Parameter Sensitivity Analysis of the Heston Stochastic Volatility Model

Artemijs Kravcenko (s3442675)

Supervisors:
Hao Wang & Henning Basold

Bachelor Thesis– Data Science and Artificial Intelligence

Leiden Institute of Advanced Computer Science (LIACS)
liacs.leidenuniv.nl

July 27, 2026

Abstract

This thesis studies how uncertainty in the parameters of the Heston stochastic volatility model affects the log-return distribution at option's maturity and related European option pricing outputs. The analysis focuses on five Heston parameters: initial variance, mean-reversion speed, long-run variance, volatility-of-volatility, and correlation between asset returns and variance. A Monte Carlo simulation is used to generate the terminal distribution, with the continuous-time model approximated on a discrete time grid using an Euler scheme. Negative simulated variance values were truncated at zero when evaluating the drift and diffusion terms. Several baseline parameter configurations are used, each representing different volatility regimes. The simulation is first validated against a density obtained from a Fokker–Planck equation associated with the Heston model.

The sensitivity analysis is carried out using two approaches. First, one-parameter perturbation experiments are used to study the changes in distributional summaries, total variation distance, tail probabilities, and European option prices. Second, Sobol indices are used as a form of global sensitivity analysis to measure how variation in multiple parameters over continuous local ranges contribute to variation in selected model outputs. The results show that parameter importance strongly depended on the baseline parameter values and on the output being measured. Given the experimental setup of this thesis, the initial variance was often influential for changes in distribution's summary values and left-tail probability metrics, while the right-tail was more sensitive to the changes in correlation. This could be mostly attributed to the fact that most baselines had negative correlation. For baseline cases with uncorrelated variance and asset price, other parameters became more important. The Sobol indices analysis showed that long-run variance level explained more variation in the terminal distribution for calm and fast mean-reverting baselines, while initial variance and volatility-of-volatility was of higher significance in high-volatility and persistent-volatility baselines. Overall, most outputs changed approximately linearly over the tested ranges, while short-maturity out-of-the-money option prices and tail threshold ratios were affected disproportionately by parameter perturbations.

Contents

1 Introduction

Derivatives are a type of financial instrument that obtain their value based on the valuation of one or multiple other securities. Options represent a major class of financial derivatives. An option gives its holder the right, but not the obligation, to buy or sell a specified quantity of an underlying asset at a predetermined price. Depending on the type of an option contract, it may be exercised either only on the expiration date or at any time up to that date. These contracts constitute a major part of global derivatives markets, with tens of billions of exchange-traded options contracts traded annually. In 2025 alone, approximately 88.65 billion options contracts were traded worldwide [13].

Three terms are relevant when describing an option, specifically its strike price, maturity, and moneyness. The strike price, denoted by K, is the predetermined price at which the underlying asset may be bought or sold. The maturity date, T, is the date on which the option expires. The period remaining until this date is referred to as the option's time to maturity. European options can be exercised only at maturity, whereas American options can be exercised at any time up to and including maturity. This thesis is focused on the European options.

The price of the underlying asset at time (t) is denoted by S_t, while S_T denotes its price at the option's maturity. The relationship between the underlying price and the strike price is referred to as the option's moneyness. A call option is in the money (ITM) when $S_t > K$, at the money (ATM) when S_t is equal to strike price K, and out of the money (OTM) when $S_t < K$. For a put option, these relationships are reversed. The option is in the money when $S_t < K$ and out of the money when $S_t > K$. Moneyness can also be expressed numerically using a ratio such as S_t/K

At maturity, the payoff of a European call option is

$$\max\{S_T - K, 0\},$$

whereas the payoff of a European put option is

$$\max\{K - S_T, 0\},$$

To provide an example, suppose the strike price is $K = 100$. If the asset price at maturity is $S_T = 120$, the call option is in the money and has a payoff of 20, whereas the put option is out of the money and has a payoff of zero. Conversely, if $S_T = 80$, the call option expires out of the money with a payoff of zero, while the put option is in the money and has a payoff of (20). If $S_T = 100$, both options are at the money and have a payoff of zero. In this thesis, the distribution of terminal log-returns is considered, which is defined as $X_T = \log(S_T/S_0)$, with S_0 being the initial asset price.

Option's price depends on the possible future price of an underlying asset, which cannot be observed at the time of valuation and must therefore be represented using a mathematical model. A number of different models exist, each having its own set of assumptions about the dynamics of the price of an underlying asset. Consequently, calculated option prices depend in part on the parameter values of a selected pricing model.

In general, options are used for protection against future uncertainty, transfer of risk, hedging of positions, and taking on additional exposure to future price movements or volatility. The most well known option pricing framework is the Black–Scholes–Merton (BSM) model developed by Fischer Black, Myron Scholes, and Robert C. Merton in 1973 [5]. One of the main assumptions of this model is that volatility is constant across maturity and strike, which is not empirically true [17].

Stochastic volatility models relax this assumption. Their core distinction is the assumption that the volatility is not constant across strike prices and maturities, but is itself a stochastic process. This thesis is concerned with the Heston model of stochastic volatility, which uses a set of five parameters controlling how the underlying asset's variance, and by extension price, evolves through time [16]. These parameters will be introduced in the background section.

The Heston model is subject to model risk, meaning that its outputs may be incorrect due to errors in the choice of model parameters stemming, for instance, from wrong calibration procedures. Even when the Heston dynamics are assumed to provide an appropriate representation of the market, its parameters are not directly observable and must be estimated from historical data or calibrated to observed option prices. Consequently, the resulting parameter values can be misspecified (i.e. estimated imperfectly), and these errors propagate to the price distribution of the underlying asset, and, by extension, to derivative prices produced by the model. Previous studies have shown that different calibrations of the Heston model can produce substantial differences in derivative prices and hedging outcomes, while the effect of parameter misspecification may be nonlinear and dependent on the parameter configuration [15, 2, 22, 6]. Identifying which parameters contribute most strongly to variation in model outputs, and whether their effects change across maturities and baseline configurations, is therefore relevant for assessing the robustness of valuations produced by the Heston model.

This thesis aims to investigate the impact that the misspecification of the parameters can have on the Heston model. The main research question of this thesis is:

How do changes in the Heston model parameters affect the terminal log-return distribution and European option prices, and how do these effects vary with maturity and the baseline parameter configuration?

This question is divided into the following subquestions:

- How does changing each Heston model parameter individually affect selected properties of the terminal log-return distribution, namely the mean, standard deviation, skewness, and kurtosis?

- How does changing each Heston model parameter individually affect the tails of the terminal log-return distribution, as measured by (i) the probabilities of returns falling more than one, two, or three standard deviations below or above the mean and (ii) selected lower- and upper-tail quantiles, such as the 1st and 99th percentiles?

- Within the selected parameter ranges, what proportion of the variance in defined model outputs is attributable to each Heston model parameter individually and to its interactions with the other parameters, as measured by first-order and total-order Sobol indices?

- How do these effects vary across half-year, one year, and two year option maturities and baseline parameter configurations representing different market conditions?

- Are the relationships between a change in Heston model parameters and the resulting effect on the features of a log-return distribution of asset prices linear or non-linear?

- To what extent can the uncertainty in the Heston model parameters translate into mispricing of European call and put options?

1.1 Thesis overview

The remainder of this thesis is organised as follows. Section 2 introduces the theoretical background on European options, stochastic volatility, the Heston model, and parameter uncertainty. Section 3 reviews previous research on Heston model calibration, parameter uncertainty, and sensitivity analysis. Section 4 describes the setup of the parameter perturbation experiments and the Sobol indices study. Section 5 presents the resulting analysis of the terminal log-return distribution and European option prices. Section 6 discusses the limitations of the methodology. Finally, Section 7 summarises the main findings, discusses the limitations of the study, and outlines possible directions for future research.

2 Background

2.1 European Options and Their Payoffs

Many different types of options exist, with majority of options being either European or American. The distinction between the two is that European options can only be exercised at their expiration date, while American options can be exercised at any moment before the expiration. This thesis focuses on European call and put options as American options generally require additional numerical procedures to account for a possibility that an option may be exercised before the maturity date.

Option prices are evaluated under the risk-neutral probability measure $\mathbb{Q}$. It is a pricing measure under which the expected return of the underlying asset equals the risk-free rate after accounting for dividends. This allows for European option values to be expressed as the discounted risk-neutral expectation of their terminal payoffs. This measure should does not represent the actual probabilities of future market outcomes.

For European options, the price is expressed as the discounted future payoff under a risk-neutral probability measure.

