{d} - \frac{1}{L}\frac{\partial L}{\partial\kappa}\right],$$

where

$$\frac{\partial^2 d}{\partial\kappa\partial\sigma} = \frac{\partial}{\partial\sigma}\left(\frac{\xi}{d}\right),$$
$$\frac{\partial^2 L}{\partial\kappa\partial\sigma} = \frac{\partial}{\partial\sigma}\left[\frac{1}{2v_0}\left(1 + \frac{\xi}{d}\right) + \cdots\right].$$

Next, we refer to the first derivative with respect to σ, already derived in Equation (E.7.2). Now, obtain the complete cross derivative:

$$\frac{\partial^2\varphi}{\partial\kappa\partial\sigma} = \frac{\partial h(\kappa)}{\partial\sigma}\varphi + h(\kappa)\frac{\partial\varphi}{\partial\sigma}.$$

Substituting Equations (E.10.1) and (E.7.2), we get:

$$\boxed{\begin{aligned}\frac{\partial^2\varphi}{\partial\kappa\partial\sigma} = \Bigg[&\frac{\bar{v}\rho\tau iu}{\sigma^2} - \frac{\partial^2 A}{\partial\kappa\partial\sigma} - \frac{4\bar{v}D}{\sigma^3} + \frac{2\bar{v}(1-2\kappa/\sigma)}{\sigma^2}\frac{\partial D}{\partial\kappa} \\ &+ \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial^2 D}{\partial\kappa\partial\sigma} + \left(-\frac{\bar{v}\rho\tau iu}{\sigma} - \frac{\partial A}{\partial\kappa} + \frac{2\bar{v}D}{\sigma^2} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\kappa}\right) \\ &\times\left(\frac{\kappa\bar{v}\rho\tau iu}{\sigma^2} - \frac{\partial A}{\partial\sigma} - \frac{4\kappa\bar{v}D}{\sigma^3} + \frac{2\kappa\bar{v}}{\sigma^2}\frac{\partial D}{\partial\sigma}\right)\Bigg]\varphi.\end{aligned}} \tag{E.10.2}$$

Appendix F

Supplementary Data Tables

This appendix compiles supplementary data tables that support the empirical analyses presented in the thesis. These tables provide detailed empirical evidence used for validating the Heston model calibration method, as well as contextualizing model performance relative to market conditions. The clearly organized tables and explanatory notes facilitate quick referencing and verification of empirical findings, ensuring transparency and thoroughness in the presentation of results.

F.1 Correlation Analysis Tables

The following table presents complete correlation results between key calibration outputs and market observables under both stable and unstable market regimes. These correlations complement the analysis presented in Section 5.10.

Table F.1.1: Correlation results for stable and unstable data (showing $|\texttt{corr}| \geq 0.30$)

Variable 1	Variable 2	Stable	Unstable
ImpliedUnderlyingPrice	FitWellUnderlyingPrice	1.00	1.00
ImpliedOptionPrice	FitWellOptionPrice	1.00	1.00
Underlying Price	ImpliedUnderlyingPrice	1.00	1.00
Underlying Price	FitWellUnderlyingPrice	1.00	1.00
GreekBLS_5	FitWellGreekBLS_5	1.00	1.00
GreekBLS_1	FitWellGreekBLS_1	1.00	1.00
RiskFreeRate	FitWellRiskFreeRate	1.00	1.00
DividendYield	FitWellDividendYield	1.00	1.00
Mid	ImpliedOptionPrice	1.00	1.00
Ask	ImpliedOptionPrice	1.00	1.00
Mid	Ask	1.00	1.00
Bid	Mid	1.00	1.00
Mid	FitWellOptionPrice	1.00	1.00
Ask	FitWellOptionPrice	1.00	1.00
Bid	ImpliedOptionPrice	1.00	0.99
Bid	FitWellOptionPrice	1.00	0.99
Bid	Ask	1.00	0.99
FitWellGreek_1	FitWellGreekBLS_1	1.00	0.99
FitWellGreek_5	FitWellGreekBLS_5	1.00	0.99
GreekBLS_1	FitWellGreek_1	1.00	0.99
GreekBLS_5	FitWellGreek_5	1.00	0.99