For a European call, this gives

$$C = e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[\max\{S_T - K, 0\}\right].$$

where S_T is the asset price at maturity, r is the risk-free interest rate, and $\mathbb{Q}$ denotes the risk-neutral measure. Similarly, the price of a European put option is given by

$$P = e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[\max\{K - S_T, 0\}\right].$$

From these expressions, it is possible to see that option prices depend on the model-implied distribution of the terminal asset price S_T. Because of the payoff profile of options, their expected value depends on the shape of the terminal price distribution. As a result, changing the shape of the distribution, including its variance, skewness, and significance of the tails, can have a considerable effect on the resulting option value.

2.2 Option Pricing Models

The payoff of an option depends on the future value of the underlying, and its future value is unknown at the time when an option is priced. Therefore, a model is required to describe the possible time-evolution of the asset's price.

The classical BSM framework provides one of the most well-known approaches to option pricing [5]. This model makes several assumptions about how an underlying behaves, with the most important one being that the asset follows a geometric Brownian motion with constant volatility. The BSM model has a parabolic partial differential equation associated with it, which provides how the price of the asset evolves through time. The model is analytically tractable and is widely used as a benchmark.

The constant volatility assumption, however, is inconsistent with patterns observed in option markets. Implied volatility often varies across strike prices and maturities, meaning that a single constant volatility parameter is generally not sufficient to describe the full pattern of market option prices [17]. Implied volatility is the volatility value that, when used as a volatility parameter in the BSM option pricing model, makes the model price equal to the observed market price of the option. An example of an implied volatility surface is shown in Figure 1.

For a fixed maturity, the variation in implied volatility across strikes is commonly described as either a volatility smile or a volatility skew. A smile is a roughly U-shaped pattern in which options with strike values far from the at-the-money strike value have higher implied volatilities than at-the-money options. This can be interpreted that options exposed to extreme price movements are priced higher than would be predicted by a BSM model calibrated to at-the-money volatility value. A skew, on the other hand, is asymmetric and appears as a slope. It should also be added that under the Black–Scholes–Merton model, the option price is strictly increasing in volatility. Consequently, there is a unique implied volatility value corresponding to that price. For example, $\sigma_{\text{imp}}(100, 1) = 0.22$ means that an option with strike $K = 100$ and one year to maturity is priced with a BSM volatility value of 22%.

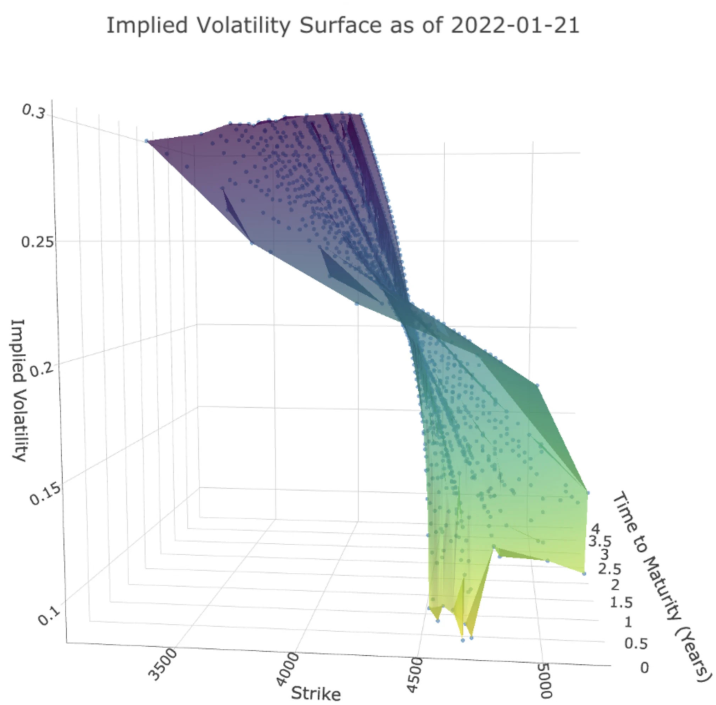

Figure 1: An empirical implied volatility surface for put options from S&P500 index on January 2022, illustrating a volatility skew. Reproduced from Luo and Yu (2023) [19], Figure 1.

2.3 Stochastic Volatility and the Heston Model

The Heston model is a stochastic volatility model introduced by Steven L. Heston in 1993 [16]. The model assumes that the asset price is driven by a variance process that evolves randomly over time. It is characterised by five parameters. The initial variance v_0 determines the starting level of variance. The parameter κ controls the speed at which the variance resets toward its long-run level. The parameter θ represents the long-run variance level, while ξ controls the volatility of the variance process process itself, and is often called volatility-of-volatility. Lastly, ρ determines the correlation between the movement of the asset price and the volatility. Depending on the chosen parameter values, different market regimes can be modelled, for instance low-volatility, high-volatility, or persistent-volatility.

The Heston model assumes that the underlying asset price S_t evolves as a stochastic process, given as

$$dS_t = (r - q)S_t \, dt + \sqrt{v_t} S_t \, dW_t^S,$$

where the volatility term $\sqrt{\nu}$ is given by a Cox-Ingersoll-Ross (CIR) process,

$$dv_t = \kappa(\theta - v_t) \, dt + \xi \sqrt{v_t} \, dW_t^v.$$

The parameters of the Heston model are:

- initial variance, v_0;
- long-run variance, θ;
- speed of mean reversion, κ;
- volatility of variance, ξ; and
- correlation between the Brownian motions driving the asset price and its variance, ρ.

Terms dW_t^S and dW_t^v are continuous random walks, or Wiener processes, with a correlation ρ:

$$dW_t^S dW_t^v = \rho \, dt.$$

A CIR process is a mean-reverting square-root diffusion. More specifically, its expected change is downward when the current variance v_t is above θ and upward when it is below θ. This follows from the drift term $\kappa(\theta - v_t)$, which is negative when $v_t > \theta$ and positive when $v_t < \theta$. Random shocks can still move variance in either direction, but the drift prevents such changes from persisting indefinitely in the absence of further shocks. This allows the Heston model to better represent changes in variance over time, which can affect the dispersion, asymmetry, and tail behaviour of the terminal asset-price distribution.

The Feller condition is

$$2\kappa\theta \geq \xi^2.$$

Its satisfaction ensures strict positivity of the continuous-time CIR variance process under the standard sufficient condition. Configurations that violate the condition are retained because empirically relevant Heston calibrations do not necessarily satisfy it. In the simulation, the resulting boundary problem is handled using the full-truncation Euler scheme described in Section 4.4.

2.4 Parameter Uncertainty and Model Risk

In the context of options, model risk refers to the risk of the model producing misleading or inaccurate outputs stemming from incorrect model assumptions, implementation errors, uncertainty in inputs, or model usage [8, 10]. Model risk carries tangible financial costs, which can be illustrated by the near-collapse of the Long-Term Capital Management (LTCM) hedge fund in 1998. The fund's risk management approach underestimated its exposure to extreme and correlated market movements, while high leverage and concentrated, relatively illiquid positions have exacerbated the losses. As a result, LTCM was rescued through a \$3.6 billion private recapitalization provided by a group of fourteen banks and brokerage firms and coordinated by the Federal Reserve Bank of New York [20].

In addition to the LTCM case, several other examples illustrate the financial impact of model risk. Cont (2006) reports that the Bank of Tokyo/Mitsubishi suffered an \$83 million loss in 1997 after an internal pricing model overvalued a portfolio of swaps and options on U.S. interest rates. In the same period, NatWest Capital Markets announced a £50 million loss linked to the mispricing of German and U.K. interest-rate options and swaptions (options on interest-rate swaps). Cont notes that these losses were attributed by contemporary observers to model risk [7].

One of the potential sources of model error with the Heston model can come from the parameter misspecification. The parameters of the Heston model must be specified or inferred from available information. They may be estimated from historical data, calibrated to observed option prices, taken from previous empirical studies, or selected to represent a particular market regime. The resulting parameter values can depend on factors such as the estimation method, calibration objective, data sample, and calibration period.

Consequently, even when the Heston model itself is accepted as the appropriate modelling framework, there is inherent uncertainty regarding the values assigned to its parameters. This parameter uncertainty constitutes one component of model risk, since changes in the parameter configuration may alter the model-implied distribution of the underlying asset and the resulting option prices.

3 Related Work

Overall, the existing literature shows that Heston model parameters must be estimated or calibrated, and that model and parameter uncertainty can materially affect derivative valuation. The present thesis contributes to this line of work by studying how estimation error of the Heston model parameters affects both the return distribution of asset prices and European option prices across several baseline parameter configurations.

After the Heston model had initially been introduced, further work has been done to develop it further. Bates (1996), for instance, has added log-normal price jumps to the model, allowing for the possibility of sudden and discontinuous movements in the underlying asset price which are driven by the Poisson process [3]. In an empirical application to Deutsche Mark options, Bates found that the default stochastic volatility model could reproduce the observed volatility smile only with parameter values that were implausible relative to the time series behaviour of implied volatility. In addition, incorporating jumps provided an additional explanation for the excess kurtosis reflected in option prices.

A related branch of literature looked at increasing the flexibility of the Heston model by allowing its parameters to vary over time. Elices (2007) presented a methodology to introduce time dependence for parameters used in a variety of models, including stochastic volatility models [12]. Benhamou, Gobet & Miri (2010) developed this direction further by deriving an analytical approximation for vanilla option prices under a time-dependent Heston model specification [4].

A related line of research focused on the estimation and calibration of Heston model parameters. The paper by Aït-Sahalia & Kimmel (2007), for instance, has looked at developing a closed-form maximum likelihood estimator for stochastic volatility (SV) models, including Heston and Bates [1]. The paper provided an efficient method to estimate SV model parameters from a discrete time series of asset prices. Gilli & Schumann (2011) have discussed the use of optimisation heuristics for financial optimisation problems which have multiple local optima, discontinuities, constraints, or otherwise poorly behaved objective functions [14]. Cui, del Baño Rollin & Germano (2017) take a different approach by expressing Heston calibration as a nonlinear least-squares problem and deriving an analytical gradient of vanilla option prices with respect to the model parameters [9]. They combine this gradient with a Levenberg–Marquardt optimisation method, which makes calibration substantially faster than approaches based on numerical gradients. Their results also show that the Heston calibration objective may have a narrow valley with a flat bottom, meaning that different parameter combinations can produce similar calibration errors. More recently, a study done by Zadgar, Fallah & Mehrdoust (2025) presented a calibration approach for the Heston model using neural networks. They find that neural-network-based method outperforms traditional calibration methods. The authors claim that the deep learning approach presented in the paper converges more quickly and generalises better that the standard optimisation [23].

The field of model risk is relatively new. Derman's 1996 report gives a practitioner's discussion of general sources of model risk, including incorrect model choice, erroneous solutions, numerical issues, and implementation problems [10]. In academic literature, Rama Cont's paper (2006) has introduced a quantitative framework for measuring model uncertainty, specifically in the context of derivatives pricing. The paper proposed two approaches, one based on a coherent risk measure over a set of pricing models calibrated to benchmark derivatives, and another based on a convex risk measure that relaxes exact calibration by penalising models according to their pricing errors [7].

A study related directly to the Heston model is Suárez-Taboada et al. (2018), who apply uncertainty quantification to the Heston model. Instead of comparing different model classes, they keep the Heston model fixed and treat selected parameters as uncertain inputs. Using stochastic collocation, they study how uncertainty in these parameters propagates to option prices, with particular attention to parameters involved in the Feller condition [22].

4 Methodology

4.1 Objective of Numerical Experiments

The objective of the Monte Carlo experiment is to study how selected outputs of the model respond to changes in its parameters. The methodology is not designed to estimate the true market uncertainty and does not attempt to fit the Heston parameters for a certain market regime. The goal of the Monte Carlo experiment is to examine how small perturbations in the Heston model parameters propagate to model outputs, given an baseline parameter configuration representing a chosen market condition.

The first part of the experiment uses one-parameter perturbations. For each baseline configuration, one Heston parameter is varied while the remaining parameters are held fixed. This makes it possible to examine the isolated response of each output to changes in that particular parameter. The second part of the experiment uses Sobol sensitivity analysis as a supplementary method. Instead of changing one parameter at a time, all parameters are varied simultaneously within specified ranges around a baseline configuration. The resulting Sobol indices are interpreted as range-dependent sensitivity measures. They describe how much of the variation in each selected model output is associated with individual parameters and their interactions within the chosen parameter ranges.

4.2 Model Setup

4.2.1 Heston Model Dynamics

We define the parameter vector Θ as

$$\Theta = (v_0, \kappa, \theta, \xi, \rho)$$

4.2.2 Terminal Log-Return

The terminal log-return of the model is defined as

$$X_T = \log(S_T/S_0),$$

where S_T is the asset price at time T and S_0 is the initial asset price. We are then able to define the change in the random variable X_T as

$$dX_t = \left(r - q - \frac{1}{2}v_t\right) dt + \sqrt{v_t}\, dW_t^S.$$

The derivation for this expression is provided in Appendix A.1.

4.3 Baseline Parameter Configuration

The experiment is conducted around six baseline vectors representing different volatility regimes, as results obtained for a single configuration may be specific to that particular configuration. Baselines A-E are intended to represent distinct volatility regimes, including low-, moderate-, high-, and persistent-volatility conditions, as well as a zero-correlation setting, whereas baseline F adopted from literature. The values of κ, θ, ξ, and ρ are adopted from Drăgulescu and Yakovenko (2002) with the initial variance v_0 set equal to the long-run variance θ [11].

The baseline parameter vectors are reported in Table 1. The parameters v_0 and θ are variances, meaning that their square roots can be interpreted as annualised volatility levels. The mean-reversion half-life is

$$\tau_{1/2} = \frac{\log 2}{\kappa},$$

which provides an additional, interpretable measure of the persistence of deviations from θ. It measures the time required for the deviation of the variance process from its long-run level to decrease by one half.

Table 1: Baseline Heston parameter configurations used in the sensitivity experiments.

Case	v_0	κ	θ	ξ	ρ	Feller	Interpretation
A	0.040	2.00	0.040	0.300	-0.70	Yes	Moderate regime
B	0.020	3.00	0.020	0.200	-0.30	Yes	Low volatility case
C	0.090	1.00	0.060	0.700	-0.80	No	High volatility case
D	0.040	0.50	0.040	0.250	-0.70	No	Persistent volatility case
E	0.040	2.00	0.040	0.300	0.00	Yes	Zero correlation case
F	0.022	11.35	0.022	0.618	0.00	Yes	Configuration adopted from literature

Baseline A has $v_0 = \theta = 0.04$, corresponding to an initial and long-run volatility of 20%. Its negative correlation parameter represents a high correlation between the variance and asset price movements, while $\kappa = 2$ implies a variance half-life of approximately 0.35 years. It is used as the central moderate-volatility configuration.

Baseline B represents a calmer regime. Its initial and long-run volatility are approximately 14.1%, while its variance process reverts more quickly and has lower volatility of variance than baseline A. Its weaker negative correlation also implies less pronounced return-distribution asymmetry.

Baseline C represents a high-volatility configuration. Its initial volatility is 30%, its long-run volatility is approximately 24.5%, and it combines high volatility of variance with strong negative correlation. The Feller condition is violated in this case, making it useful for examining sensitivity in a regime where the variance process can approach zero and the numerical treatment of the variance boundary becomes more important.

Baseline D represents persistent variance. It retains the 20% initial and long-run volatility levels but uses $\kappa = 0.5$, corresponding to a mean-reversion half-life of approximately 1.39 years. Variance shocks and the initial variance level are therefore expected to remain influential over a larger part of the investigated maturity range.

Baseline E is identical to baseline A except that $\rho = 0$. The comparison between A and E therefore isolates the effect of return–variance correlation while keeping all other model parameters fixed.

The values $\kappa = 11.35$, $\theta = 0.022$, $\xi = 0.618$, and $\rho = 0$ used in baseline F are based on the annualised estimates reported by Drăgulescu and Yakovenko [11]. Their estimates were obtained by fitting the Heston return distribution to Dow–Jones Industrial Average data. Since their specification does not provide a fixed initial variance for the present conditional simulation, v_0 is set equal to the fitted long-run variance θ. Baseline F should therefore be interpreted as literature-informed rather than as a direct reproduction of their empirical experiment or as an option-market calibration.

4.4 Monte Carlo Simulation Procedure

The Heston model is simulated on a discrete time grid. Let the maturity be T, and let N denote the number of time steps. The time step size is

$$\Delta t = \frac{T}{N},$$

with grid points

$$t_n = n\Delta t, \qquad n = 0, \dots, N.$$

Since the variance process appears inside square roots, the simulation uses the truncated variance

$$v_n^+ = \max(v_n, 0).$$

This prevents negative simulated variance values from being used inside the diffusion terms. Other truncation schemes exist, such as reflection, where the negative variance value is replaced with its absolute value. Full truncation was selected as it has been found to produce lower bias and RMS error in comparison with other truncation schemes in simulations of the Heston model [18].