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
GreekBLS_4	FitWellGreekBLS_4	1.00	0.96
FitWellGreek_4	FitWellGreekBLS_4	0.965	0.81
FitWellGreek_4	GreekBLS_4	0.99	0.82
Last	ImpliedOptionPrice	0.99	0.97
Last	FitWellOptionPrice	0.99	0.97
Ask	Last	0.99	0.97
Mid	Last	0.99	0.97
Bid	Last	0.99	0.96
Bid	Intrinsic Value	0.99	0.99
Mid	Intrinsic Value	0.99	0.99
Intrinsic Value	ImpliedOptionPrice	0.99	0.99
Intrinsic Value	FitWellOptionPrice	0.99	0.99
Ask	Intrinsic Value	0.99	0.99
Last	Intrinsic Value	0.98	0.96
GreekBLS_2	FitWellGreekBLS_2	0.96	0.91
GreekBLS_3	FitWellGreekBLS_3	0.64	0.95
IV	FitWellIV	0.94	0.75
FitWellGreek_4	FitWellGradiant_2	0.94	0.92
FitWellGreek_3	FitWellGreeksBS_3	0.41	0.95
FitWellGreek_3	GreeksBS_3	0.44	0.94
FitWellGradiant_2	FitWellGreekBLS_4	0.93	0.60
GreekBLS_4	FitWellGradiant_2	0.92	0.57
FitWellGreek_3	FitWellGreekBLS_1	0.05	0.92
FitWellGreek_1	FitWellGreek_3	0.03	0.91
GreekBLS_1	FitWellGreek_3	-0.05	-0.91
GreekBLS_3	FitWellGreek_1	0.06	0.88
FitWellGreek_1	FitWellGreekBLS_3	0.08	0.88
GreekBLS_3	FitWellGreekBLS_1	0.07	0.87
FitWellGreek_4	Extrinsic Value	0.86	0.50
Extrinsic Value	FitWellGreekBLS_4	0.85	0.51
GreekBLS_1	GreekBLS_3	0.07	0.85
FitWellGreekBLS_1	FitWellGreekBLS_3	0.09	0.85
Intrinsic Value	FitWellGreekBLS_1	0.31	-0.85
Intrinsic Value	FitWellGreek_1	0.30	-0.85
Extrinsic Value	GreekBLS_4	0.84	0.43
GreekBLS_1	FitWellGreekBLS_3	0.09	0.84
Intrinsic Value	GreekBLS_1	0.31	-0.84
Last	FitWellGreekBLS_1	0.33	-0.83
Last	FitWellGreek_1	0.33	-0.83
Ask	FitWellGreek_1	0.33	-0.83
Bid	IV	0.24	0.83
Intrinsic Value	FitWellGreek_3	0.17	0.83
ImpliedOptionPrice	FitWellGreekBLS_1	0.33	-0.82
ImpliedOptionPrice	FitWellGreek_1	0.33	-0.82
Mid	FitWellGreekBLS_1	0.33	-0.82
Ask	FitWellGreekBLS_1	0.33	-0.82
Bid	FitWellGreek_1	0.33	-0.82
Mid	FitWellGreek_1	0.33	-0.82
FitWellOptionPrice	FitWellGreek_1	0.33	-0.82
Mid	IV	0.24	0.82
FitWellParameters_2	FitWellGradiant_4	0.81	0.62
IV	ImpliedOptionPrice	0.24	0.81
IV	FitWellOptionPrice	0.24	0.81
Intrinsic Value	IV	0.27	0.81
Ask	IV	0.23	0.81
Last	GreekBLS_1	0.34	-0.81
Ask	GreekBLS_1	0.34	-0.81
Bid	FitWellGreekBLS_1	0.34	-0.81
FitWellOptionPrice	FitWellGreekBLS_1	0.33	-0.81
Extrinsic Value	FitWellGradiant_2	0.80	0.36