At each time step, two independent standard normal random variables are generated:

$$Z_n^S, Z_n^\perp \sim \mathcal{N}(0, 1).$$

The Heston model requires the Brownian shocks driving the asset price and the variance process to have correlation ρ. To impose this correlation, the variance shock is constructed as

$$Z_n^v = \rho Z_n^S + \sqrt{1 - \rho^2} Z_n^\perp.$$

Since Z_n^S and $Z_n^\perp$ are independent standard normal variables,

$$\mathrm{Var}(Z_n^v) = \rho^2 + (1 - \rho^2) = 1, \qquad \mathrm{Cov}(Z_n^S, Z_n^v) = \rho.$$

Consequently, $Z_n^v \sim \mathcal{N}(0, 1)$ and

$$\mathrm{Corr}(Z_n^S, Z_n^v) = \rho.$$

This construction is valid for all $\rho \in [-1, 1]$, including negative values of ρ.

The variance process is then updated using the Euler scheme with full truncation of the variance value:

$$v_{n+1} = v_n + \kappa(\theta - v_n^+)\Delta t + \xi \sqrt{v_n^+}\sqrt{\Delta t}\, Z_n^v.$$

The terminal log-return process is updated as

$$X_{n+1} = X_n + \left(r - q - \frac{1}{2}v_n^+\right)\Delta t + \sqrt{v_n^+}\sqrt{\Delta t}\, Z_n^S.$$

The initial values are

$$X_0 = 0, \qquad v_0 > 0,$$

where v_0 is given by the chosen initial variance value and $X_0 = 0$ follows from

$$X_0 = \log(S_0/S_0) = 0.$$

After repeating this procedure for M simulated paths, the result is a sample of terminal log-returns

$$X_T^{(1)}, X_T^{(2)}, \dots, X_T^{(M)}.$$

This simulated sample is then used to estimate distributional summaries and tail probabilities of X_T reported in sections 5.4 and 5.5, respectively.

4.5 Parameter Perturbation Design

For each experiment, a perturbed parameter vector is constructed and the Monte Carlo simulation is repeated under this new parameter set. The resulting terminal log-return sample is then compared with the baseline sample using the metrics defined in the following section.

4.5.1 One-Parameter Perturbations

The main perturbation design changes one Heston parameter at a time while holding all other parameters fixed at their baseline values. This allows the effect of each individual parameter to be isolated.

A one-parameter perturbation of parameter θ_j is defined as

$$\Theta_\varepsilon^{(j)} = (\theta_1, \dots, \theta_j(1+\varepsilon), \dots, \theta_d),$$

where ε is the relative perturbation size. For example, a positive perturbation of the volatility of volatility parameter ξ gives

$$\Theta_\varepsilon^{(\xi)} = (v_0, \kappa, \theta, \xi(1+\varepsilon), \rho).$$

For each perturbed parameter vector, the same Monte Carlo procedure is applied to generate a new sample of terminal log-returns. This sample is then compared against the baseline distribution generated under the corresponding set of parameters.

4.5.2 Perturbation Sizes

The perturbation sizes are chosen symmetrically around the baseline value. The relative perturbation grid ranges from -0.20 to 0.20, with increments of 0.02:

$$\varepsilon \in \{-0.20, -0.18, \dots, 0, \dots, 0.18, 0.20\}.$$

Thus, each parameter is varied from 80% to 120% of its baseline value in increments of 2%, while all remaining parameters are held fixed. This perturbation range was chosen to assess the sensitivity of the distributional properties of log-returns and option prices to relatively small changes in the model parameters. The aim is to examine whether modest parameter variations can generate disproportionately large or nonlinear effects on these variables.

For strictly positive parameters such as v_0, κ, θ, and ξ, multiplicative perturbations preserve positivity as long as $1+\varepsilon > 0$. For the correlation parameter ρ, additive perturbations are used instead:

$$\rho_\varepsilon = \rho + \varepsilon,$$

with ε perturbation sizes kept the same and with values satisfying $\rho_\varepsilon \in [-1, 1]$ only being considered.

4.6 Metrics for Distributional Analysis

4.6.1 Distribution Summaries

All distributional metrics introduced below depend on the maturity T and the Heston parameter vector. For readability, these arguments are omitted unless different parameter configurations are compared explicitly. For each parameter configuration Θ, the simulated terminal log-return sample is denoted by

$$X_T^{(1)}, X_T^{(2)}, \ldots, X_T^{(M)}.$$

The distribution of X_T is first summarised using its mean, standard deviation, skewness, and excess kurtosis. These quantities describe the location, dispersion, asymmetry, and tail-heaviness of the terminal log-return distribution.

The mean of X_T under parameter set Θ is defined as

$$\mu_T = \mathbb{E}_{\mathbb{Q}}[X_T].$$

In the Monte Carlo experiment, this is estimated by the sample mean

$$\hat{\mu}_T = \frac{1}{M} \sum_{m=1}^{M} X_T^{(m)}.$$

The standard deviation is defined as

$$\sigma_T = \sqrt{\mathrm{Var}_{\mathbb{Q}}(X_T)},$$

which is estimated by

$$\hat{\sigma}_T = \sqrt{\frac{1}{M-1} \sum_{m=1}^{M} \left(X_T^{(m)} - \hat{\mu}_T \right)^2}.$$

Skewness is used to measure the asymmetry of the terminal log-return distribution:

$$\mathrm{Skew}(X_T) = \frac{\mathbb{E}_{\mathbb{Q}}\left[(X_T - \mu_T)^3\right]}{(\sigma_T)^3}.$$

The Monte Carlo estimate is

$$\widehat{\mathrm{Skew}}(X_T) = \frac{\frac{1}{M} \sum_{m=1}^{M} \left(X_T^{(m)} - \hat{\mu}_T \right)^3}{(\hat{\sigma}_T)^3}.$$

Kurtosis is used to measure the heaviness of the tails relative to the spread of the distribution:

$$\mathrm{Kurt}(X_T) = \frac{\mathbb{E}_{\mathbb{Q}}\left[(X_T - \mu_T)^4\right]}{(\sigma_T)^4}.$$

Since a Gaussian distribution has kurtosis equal to 3, excess kurtosis is reported as

$$\mathrm{ExKurt}(X_T) = \mathrm{Kurt}(X_T) - 3.$$

The Monte Carlo estimate is

$$\widehat{\mathrm{ExKurt}}(X_T) = \frac{\frac{1}{M}\sum_{m=1}^{M}\left(X_T^{(m)} - \hat{\mu}_T\right)^4}{(\hat{\sigma}_T)^4} - 3.$$

4.6.2 Metrics for Evaluating the Impact of Parameter Perturbation on Tail Probabilities

Tail probabilities are used to measure how parameter perturbations affect the probability of extreme underlying asset prices at option's maturity. Throughout this section, a superscript 0 denotes the baseline parameter configuration, while a superscript (j, ε) denotes the configuration obtained by perturbing parameter j by ε.

For a fixed maturity T, let

$$F_T(x) = \mathbb{P}_{\mathbb{Q}}(X_T \leq x)$$

denote the cumulative distribution function of the terminal log-return under a given parameter configuration. The baseline distribution is denoted by F_T^0.

For a tail level $\alpha \in (0, 1)$, the left and right tail thresholds are defined from the baseline distribution as

$$c_L(T, \alpha) = \left(F_T^0\right)^{-1}(\alpha),$$

and

$$c_R(T, \alpha) = \left(F_T^0\right)^{-1}(1 - \alpha).$$

These thresholds are fixed using the baseline distribution. The same thresholds are then used when evaluating the perturbed distributions. This ensures that the baseline and perturbed models are compared on the probability of the same tail event.

For a perturbation of parameter j by ε, the left tail probabilities are

$$p_L^0(T, \alpha) = \mathbb{P}_{\mathbb{Q}}^0\left(X_T \leq c_L(T, \alpha)\right),$$

and

$$p_L^{j,\varepsilon}(T, \alpha) = \mathbb{P}_{\mathbb{Q}}^{j,\varepsilon}\left(X_T \leq c_L(T, \alpha)\right).$$

Similarly, the right-tail probabilities are

$$p_R^0(T, \alpha) = \mathbb{P}_{\mathbb{Q}}^0\left(X_T \geq c_R(T, \alpha)\right),$$

and

$$p_R^{j,\varepsilon}(T,\alpha) = \mathbb{P}_{\mathbb{Q}}^{j,\varepsilon}\left(X_T \geq c_R(T,\alpha)\right).$$

In addition to quantile-based thresholds, a fixed standard-deviation threshold is used to compare the baseline distribution with a distribution resulting from perturbing parameters. For a chosen constant $k > 0$, we define the baseline right-tail threshold as

$$c_{R,k}(T) = \mu_T^0 + k\sigma_T^0.$$

The corresponding right-tail probability ratio is

$$R_{R,k}^{j,\varepsilon}(T) = \frac{\mathbb{P}_{\mathbb{Q}}^{j,\varepsilon}\left(X_T \geq c_{R,k}(T)\right)}{\mathbb{P}_{\mathbb{Q}}^{0}\left(X_T \geq c_{R,k}(T)\right)}.$$

This metric measures how much more or less likely a baseline-defined k-standard-deviation right-tail event becomes after perturbing one model parameter.