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
Bid	GreekBLS_1	0.34	-0.80
ImpliedOptionPrice	GreekBLS_1	0.34	-0.80
Mid	GreekBLS_1	0.34	-0.80
FitWellOptionPrice	GreekBLS_1	0.34	-0.80
Last	FitWellGreek_3	0.15	0.80
Ask	FitWellGreek_3	0.15	0.80
Intrinsic Value	FitWellGreekBLS_5	0.17	-0.80
Last	IV	0.24	0.79
FitWellGreek_1	FitWellGreekBLS_5	0.67	0.79
GreekBLS_5	FitWellGreek_1	0.67	0.79
FitWellGradiant_3	FitWellGreekBLS_4	0.32	0.79
Mid	FitWellGreek_3	0.15	0.79
ImpliedOptionPrice	FitWellGreek_3	0.15	0.79
Layer2Residual	FitWellResidual	0.01	0.79
FitWellOptionPrice	FitWellGreek_3	0.15	0.79
Intrinsic Value	GreekBLS_5	0.16	-0.79
Bid	FitWellGreek_3	0.15	0.79
FitWellParameters_4	FitWellParameters_5	0.78	-0.38
GreekBLS_1	FitWellGreekBLS_5	0.67	0.78
GreekBLS_1	GreekBLS_5	0.67	0.78
FitWellGreekBLS_1	FitWellGreekBLS_5	0.67	0.78
FitWellGreekBLS_1	GreekBLS_5	0.67	0.78
Ask	FitWellGreekBLS_5	0.19	-0.78
Last	GreekBLS_3	0.02	-0.78
Mid	FitWellGreekBLS_5	0.19	-0.78
Intrinsic Value	FitWellGreek_5	0.16	-0.78
FitWellGradiant_1	FitWellGradiant_5	-0.77	-0.25
FitWellFeller	FitWellParameters_5	-0.77	-0.42
FitWellGreek_1	FitWellGreek_5	0.67	0.77
Ask	GreekBLS_5	0.20	-0.77
Intrinsic Value	Layer2Residual	0.09	0.77
ImpliedOptionPrice	FitWellGreekBLS_5	0.19	-0.77
Last	Layer2Residual	0.09	0.77
Mid	GreekBLS_5	0.19	-0.77
Bid	FitWellGreekBLS_5	0.2	-0.77
Ask	Layer2Residual	0.083	0.77
FitWellOptionPrice	FitWellGreekBLS_5	0.19	-0.77
Mid	Layer2Residual	0.08	0.77
ImpliedOptionPrice	GreekBLS_5	0.19	-0.77
Intrinsic Value	GreekBLS_3	0.01	-0.77
Ask	FitWellGreek_5	0.20	-0.76
IV	FitWellGreek_3	0.17	0.76
Bid	GreekBLS_5	0.19	-0.76
ImpliedOptionPrice	Layer2Residual	0.09	0.76
Layer2Residual	FitWellGreekBLS_5	0.017	-0.76
Mid	GreekBLS_3	0.02	-0.76
Bid	Layer2Residual	0.08	0.76
Bid	GreekBLS_3	0.02	-0.76
Last	FitWellGreekBLS_5	0.19	-0.76
FitWellOptionPrice	GreekBLS_5	0.19	-0.76
Ask	GreekBLS_3	0.019	-0.76
Mid	FitWellGreek_5	0.20	-0.76
ImpliedOptionPrice	GreekBLS_3	-0.02	-0.76
Last	GreekBLS_5	0.19	-0.76
ImpliedOptionPrice	FitWellGreek_5	0.19	-0.76
Layer2Residual	GreekBLS_5	0.017	-0.76
Layer2Residual	FitWellOptionPrice	0.02	0.75
Layer2Residual	FitWellGreek_1	0.046	-0.75
FitWellOptionPrice	GreekBLS_3	0.02	-0.75
Bid	FitWellGreek_5	0.196	-0.75

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
Layer2Residual	FitWellGreekBLS_3	-0.034	-0.75
Layer2Residual	FitWellGreek_5	0.017	-0.75
FitWellOptionPrice	FitWellGreek_5	0.19	-0.75
GreekBLS_1	FitWellGreek_5	0.67	0.75
Last	FitWellGreek_5	0.19	-0.75
Last	FitWellGreekBLS_3	-0.15	-0.75
GreekBLS_4	FitWellGradiant_3	0.31	0.74
Layer2Residual	FitWellGreekBLS_1	0.048	-0.74
Intrinsic Value	FitWellGreekBLS_3	0.11	-0.73
IV	GreekBLS_3	-0.54	-0.73
Ask	FitWellGreekBLS_3	0.11	-0.73
Mid	FitWellGreekBLS_3	-0.15	-0.72
Layer2Residual	GreekBLS_1	0.05	-0.72
ImpliedOptionPrice	FitWellGreekBLS_3	0.11	-0.72
Layer2Residual	GreekBLS_3	-0.009	-0.72
Bid	FitWellGreekBLS_3	0.11	-0.72
FitWellGradiant_3	FitWellGradiant_5	-0.51	-0.72
Layer2Residual	FitWellGreek_3	-0.03	0.72
FitWellOptionPrice	FitWellGreekBLS_3	0.11	-0.72
IV	FitWellGreekBLS_1	0.12	-0.70
Strike	LogMoneyness	0.69	0.57
FitWellGreek_4	FitWellGradiant_5	-0.69	-0.33
IV	FitWellGreek_1	0.11	-0.69
IV	GreekBLS_1	0.12	-0.68
FitWellGradiant_5	FitWellGreekBLS_4	-0.67	-0.68
FitWellGreek_2	FitWellGreekBLS_2	0.67	0.076
FitWellGreek_5	FitWellGreekBLS_1	0.67	0.75
GreekBLS_4	FitWellGradiant_5	-0.67	-0.63
GreekBLS_2	FitWellGreek_2	0.65	0.069
FitWellGreek_3	FitWellGreekBLS_5	0.03	0.65
GreekBLS_5	FitWellGreek_3	0.034	0.65
IV	Layer2Residual	0.067	0.64
FitWellGreekBLS_3	FitWellGreekBLS_5	-0.03	-0.66
FitWellGreek_5	FitWellGreekBLS_3	0.05	0.64
LogMoneyness	FitWellUnderlyingPrice	-0.25	-0.63
FitWellResidual	FitWellGreekBLS_3	-0.12	-0.63
GreekBLS_5	FitWellGreekBLS_3	0.050	0.63
LogMoneyness	ImpliedUnderlyingPrice	-0.25	-0.63
FitWellGreek_3	FitWellGreek_5	-0.03	-0.62
IV	FitWellGreekBLS_3	-0.293	-0.62
Underlying Price	LogMoneyness	-0.25	-0.62
FitWellGradiant_1	FitWellGradiant_4	0.54	0.61
Days To Expiration	GreekBLS_4	0.60	0.53
Last	FitWellResidual	-0.04	0.60
Ask	FitWellResidual	-0.04	0.60
Mid	FitWellResidual	-0.04	0.60
Intrinsic Value	FitWellResidual	-0.03	0.60
ImpliedOptionPrice	FitWellResidual	-0.04	0.60
FitWellResidual	FitWellOptionPrice	-0.04	0.60
Days To Expiration	FitWellGreekBLS_4	0.59	0.55
FitWellGreek_4	Days To Expiration	0.59	0.38
Bid	FitWellResidual	-0.04	0.59
Ask	BidAskSpread	0.58	0.24
FitWellParameters_2	FitWellParameters_4	-0.11	-0.58
Days To Expiration	FitWellGradiant_2	0.57	0.20
Last	BidAskSpread	0.57	0.23
Mid	BidAskSpread	0.57	0.184
BidAskSpread	ImpliedOptionPrice	0.57	0.23
BidAskSpread	FitWellOptionPrice	0.57	0.23
FitWellGreek_4	FitWellGradiant_1	0.57	0.14

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
GreekBLS_3	FitWellGreekBLS_5	0.03	0.57
FitWellResidual	FitWellGreek_5	0.02	-0.57
FitWellResidual	FitWellGreekBLS_5	0.02	-0.57
FitWellResidual	GreekBLS_5	0.02	-0.57
Extrinsic Value	FitWellGradiant_5	-0.56	-0.45
FitWellGradiant_4	FitWellGradiant_5	-0.56	-0.52
RiskFreeRate	FitWellParameters_5	-0.56	0.25
FitWellParameters_5	FitWellRiskFreeRate	-0.56	0.25
GreekBLS_3	GreekBLS_5	0.031	0.56
GreekBLS_3	FitWellGreek_5	0.036	0.56
FitWellResidual	FitWellGreek_1	-0.02	-0.56
FitWellParameters_2	FitWellGradiant_2	-0.14	-0.56
Bid	BidAskSpread	0.55	0.13
FitWellGradiant_1	FitWellGreekBLS_4	0.55	0.23
GreekBLS_4	FitWellGradiant_1	0.55	0.237
LogMoneyness	FitWellGradiant_3	0.54	0.04
FitWellGreek_2	FitWellGreek_3	-0.56	-0.16
FitWellResidual	FitWellGreek_3	0.073	0.54
FitWellResidual	GreekBLS_3	-0.102	-0.54
Extrinsic Value	FitWellGreek_3	-0.13	-0.54
Intrinsic Value	BidAskSpread	0.53	0.24
FitWellIV	FitWellGreekBLS_2	-0.12	-0.53
Extrinsic Value	GreekBLS_1	0.12	0.53
FitWellResidual	FitWellGreekBLS_1	-0.007	-0.53
FitWellResidual	GreekBLS_1	-0.007	-0.53
Bid Size	Ask Size	0.41	0.52
Days To Expiration	FitWellGradiant_3	0.11	0.52
FitWellGradiant_3	FitWellGradiant_5	-0.51	-0.72
Extrinsic Value	BidAskSpread	0.17	-0.51
RiskFreeRate	DividendYield	-0.31	-0.51
DividendYield	FitWellRiskFreeRate	-0.31	-0.51
FitWellFeller	FitWellParameters_4	-0.50	0.29
FitWellGradiant_2	FitWellGradiant_5	-0.50	-0.052
Extrinsic Value	FitWellGreekBLS_1	0.12	0.50
FitWellParameters_5	FitWellGradiant_2	0.035	-0.50
Layer1Residual	FitLayerUsed	-0.014	-0.50
FitWellParameters_4	FitWellGradiant_2	0.16	0.50
LogMoneyness	FitWellGreekBLS_1	-0.49	-0.143
LogMoneyness	FitWellGreek_1	-0.49	-0.14
LogMoneyness	GreekBLS_1	-0.49	-0.152
Extrinsic Value	FitWellGreek_1	0.14	0.49
Underlying Price	FitWellIV	0.073	-0.49
IV	FitWellResidual	0.05	0.49
FitWellUnderlyingPrice	FitWellIV	0.074	-0.49
ImpliedUnderlyingPrice	FitWellIV	0.074	-0.49
IV	FitWellGreekBLS_5	0.05	-0.49
Strike	Underlying Price	0.48	0.264
Strike	ImpliedUnderlyingPrice	0.48	0.252
Strike	FitWellUnderlyingPrice	0.48	0.245
FitWellIV	GreekBLS_3	-0.48	-0.242
DividendYield	FitWellUnderlyingPrice	0.48	0.023
DividendYield	ImpliedUnderlyingPrice	0.48	0.02
Underlying Price	DividendYield	0.48	0.019
FitWellOptionPrice	FitWellIV	0.16	0.48
RiskFreeRate	FitWellDividendYield	-0.3	-0.48
IV	GreekBLS_5	0.047	-0.48
FitWellRiskFreeRate	FitWellDividendYield	-0.3	-0.48
FitWellGradiant_3	FitWellGradiant_4	0.48	0.48
ImpliedOptionPrice	FitWellIV	0.16	0.48
Intrinsic Value	FitWellIV	0.18	0.48