4.6.3 Sobol Indices

Sobol indices are variance-based global sensitivity measures that decompose the variance of a scalar model output into contributions from individual input parameters and their interactions [21].

For each baseline configuration of the model, Sobol indices are computed over a local parameter region around the baseline. Positive parameters are varied multiplicatively around their baseline values, while the correlation parameter is varied additively subject to the constraint $\rho \in [-1, 1]$. This means that the Sobol indices should be interpreted as sensitivity measures conditional on selected baselines.

Table 2 lists the scalar outputs Y used in the Sobol sensitivity analysis. Each output is computed from the Monte Carlo terminal log-return sample generated under a sampled Heston parameter vector. The Sobol indices are then computed separately for each output, so the resulting parameter importance is output-dependent.

The first-order Sobol index of parameter Θ_i is defined as

$$S_i = \frac{\mathrm{Var}\left(\mathbb{E}[Y \mid \Theta_i]\right)}{\mathrm{Var}(Y)}.$$

This measures the fraction of output variance explained by a single parameter. A large first-order index means that changing this parameter by itself explains a large share of the variation in the output.

The total-order Sobol index is defined as

$$S_{T_i} = 1 - \frac{\mathrm{Var}\left(\mathbb{E}[Y \mid \Theta_{-i}]\right)}{\mathrm{Var}(Y)},$$

For each sampled parameter vector, the same Monte Carlo procedure is used to generate terminal log-return samples. The selected scalar outputs are then computed and passed to estimate Sobol indices. The resulting indices are reported separately for each output metric and each baseline configuration. In the interpretation of the results, total-order indices are used as the main ranking measure, since they include interaction effects.

Output metric	Notation	Description
Mean log-return at expiration	$\mathbb{E}[X_T]$	Average terminal log-return under the sampled Heston parameter vector.
Standard deviation	$\mathrm{Std}(X_T)$	Dispersion of the terminal log-return distribution.
Skewness	$\mathrm{Skew}(X_T)$	Measures asymmetry of the terminal log-return distribution. Negative skewness indicates a stronger left tail.
Excess kurtosis	$\mathrm{ExKurt}(X_T)$	Measures tail-heaviness relative to a Gaussian distribution. A Gaussian distribution has excess kurtosis equal to zero.
Left k-sigma tail probability	$\mathbb{P}(X_T \leq \mu_T^0 - k\sigma_T^0)$	Probability of falling below a baseline-defined left-tail threshold based on the baseline mean and standard deviation.
Right k-sigma tail probability	$\mathbb{P}(X_T \geq \mu_T^0 + k\sigma_T^0)$	Probability of exceeding a baseline-defined right-tail threshold based on the baseline mean and standard deviation.
Left quantile tail probability	$\mathbb{P}(X_T \leq c_L(T, \alpha))$	Probability of falling below the baseline α-quantile threshold. In the experiment, this is used to study downside tail sensitivity.
Right quantile tail probability	$\mathbb{P}(X_T \geq c_R(T, \alpha))$	Probability of exceeding the baseline $(1-\alpha)$-quantile threshold. In the experiment, this is used to study upside tail sensitivity.
Total variation distance	$d_{\mathrm{TV}}(F_T^{\Theta}, F_T^{\Theta_0})$	Distance between the perturbed and baseline terminal log-return distributions, estimated from histogram bin probabilities.
Put option price	$P(0.9S_0, T)$	European put price with strike $K = 0.9S_0$, estimated from simulated terminal asset prices.
Call option price	$C(1.1S_0, T)$	European call price with strike $K = 1.1S_0$, estimated from simulated terminal asset prices.
Relative put price change	$\frac{P_{\Theta}(0.9S_0,T) - P_{\Theta_0}(0.9S_0,T)}{P_{\Theta_0}(0.9S_0,T)}$	Relative change in the $K = 0.9S_0$ put price compared with the baseline price.
Relative call price change	$\frac{C_{\Theta}(1.1S_0,T) - C_{\Theta_0}(1.1S_0,T)}{C_{\Theta_0}(1.1S_0,T)}$	Relative change in the $K = 1.1S_0$ call price compared with the baseline price.

Table 2: Scalar output metrics used in the Sobol sensitivity analysis. Each output is computed for sampled Heston parameter vectors and used as the scalar quantity $Y = G(\Theta)$ in the Sobol variance decomposition.

4.6.4 Impact on European Call/Put Prices

In addition to distributional metrics of the terminal log-return, the experiment also studies how parameter perturbations affect European option prices.

For each simulation run, the terminal log-return samples are converted into terminal asset prices by

$$S_T^{(m)} = S_0 \exp(X_T^{(m)}), \qquad m = 1, \dots, M.$$

For strike K and maturity T, the European call payoff is

$$(S_T - K)^+,$$

and the European put payoff is

$$(K - S_T)^+.$$

The baseline Monte Carlo call and put prices are estimated as

$$\widehat{C}_0(K,T) = e^{-rT} \frac{1}{M} \sum_{m=1}^{M} \left(S_{T,0}^{(m)} - K \right)^+,$$

and

$$\widehat{P}_0(K,T) = e^{-rT} \frac{1}{M} \sum_{m=1}^{M} \left(K - S_{T,0}^{(m)} \right)^+.$$

Here $S_{T,0}^{(m)}$ denotes the terminal asset price from the baseline simulation.

For a perturbation of parameter j with perturbation size ε, the corresponding perturbed terminal asset prices are denoted by

$$S_{T,j,\varepsilon}^{(m)}.$$

The call and put prices obtained after a parameter has been changed are estimated as

$$\widehat{C}_{j,\varepsilon}(K,T) = e^{-rT} \frac{1}{M} \sum_{m=1}^{M} \left(S_{T,j,\varepsilon}^{(m)} - K \right)^+,$$

and

$$\widehat{P}_{j,\varepsilon}(K,T) = e^{-rT} \frac{1}{M} \sum_{m=1}^{M} \left(K - S_{T,j,\varepsilon}^{(m)} \right)^+.$$

The absolute call and put price impacts are then computed as

$$\Delta C_{j,\varepsilon}(K,T) = \widehat{C}_{j,\varepsilon}(K,T) - \widehat{C}_0(K,T),$$

and

$$\Delta P_{j,\varepsilon}(K,T) = \widehat{P}_{j,\varepsilon}(K,T) - \widehat{P}_0(K,T).$$

Relative price changes are also computed:

$$R_C^{j,\varepsilon}(K,T) = \frac{\widehat{C}_{j,\varepsilon}(K,T) - \widehat{C}_0(K,T)}{\widehat{C}_0(K,T)},$$

and

$$R_P^{j,\varepsilon}(K,T) = \frac{\widehat{P}_{j,\varepsilon}(K,T) - \widehat{P}_0(K,T)}{\widehat{P}_0(K,T)}.$$

Since relative price changes can become unstable when the baseline option price is close to zero, log-price ratios are also reported:

$$L_C^{j,\varepsilon}(K,T) = \log\left(\frac{\widehat{C}_{j,\varepsilon}(K,T)}{\widehat{C}_0(K,T)}\right),$$

and

$$L_P^{j,\varepsilon}(K,T) = \log\left(\frac{\widehat{P}_{j,\varepsilon}(K,T)}{\widehat{P}_0(K,T)}\right).$$

The log-price ratio is useful when comparing sensitivities across strikes and maturities as it measures proportional price changes while being more interpretable than a raw percentage changes when option prices vary greatly in magnitude.

To compare strikes across experiments, strike levels are expressed using the moneyness measure, defined as

$$m = \frac{K}{S_0}.$$

Thus $m < 1$ corresponds to strikes below the initial asset price, $m = 1$ is at-the-money, and $m > 1$ corresponds to strikes above the initial asset price. Put prices are especially relevant for downside terminal moves, while call prices are more sensitive to upside terminal moves. Therefore, the option price analysis complements the left-tail and right-tail probability analysis.

4.7 Validation of the Monte Carlo Simulation Procedure

4.7.1 Fokker-Planck Validation

The Monte Carlo simulation is validated by comparing the simulated terminal log-return distribution with a density obtained from the Fokker–Planck equation of the Heston model. The Fokker–Planck equation describes the time evolution of the probability density associated with the continuous-time stochastic process. It can be used as a benchmark for checking whether the discretised simulation reproduces the distribution implied by the model.

The validation utilises the derivation done by Dragulescu and Yakovenko from their 2002 paper [11]. Since their analytical expression is defined for the centered log-return, we write

$$Y_t = X_t - (r-q)t.$$

Then Y_t should satisfy

$$dY_t = -\frac{1}{2}v_t\,dt + \sqrt{v_t}\,dW_t^S,$$

while the variance process remains

$$dv_t = \kappa(\theta - v_t)\,dt + \xi\sqrt{v_t}\,dW_t^v.$$

Let $p(t,y,v)$ denote the joint density of (Y_t, v_t). The corresponding Fokker–Planck equation is

$$\frac{\partial p}{\partial t} = \kappa\frac{\partial}{\partial v}\left[(v-\theta)p\right] + \frac{1}{2}\frac{\partial}{\partial y}(vp) + \rho\xi\frac{\partial^2}{\partial y\partial v}(vp) + \frac{1}{2}\frac{\partial^2}{\partial y^2}(vp) + \frac{\xi^2}{2}\frac{\partial^2}{\partial v^2}(vp).$$

The initial condition is

$$p(0,y,v) = \delta(y)\delta(v - v_0).$$

The terminal marginal density of the centered log-return is obtained by integrating out the variance variable,

$$f_{Y_T}^{FP}(y) = \int_0^\infty p(T,y,v)\,dv.$$

This marginal density is evaluated numerically using the Fourier representation of the Fokker–Planck solution. The Monte Carlo sample is converted to centred log-returns by

$$Y_T^{(m)} = X_T^{(m)} - (r-q)T, \qquad m = 1,\dots,M.$$

The empirical distribution of $Y_T^{(m)}$ is then compared with $f_{Y_T}^{FP}$. For histogram bins $B_i = [a_i, a_{i+1})$, the simulated probability of a path ending up at bin B_i is

$$\hat{p}_i^{MC} = \frac{1}{M}\sum_{m=1}^{M}\mathbf{1}_{\{Y_T^{(m)}\in B_i\}},$$

while the Fokker–Planck bin probabilities are

$$p_i^{FP} = \int_{B_i} f_{Y_T}^{FP}(y)\,dy.$$

The discrepancy between the simulated distribution and the one produced by the Fokker–Planck will be measured using the integrated absolute error, defined as

$$\text{IAE} = \sum_i \left|\hat{p}_i^{MC} - p_i^{FP}\right|.$$

A visual comparison is also made by plotting the Monte Carlo histogram and the Fokker–Planck density on the same axes. In addition, semilog plots are used to compare the agreement in the tails of the distribution. Agreement between the two densities indicates that the Monte Carlo scheme is reproducing the distribution of the continuous-time Heston model. Differences may arise from finite sample error, time-discretisation error in the Euler scheme, and numerical error in evaluating the Fourier integral. For a fixed maturity T, we expect that as the empirical distribution of log-returns will approach the Fokker–Planck density as the number of paths $M \to \infty$ and the time step $\Delta t \to 0$. This should be evidenced by the IAE measure decreasing as the number of simulated paths increases.

4.8 Hypotheses and Expected Outcomes

The experimental hypotheses are based on the two stochastic differential equations defining the Heston model, described above. They concern the expected direction and relative strength of parameter effects rather than the estimation of empirical market relationships.

Under the SDE for the variance process, the expected instantaneous variance at time t is

$$\mathbb{E}^{\mathbb{Q}}[v_t] = \theta + (v_0 - \theta)e^{-\kappa t}.$$

Consequently, the expected accumulated variance over the interval $[0, T]$ is

$$A_T(\boldsymbol{\Theta}) = \mathbb{E}^{\mathbb{Q}}\left[\int_0^T v_t\,dt\right] = \theta T + (v_0 - \theta)\frac{1 - e^{-\kappa T}}{\kappa}.$$

Since the terminal log-return satisfies

$$X_T = (r - q)T - \frac{1}{2}\int_0^T v_t\,dt + \int_0^T \sqrt{v_t}\,dW_t^S,$$

its expected value is

$$\mathbb{E}^{\mathbb{Q}}[X_T] = (r - q)T - \frac{1}{2}A_T(\boldsymbol{\Theta}).$$

Derivations of expressions above are provided in Appendix A.2. Based on these, the following hypotheses can be made.

H1: Effects of the variance-level parameters. An increase in either the initial variance v_0 or the long-run variance θ is expected to increase accumulated variance. Consequently, it should generally increase the standard deviation and tail probabilities of X_T, while decreasing its mean through the term $-\frac{1}{2}A_T(\boldsymbol{\Theta})$. Because European option payoffs are convex, increases in these variance-level parameters are also expected to increase call and put values in most of the investigated configurations.

H2: Dependence on maturity and mean reversion. The influence of v_0 is expected to be strongest at shorter maturities and under slow mean reversion, since the initial variance remains influential for longer. The influence of θ is expected to increase with maturity and with the speed of mean reversion, as the variance process becomes more strongly anchored to its long-run level.

When $v_0 = \theta$, the expected accumulated variance simplifies to

$$A_T(\mathbf{\Theta}) = \theta T,$$

and is independent of κ. In such configurations, perturbations of κ are not expected to have a substantial effect on the mean log-return. When $v_0 > \theta$, as in baseline C, increasing κ is expected to reduce accumulated variance by pulling the variance more quickly toward its lower long-run level.

H3: Impact of the volatility-of-volatility parameter. The parameter ξ does not enter the expression for expected accumulated variance and is not expected to have a first-order effect on the mean of X_T. Increasing ξ, however, makes the variance process more variable. It is expected that increasing the value of this parameter will affect kurtosis of the distribution and tail ratios, particularly in configurations with high variance or correlation cases.

H4: Impact of the correlation parameter. The correlation parameter ρ is expected primarily to affect the asymmetry of the terminal log-return distribution. More negative values of ρ couple negative asset-price shocks with positive variance shocks and are therefore expected to result in a more skewed distribution of terminal asset prices and by extension more extreme outcomes on one side of the distribution. Baselines A and E are useful here as they differ only in the value of ρ.

H5: Local linearity of parameter responses. For sufficiently small perturbations, smooth output metrics are expected to respond linearly to changes in a single parameter. Non-linear changes are expected for skewness, excess kurtosis, tail probabilities, and out-of-the-money option prices since these quantities depend on relatively small regions of the terminal distribution. Nonlinearity is also expected to be more visible in high volatility baselines or those with high ξ.

H6: Regime-dependent parameter importance. No single parameter is expected to be most impactful for all model outputs and across all baseline configurations. The relative importance of v_0 and θ is expected to depend on maturity and mean reversion. Further, it is expected that ξ and ρ parameters will be more important for skewness and tail-related outputs. Differences between the first-order and total-order Sobol indices are expected to be larger for tail and option pricing outputs.

5 Experiments and Results

This chapter presents the numerical results of the study. The Monte Carlo implementation is first validated against the Fokker–Planck density. Subsequent sections present the impact of Heston parameter perturbations on the terminal log-return distribution, tail probabilities, and European option prices across the baseline configurations and maturities are presented. Finally, results of variance decomposition using Sobol indices is used to assess the relative importance of the parameters and their interactions when they are varied jointly.

5.1 Experimental Setup

The perturbation experiment was conducted around the six baseline parameter configurations introduced in Section 4.3. For each baseline and maturity 1,000,000 Monte Carlo paths were generated using the Euler scheme. The maturity grid is

$$T \in \{0.05,\ 0.10,\ 0.25,\ 0.50,\ 1.00,\ 2.00\},$$

with 252 simulation steps per year corresponding to the amount of trading days in one calendar year. This gives

$$N \in \{13,\ 25,\ 63,\ 126,\ 252,\ 504\}$$

for the respective maturities. Throughout the experiment, The starting asset price S_0 is set to 100 and interest rate is set to 0: $r = q = 0$.

Perturbations of ρ were additive over the same numerical grid. The $\rho = -1$ boundary implied by the -0.20 perturbation in case C was excluded. The resulting design contained 3,594 admissible parameter-perturbation experiments. Common random numbers were used for the baseline and perturbed simulations within each case and maturity, reducing Monte Carlo noise in the estimated changes.

Distributional effects were evaluated using total variation distance, together with changes in the probabilities of crossing the baseline one-percent left- and right-tail thresholds. European calls and puts were valued across different levels of moneyness, defined as the ratio of the strike price K to the initial underlying asset price S_0. A call option is in the money (ITM) when $K/S_0 < 1$ and out of the money (OTM) when $K/S_0 > 1$. For a put option, the relationship is reversed. It is in the money when $K/S_0 > 1$ and out of the money when $K/S_0 < 1$. Options with $K/S_0 = 1$ are referred to as at the money.