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
Bid	FitWellIV	0.159	0.48
ImpliedUnderlyingPrice	FitWellDividendYield	0.47	0.016
FitWellUnderlyingPrice	FitWellDividendYield	0.47	0.02
Underlying Price	FitWellDividendYield	0.47	0.016
Mid	FitWellIV	0.16	0.47
Ask	FitWellIV	0.16	0.47
FitWellIV	GreekBLS_2	-0.117	-0.47
Days To Expiration	Extrinsic Value	0.46	0.295
GreekBLS_4	FitWellGreek_5	-0.042	-0.46
Layer1Residual	FitWellGradiant_3	-0.064	-0.46
IV	FitWellGreek_5	0.04	-0.46
LogMoneyness	FitWellGreekBLS_2	0.179	-0.45
Last	FitWellIV	0.16	0.45
RiskFreeRate	FitWellFeller	0.44	-0.33
FitWellParameters_2	FitWellParameters_3	0.133	-0.44
FitWellParameters_4	FitWellGradiant_1	-0.294	-0.44
FitWellFeller	FitWellRiskFreeRate	0.43	-0.33
FitWellParameters_4	FitWellRiskFreeRate	-0.40	-0.43
RiskFreeRate	FitWellParameters_4	-0.40	-0.43
Days To Expiration	FitWellGradiant_5	-0.33	-0.43
Strike	FitWellDividendYield	0.42	0.39
Days To Expiration	FitWellGradiant_4	0.201	0.42
LogMoneyness	FitWellParameters_3	0.035	0.42
Volume	FitWellGreek_2	0.41	0.003
Extrinsic Value	FitWellGradiant_1	0.41	0.034
FitWellGradiant_3	FitWellGreekBLS_5	-0.41	-0.182
GreekBLS_5	FitWellGradiant_3	-0.41	-0.182
FitWellGreek_5	FitWellGradiant_3	-0.41	-0.260
Strike	DividendYield	0.41	0.36
Extrinsic Value	GreekBLS_5	0.138	0.41
Extrinsic Value	FitWellGreekBLS_5	0.14	0.41
LogMoneyness	GreekBLS_2	0.18	-0.41
Strike	FitWellGreekBLS_2	0.079	-0.41
FitWellParameters_2	FitWellParameters_5	0.09	0.41
FitWellIV	FitWellGreekBLS_3	-0.4	-0.222
FitWellIV	FitWellGreek_3	0.04	0.40
FitWellParameters_3	FitWellParameters_5	0.231	-0.40
LogMoneyness	FitWellIV	-0.263	0.40
FitWellParameters_3	FitWellGradiant_5	0.071	0.40
FitWellGreek_3	FitWellGreekBLS_2	-0.37	-0.25
FitWellParameters_2	FitWellGradiant_5	-0.39	-0.31
FitWellParameters_1	FitWellRiskFreeRate	-0.39	-0.33
RiskFreeRate	FitWellParameters_1	-0.39	-0.33
GreekBLS_4	FitWellGreekBLS_5	-0.087	-0.39
GreekBLS_2	FitWellGreek_3	-0.40	-0.25
GreekBLS_4	GreekBLS_5	-0.087	-0.38
DividendYield	FitWellGreekBLS_2	-0.046	-0.38
FitWellDividendYield	FitWellGreekBLS_2	-0.045	-0.38
GreekBLS_2	GreekBLS_4	0.017	0.38
FitWellResidual	FitWellIV	0.05	0.38
FitWellGradiant_4	FitWellGreekBLS_4	0.176	0.38
Strike	FitLayerUsed	0.099	0.38
Intrinsic Value	Extrinsic Value	-0.31	-0.37
FitWellGreekBLS_2	FitWellGreekBLS_4	0.04	0.37
Strike	Layer1Residual	-0.11	-0.37
Strike	GreekBLS_2	0.083	-0.37
LogMoneyness	FitLayerUsed	0.093	0.37
Extrinsic Value	FitWellGreek_5	0.18	0.37
GreekBLS_4	FitWellGreekBLS_2	0.37	0.37
LogMoneyness	Layer1Residual	-0.09	-0.37