For this thesis, options with moneyness ratios of

$$\frac{K}{S_0} \in \{0.8,\ 0.9,\ 1.0,\ 1.1,\ 1.2\}$$

have been considered.

To provide a compact comparison, the subsequent discussion in Section 5.6 focuses on two moneyness ratios. of At $K = 0.90S_0$, the call is ITM and the put is OTM, whereas at $K = 1.10S_0$, the call is OTM and the put is in-the-money. These options are examined at maturities of $T = 0.05$, 1, and 2.

Although terminal log-returns were simulated, option values are calculated from the terminal price levels,

$$S_T = S_0 \exp(X_T).$$

Consequently, the resulting option values are ordinary linear prices rather than log prices. Percentage price changes are used for economic interpretation in the main text, while log-price ratios are retained as a proportional sensitivity measure and robustness validation.

5.2 Validation Against the Fokker–Planck Density

In the validation run reported below, we set $r = q = 0$, as a result this transformation does not shift the simulated sample. The comparison is therefore made between the empirical distribution of Y_T and the Fokker–Planck density evaluated at the same maturity T.

For the validation experiment, the baseline parameter vector is

$$(v_0, \kappa, \theta, \xi, \rho) = (0.04, 1.5, 0.08, 0.4, -0.4),$$

with maturity $T = 1$, $M = 10^6$ Monte Carlo paths, and $N = 504$ time steps. The simulation uses the same full truncation Euler scheme as in the main perturbation experiments.

The validation results are shown in Table 3. The Monte Carlo histogram and the Fokker–Planck density are visually close, and the estimated total variation distance is small. For the main validation run with $M = 10^6$ paths, the integrated absolute error is 0.00869, corresponding to a total variation distance of 0.004345. Thus, the two distributions differ by less than approximately 0.5% in total variation distance.

The probability mass covered by the numerical grid is also close to one for both methods. The Monte Carlo sample has 0.999917 of its mass inside the grid, while the Fokker–Planck density integrates to approximately 1.000183 over the same grid. This indicates that the grid captures essentially all relevant probability mass and that the numerical Fourier inversion is well-normalised up to small numerical integration error.

M	N	IAE	$\widehat{d}_{\mathrm{TV}}$	MC mass on grid	FP mass on grid
100,000	504	0.019890	0.009945	0.999860	1.000183
1,000,000	504	0.008690	0.004345	0.999917	1.000183
2,000,000	504	0.007463	0.003731	0.999893	1.000183
10,000,000	504	0.006361	0.003180	0.999902	1.000183

Table 3: Validation of the Monte Carlo terminal log-return distribution against the Fokker–Planck density. The discrepancy decreases as the number of Monte Carlo paths increases, while the time discretisation is kept fixed at $N = 504$.

The decrease in both IAE and total variation distance as M increases is consistent with Monte Carlo sampling error decreasing with the number of simulated paths. Overall, the validation supports that the Monte Carlo implementation reproduces the terminal return distribution implied by the continuous-time Heston model to a sufficient degree of accuracy for the subsequent perturbation experiments.

5.3 Baseline Terminal Log-Return Distribution

Table 4 summarises the baseline terminal log-return distributions at $T = 1$. These distributions provide the reference points relative to which the parameter-perturbation effects are evaluated.

The strongest contrast between baselines can be seen when comparing cases B and C. Case B has the smallest dispersion, with a standard deviation of 0.1424, and only moderate departures from Gaussianity. Case C has the largest standard deviation, 0.3121, together with skewness of -2.328 and excess kurtosis of 8.420. Its one-percent quantile is -1.2311, whereas its 99-percent quantile is

Table 4: Summary statistics of the baseline terminal log-return distributions at $T = 1$. Each row is estimated using 1,000,000 Monte Carlo paths and 252 time steps. The final column reports the fraction of paths with a negative raw Euler-updated terminal variance before full truncation was applied.

Case	Mean	Std. dev.	Skewness	Ex. kurt.	$q_{0.01}$	$q_{0.99}$	Neg. raw v_T frac.
A	−0.0200	0.2060	−0.921	1.578	−0.6396	0.3458	0.00013
B	−0.0101	0.1424	−0.308	0.461	−0.3804	0.3000	0.00000
C	−0.0394	0.3121	−2.328	8.420	−1.2311	0.3556	0.14079
D	−0.0198	0.2074	−1.146	2.284	−0.6741	0.3308	0.00689
E	−0.0201	0.2002	−0.063	0.638	−0.5193	0.4635	0.00016
F	−0.0111	0.1483	−0.025	0.354	−0.3695	0.3421	0.01002

only 0.3556. Case C therefore represents both a higher-dispersion regime and a substantially more asymmetric, heavy-tailed return distribution.

Cases A and E best illustrate the impact of the correlation parameter as all other parameters are identical. Their standard deviations are relatively similar, at 0.2060 and 0.2002, but their skewness differs considerably: −0.921 in case A compared with −0.063 in case E. The one-percent quantile is also more negative in case A, at −0.6396, compared with −0.5193 in case E. This is consistent with negative return-variance correlation producing greater left-tail asymmetry while having a smaller effect on overall dispersion.

The final column in Table 4 records the fraction of paths for which the raw Euler-updated terminal variance was negative before full truncation was applied. These raw values were not inserted into the drift or diffusion terms, which use

$$v_t^+ = \max(v_t, 0).$$

The fraction is negligible in most cases but reaches 14.08% in case C. This is quite high and is considered as a numerical limitation when interpreting results for the stressed configuration.

5.4 Effect of Parameter Perturbations on Distribution Summaries

For each admissible one-parameter perturbation, the resulting terminal log-return distribution is compared with its corresponding baseline using the total variation distance defined in Section 5.4. The expanded experiment contains 3,594 feasible combinations across the six baselines, six maturities, five parameters and the non-zero perturbation grid. The perturbation $\epsilon = -0.20$ was omitted for ρ in case C because it reaches the boundary value $\rho = -1$. Table 5 reports the mean total variation distance for each baseline and perturbed parameter, averaged over all tested maturities and feasible perturbation sizes. The main measure used in this section to measure the discrepancy between two distributions is the total variation distance. For two distributions with bin probabilities p_i and q_i, the empirical total variation distance is computed as

$$d_{\mathrm{TV}}(p, q) = \frac{1}{2} \sum_i |p_i - q_i|.$$

This gives a direct measure of how much probability mass must be moved to transform one distribution into the other. A value close to zero indicates that the perturbed distribution is close to the baseline distribution, while larger values indicate a stronger distributional effect.

Table 5: Mean total variation distance between the baseline and perturbed terminal log-return distributions. Values are averaged over the tested maturities and feasible perturbation sizes. Bold entries identify the largest value within each baseline.

Case	v_0	κ	θ	ξ	ρ
A	**0.0218**	0.0073	0.0144	0.0139	0.0179
B	**0.0192**	0.0065	0.0153	0.0102	0.0139
C	0.0322	0.0090	0.0122	0.0285	**0.0421**
D	**0.0281**	0.0050	0.0085	0.0163	0.0216
E	**0.0209**	0.0062	0.0135	0.0105	0.0155
F	0.0137	0.0095	**0.0214**	0.0119	0.0176

Note: Each cell contains 120 observations, except the case C correlation cell, which contains 114. The perturbation $\epsilon = -0.20$ was excluded for this cell because it would set $\rho = -1$.

Several findings can be highlighted based on the Table 5. Initial variance v_0 produces the largest mean distributional change in cases A, B, D and E. Correlation is more important in the stressed case C, where its mean total variation distance of 0.0421 is the largest entry in the table. In the fast mean-reverting empirical benchmark F, the long-run variance θ is more significant. Mean-reversion speed κ produces the smallest mean distance in every baseline. Volatility of volatility ξ is not a meaningful parameter in any of the six regimes.

Further, Table 6 shows how the parameter importance changes depending which maturity is being considered. At short maturities, v_0 produces the largest distributional change in all six baseline configurations. its case-averaged total variation distance is 0.0277 at $T = 0.05$, compared with 0.0125 for ρ and 0.0058 for θ. Its influence subsequently declines, reaching 0.0144 at $T = 2$. By contrast, the effect of θ increases from 0.0058 to 0.0244, while that of ρ increases to 0.0261. These results are consistent with the diminishing impact of the initial variance and the increasing importance of long-run variance as the horizon increases.

Table 6: Mean total variation distance by maturity and perturbed parameter. Values are first averaged over perturbation sizes within each baseline and then averaged equally across the six baselines. Bold entries indicate the largest value at each maturity.