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
Ask Size	GreekBLS_3	0.04	0.37
Bid	Ask Size	-0.163	-0.37
FitWellParameters_1	FitWellParameters_5	0.36	0.069
Strike	FitWellGreekBLS_1	-0.36	-0.117
Strike	FitWellGreek_1	-0.36	-0.107
Strike	GreekBLS_1	-0.36	-0.121
Mid	Ask Size	-0.16	-0.36
ImpliedUnderlyingPrice	FitWellParameters_3	0.175	-0.36
Ask	Ask Size	-0.165	-0.36
FitWellGreek_4	FitWellParameters_2	0.06	-0.364
Ask Size	ImpliedOptionPrice	-016	-0.36
Ask Size	Last	-0.17	-0.36
FitWellParameters_4	FitWellGradiant_4	-0.115	-0.36
Ask Size	FitWellOptionPrice	-0.16	-0.36
GreekBLS_4	FitWellGradiant_4	0.178	0.36
FitWellUnderlyingPrice	FitWellParameters_3	0.175	-0.36
Strike	FitWellGradiant_3	0.35	0.064
FitWellGreek_4	FitWellParameters_4	0.06	0.35
FitWellParameters_2	FitWellGradiant_1	0.35	0.264
Ask Size	Intrinsic Value	-0.17	-0.35
Bid Size	Mid	-0.122	-0.35
Bid Size	Ask	-0.123	-0.35
Bid	Bid Size	-0.122	-0.35
Bid Size	ImpliedOptionPrice	-0.12	-0.35
Underlying Price	FitWellParameters_3	0.173	-0.35
Bid Size	Last	-0.12	-0.35
ImpliedUnderlyingPrice	FitWellFeller	-0.34	-0.091
FitWellUnderlyingPrice	FitWellFeller	-0.34	-0.085
Underlying Price	FitWellFeller	-0.34	-0.087
Open Interest	Volume	0.34	0.198
Bid Size	FitWellOptionPrice	-0.12	-0.34
FitWellGreek_4	GreeksBLS_3	0.05	-0.34
Bid Size	Intrinsic Value	-0.08	-0.34
Ask Size	FitWellGreekBLS_3	-0.02	0.34
Layer2Residual	FitWellIV	0.045	0.34
FitLayerUsed	FitWellParameters_3	0.056	0.34
Extrinsic Value	FitWellGradiant_3	0.17	0.34
GreekBLS_2	FitWellGreekBLS_4	0.024	0.34
Strike	FitWellGradiant_1	0.33	0.198
Volume	FitWellGreekBLS_2	0.33	0.055
LogMoneyness	FitWellGradiant_5	-0.33	0.256
Volume	GreekBLS_2	0.33	0.048
FitWellIV	FitWellGreek_2	-0.33	-0.068
FitWellGradiant_3	FitWellGreekBLS_1	-0.33	0.146
FitWellGreek_1	FitWellGradiant_3	-0.33	0.09
GreekBLS_1	FitWellGradiant_3	-0.33	0.032
GreekBLS_4	FitWellGreekBLS_3	0.007	-0.33
Ask Size	FitWellGreek_1	-0.06	0.33
Ask Size	FitWellGreekBLS_1	-0.07	0.33
Bid Size	GreekBLS_3	0.03	0.33
FitWellIV	FitWellGradiant_1	-0.24	-0.33
DividendYield	GreekBLS_2	-0.050	-0.33
GreekBLS_2	FitWellDividendYield	-0.048	-0.33
DividendYield	GreekBLS_4	0.32	-0.087
IV	FitWellGreek_2	-0.32	-0.046
LogMoneyness	GreekBLS_5	-0.32	-0.172
LogMoneyness	FitWellGreek_5	-0.32	-0.163
LogMoneyness	FitWellGreekBLS_5	-0.32	-0.169
GreekBLS_4	FitWellDividendYield	0.32	-0.09
GreekBLS_2	FitWellGreekBLS_3	-0.19	-0.32

Continued on next page

(Table F.1.1 continued)

Variable 1	Variable 2	Stable	Unstable
Extrinsic Value	RiskFreeRate	-0.022	-0.32
Bid Size	FitWellGreekBLS_1	-0.04	0.32
Extrinsic Value	FitWellRiskFreeRate	-0.02	-0.32
Bid Size	FitWellGreek_1	-0.04	0.32
FitWellGreek_5	FitWellGreekBLS_4	-0.05	-0.32
Ask Size	GreekBLS_1	-0.07	0.32
Ask Size	Layer2Residual	0.005	-0.32
DividendYield	FitWellGreekBLS_4	0.31	-0.114
FitWellDividendYield	FitWellGreekBLS_4	0.31	-0.12
Volume	FitWellGreek_3	-0.32	- 0.067
FitWellParameters_5	FitWellGradiant_1	-0.31	-0.22
Intrinsic Value	LogMoneyness	-0.31	0.184
Days To Expiration	DividendYield	0.31	0.246
FitWellGreek_2	FitWellGradiant_1	0.31	0.072
Days To Expiration	RiskFreeRate	0.043	-0.31
Days To Expiration	FitWellRiskFreeRate	0.046	-0.31
FitWellParameters_3	FitWellGradiant_3	-0.176	-0.31
Bid Size	GreekBLS_1	-0.04	0.31
Open Interest	FitWellGreek_3	0.118	0.31
IV	FitWellGreekBLS_2	-0.14	-0.31
Bid Size	Layer2Residual	-0.02	-0.31
FitWellIV	GreekBLS_1	0.123	-0.31
Extrinsic Value	FitWellGreekBLS_3	-0.02	-0.31
FitWellParameters_3	FitWellGradiant_2	-0.042	0.31
Extrinsic Value	Layer2Residual	-0.053	-0.31
Bid Size	FitWellGreekBLS_3	0.02	0.31
Days To Expiration	FitWellDividendYield	0.30	0.243
FitWellGradiant_1	FitWellDividendYield	0.30	0.256
FitWellGradiant_2	FitWellGradiant_3	0.30	0.184
DividendYield	FitWellGradiant_1	0.30	0.248
Open Interest	FitWellGreek_1	0.015	0.30
Bid Size	FitWellGreek_3	0.03	-0.30
Open Interest	FitWellGreekBLS_1	0.012	0.30
GreekBLS_2	GreekBLS_3	-0.112	-0.30
Days To Expiration	FitWellGradiant_1	0.272	0.30
Last	Open Interest	-0.2	-0.30
FitWellGradiant_5	FitWellGreekBLS_2	-0.08	-0.30
Ask	Open Interest	-0.197	-0.30
Open Interest	GreekBLS_1	0.011	0.30
FitWellGreek_4	FitWellGradiant_3	0.35	0.444
FitWellGreek_4	DividendYield	0.30	-0.06
FitWellGreek_4	FitWellDividendYield	0.31	-0.07
FitWellGreek_4	FitWellRiskFreeRate	-0.06	-0.3
FitWellGreek_4	RiskFreeRate	-0.06	-0.3

F.2 Event-Based Context for Residual Spikes

Table F.2.1 lists notable Apple-related events and broader market contexts corresponding to months with significant calibration residual spikes (defined here as residual > 0.5). These annotations help explain the timing and cause of large model errors.

Table F.2.1: Months with large residual spikes (residual > 0.5) and associated events for Apple and the market

Month	*Residual Count*	*Key Apple Events or Market Context*
Jan 2020	226	*Earnings Q1 FY20 (Jan 28)*: Record $91.8B revenue amid early coronavirus concerns Apple Inc. (2020b); *COVID-19*: WHO declares global emergency (Jan 30) World Health Organization (2020); initial market volatility spike Gaetano (2020).
Mar 2020	58	*Market*: COVID crash bottomed (Mar 23); VIX at record highs (World Economic Forum, 2020; CNBC, 2020); *Product Launch*: Apple introduces 4th-generation iPad Pro with LiDAR and Magic Keyboard support (Mar 18) and MacBook Air with Magic Keyboard refresh (Apple Inc., 2020d,f).
Apr 2020	47	*Earnings Q2 FY20 (Apr 30)*: Revenue $58.3B, pandemic impacts evident (Apple Inc., 2020c); *Product*: iPhone SE launched (Apr 24) Apple Inc. (2020e).
May 2020	9	*API Launch (May 20)*: Apple-Google COVID-19 contact tracing released; privacy concerns impacted adoption BBC News (2020); Etherington (2020).
Jun 2020	12	*WWDC 2020 (Jun 22)*: Transition to Apple Silicon announced, significant strategic shift Apple Inc. (2020a); Murray (2020).
Jan 2021	14	*Earnings Q1 FY21 (Jan 27)*: Record quarterly revenue of $111.4 B—first-ever above $100B—driven by strong iPhone 12 demand (Axios, 2021; Slivka, 2021; Apple Inc., 2021a).
Apr 2021	6	*Earnings Q2 FY21 (Apr 28)*: Record $89.6B revenue, driven by 5G iPhone sales and post-COVID consumer optimism Apple Inc. (2021b); Canon and Jolly (2021).
Jul 2021	6	*WWDC 2021 Aftermath*: iOS 15 developer beta rollout begins; market anticipation for iPhone 13 Apple Inc. (2021a); Wacker (2021).
Oct 2021	5	*Product Launch (Oct 18)*: M1 Pro/Max MacBook Pros, AirPods 3 unveiled Apple Inc. (2021c); Gibbs (2021).
Jan 2022	9	*Earnings Q1 FY22 (Jan 27)*: Apple delivers record quarterly revenue of USD 123.9B despite supply constraints; macroeconomic inflation concerns emerge (Apple Inc., 2022a; The Verge, 2022).
Apr 2022	18	*Earnings Q2 FY22 (Apr 28)*: $97.3B revenue, China COVID lockdowns affected supply; Self Service Repair program initiated (Apr 27) Apple Inc. (2022c); Reuters (2022).
Jul 2022	12	*Post-WWDC Effects*: Apple releases redesigned MacBook Air powered by the M2 chip, and investor sentiment improves despite persistent macroeconomic volatility (Apple Inc., 2022b; Nellis, 2022).
Oct 2022	20	*Product Updates & Market Volatility*: Apple unveils 6th-generation iPad Pro with M2 chip alongside redesigned 10th-generation iPad; market volatility remains elevated during global inflation concerns (Rossignol, 2022; Krauskopf and Davies, 2022).
Jan 2023	15	*Market*: Apple stock fell 3.7% after downgrade; market cap dropped below $2 trillion—the first time since 2021 (Balu and Randewich, 2023).
Feb 2023	8	*AI Development & Earnings*: Apple holds its first in-person internal AI summit since the pandemic, reflecting renewed focus on generative AI efforts (Neely, 2023). Days earlier, it reported a Q1 FY23 revenue decline—its first since 2019—and forecasted a further dip, while confirming recovery in iPhone production after prior supply chain disruptions in China (Nellis, 2023a).
May 2023	21	*Product*: Apple Pay Later US launch; critical security updates (iOS 16.5) Apple Inc. (2023a).
Jul 2023	29	*WWDC 2023 (Jun 5–9)*: Apple unveils the Vision Pro mixed-reality headset and previews updates including macOS Sonoma and iOS17 (Apple Inc., 2023c; Wired Staff, 2023).
Aug 2023	24	*Financial & Supply Signals*: Apple's Q3 earnings beat expectations driven by services, but iPhone sales fell; signs of shipment constraints surfaced as demand weakened (Nellis, 2023b).
Oct 2023	16	*Product Launch (Oct 30)*: Apple unveils M3 chip family powering updated MacBook Pro and iMac models during the "Scary Fast" event (The Verge, 2023; Apple Inc., 2023b).
Nov 2023	15	*Supply Constraints*: Tim Cook confirms high demand and supply constraints for iPhone 15 Pro and Pro Max during fiscal Q4 earnings reporting (Reuters, 2023).
Jan 2024	39	*China Demand Weakness*: Apple's iPhone sales dropped sharply in China during the first week of 2024, signalling early signs of broader market headwinds (Mo and Goh, 2024).
Feb 2024	10	*Product Launch*: Apple Vision Pro becomes publicly available; CEO Tim Cook's first public appearance wearing the headset signals full commitment to spatial computing as Apple's next platform (Bilton, 2024).
Apr 2024	39	*Earnings Q2 FY24 (May 2)*: Significant decline in iPad and Mac sales (-17%); large share buyback announced Apple Inc. (2024b).
May 2024	22	*AI & Product Roadmap*: Leaks and discussions about upcoming M4 chip series Apple Inc. (2024a).