T	v_0	κ	θ	ξ	ρ
0.05	**0.0277**	0.0026	0.0058	0.0100	0.0125
0.10	**0.0269**	0.0039	0.0080	0.0119	0.0154
0.25	**0.0253**	0.0063	0.0117	0.0152	0.0214
0.50	0.0228	0.0082	0.0154	0.0176	**0.0257**
1.00	0.0189	0.0103	0.0201	0.0185	**0.0275**
2.00	0.0144	0.0122	0.0244	0.0180	**0.0261**

Another finding is the impact of perturbing the correlation parameter ρ in a high volatility baseline C. The largest observed total variation distance is 0.1559, obtained at $T = 1$ when a perturbation of -0.18 changes ρ from -0.8 to -0.98. At the positive perturbation of $+0.18$, the distance is only 0.0756. The negative perturbation therefore produces approximately twice the distributional change of the positive perturbation of equal magnitude. Thus, there appears to be asymmetric and non-linear response as correlation approaches its lower boundary, however, a general non-linearity cannot be claimed as this applies only to one particular baseline in this experiment.

Overall, the results indicate that the ranking of parameter importance differs depending on the maturity. Initial variance is most important at short horizons, while long-run variance and correlation become more influential at longer horizons. Correlation is particularly important in stressed and persistent-volatility regimes, whereas κ remains comparatively weak throughout the experiment. Since total variation measures the magnitude but not the direction of a distributional change, the left- and right-tail results in the following subsection are used to identify where the probability mass is redistributed.

5.5 Effect of Parameter Perturbations on Tail Probabilities

Metrics related to tail probabilities are evaluated using the fixed baseline thresholds defined in the Section 4.6.2. Consequently, the baseline probability of exceeding either threshold is $p_L^0 = p_R^0 = 0.01$. The probability ratios

$$R_L(\epsilon) = \frac{p_L^\epsilon}{p_L^0}, \qquad R_R(\epsilon) = \frac{p_R^\epsilon}{p_R^0}$$

measure whether perturbing a chosen parameter increases the probability of observing a price path that produces a more extreme terminal asset price. A ratio above one represents an increase in more extreme outcomes, whereas a ratio below one represents a decrease.

Table 7 reports the median probability ratios at perturbations of $\epsilon = -0.10$ and $+0.10$. The medians are calculated over the six baselines and six maturities. This perturbation magnitude is sufficiently large to reveal the direction of the response while avoiding the zero-count observations produced by the most extreme correlation perturbation.

Table 7: Median left- and right-tail probability ratios at $\epsilon = \pm 0.10$. For v_0, κ, θ, and ξ, the perturbation is relative; for ρ, it is additive.

Parameter	$R_L(-0.10)$	$R_L(+0.10)$	$R_R(-0.10)$	$R_R(+0.10)$
v_0	0.8569	1.1518	0.7908	1.2495
κ	1.0190	0.9820	0.9996	1.0013
θ	0.9530	1.0488	0.9195	1.0909
ξ	0.9371	1.0610	1.0399	0.9626
ρ	1.1149	0.8883	0.8251	1.1747

Increasing v_0 raises both tail probabilities in every tested configuration. At $\epsilon = 0.10$, the median left-tail probability ratio increases by approximately 15.2%, while the median right-tail ratio increases by approximately 24.9%. Increasing θ also raises both tail probabilities, although its typical effect is smaller when short and long maturities are aggregated across the six maturities to obtain an overall summary. By contrast, perturbations to κ produce ratios close to one, confirming that its effect on the tail probabilities is comparatively weak.

Changing the correlation parameter does not influence the left and right tail probability ratios equally. Increasing ρ reduces the median left-tail ratio by approximately 11.2% and increases the median right-tail ratio by approximately 17.5%. On the other hand, Decreasing ρ produces the opposite effect. This direction is observed for every tested baseline, maturity and feasible perturbation magnitude. Increasing ξ also consistently raises the left-tail probability ratio. The baseline correlation value plays a significant role: in the negatively correlated cases A-D, increasing ξ reduces the right-tail probability, whereas in the zero-correlation cases E and F it increases both tails.

Table 8 demonstrates that parameter importance changes substantially depending on maturity. Initial variance v_0 produces the largest response in both tails at $T = 0.05$, however, its influence declines as maturity increases. Conversely, the effect of θ strengthens with maturity and is largest for the left tail at $T = 2$. Correlation becomes particularly important at medium and long maturities and produces the largest right-tail response at both $T = 0.5$ and $T = 2$.

Table 8: Baseline-averaged tail probability ratios following positive parameter perturbations at selected maturities

	$T = 0.05$		$T = 0.5$		$T = 2$	
Parameter	$\bar{R}_L$	$\bar{R}_R$	$\bar{R}_L$	$\bar{R}_R$	$\bar{R}_L$	$\bar{R}_R$
v_0	**1.255**	**1.348**	1.131	1.227	1.064	1.106
θ	1.016	1.019	1.084	1.126	**1.162**	1.231
κ	0.997	0.999	0.976	1.003	0.951	1.032
ξ	1.044	0.957	1.085	0.950	1.086	0.941
ρ	0.919	1.140	**0.868**	**1.379**	0.865	**1.396**

Notes: $\bar{R}_L$ and $\bar{R}_R$ denote tail probability ratios averaged across the six baseline cases. A value of one indicates no change relative to the baseline tail probability. Bold values identify the largest absolute deviation from one within each combination. For v_0, θ, κ, and ξ, the perturbation is a 10% relative increase; for ρ, it is an additive increase of 0.10.

In baseline C, When ρ is reduced from -0.8 to -0.98, no right-tail exceedances have been observed among the 10^6 paths at $T \in \{0.5, 1, 2\}$. This should not be interpreted as proof that the event is impossible as it may be the case that 10^6 paths is not enough for a rare terminal asset price to appear in the generated data. In the opposite direction, increasing ρ to -0.62 raises the right-tail probability at $T = 2$ from 1% to 2.795%. These results mainly illustrate the impact of the initial correlation value on the derived results.

5.6 Effect of Parameter Perturbations on Option Prices

To measure the impact of parameter perturbations on option prices, the absolute and percentage price changes will be used. These two measures must be interpreted together, since percentage changes can become very large when the baseline option price is close to zero. The uncertainty of ΔV is measured using the paired Monte Carlo standard error.

Table 9 reports percentage changes for an OTM put with $K/S_0 = 0.9$ and an OTM call with $K/S_0 = 1.1$. These contracts are used primarily used for comparisons as their Monte Carlo standard errors remain below 5% in all six cases and at every maturity.

Table 9: Case-averaged percentage option price changes following positive parameter perturbations

	$T = 0.05$		$T = 0.5$		$T = 2$	
Parameter	OTM put	OTM call	OTM put	OTM call	OTM put	OTM call
v_0	**32.46**	**47.59**	**6.21**	**8.69**	2.30	2.87
θ	2.21	2.74	4.56	5.49	**5.59**	**6.32**
κ	-0.57	-0.30	-0.20	0.47	0.25	0.95
ξ	7.68	-2.69	0.24	-3.16	-0.85	-2.24
ρ	-12.85	22.32	-2.95	5.67	-0.45	2.74

Notes: Entries are percentage changes relative to the corresponding baseline option price, averaged across the six baseline cases. The OTM put and call have $K/S_0 = 0.9$ and $K/S_0 = 1.1$, respectively. For parameters v_0, θ, κ, and ξ, the perturbation is a 10% relative increase. For ρ, it is an additive increase of 0.10. Bold values indicate the largest absolute percentage change within each maturity and option type.

Initial variance v_0 has the greatest proportional effect at short maturities. At $T = 0.05$, a 10% increase in v_0 raises the case-averaged OTM put and call prices by 32.46% and 47.59%, respectively. Its influence declines with maturity, with corresponding changes of only 2.30% and 2.87% at $T = 2$. The long-run variance θ, on the other hand, drives a larger change in the option price as maturity increases. Its effects increase from 2.21% and 2.74% at $T = 0.05$ to 5.59% and 6.32% at $T = 2$, where it becomes the most influential parameter for both contracts.

A different situation can be observed for the correlation parameter. At $T = 0.05$, an additive increase of 0.10 in ρ reduces the OTM put price by 12.85%, while increasing the OTM call price by 22.32%. This is primarily due to the fact that most baselines had a negative initial correlation value. As a result, increasing the correlation parameter shifts the distribution of terminal asset prices to the right, thereby making the left tail of the distribution less significant. Consequently, the OTM put option prices decrease for both all three maturities shown in Table 9.

The impact from perturbing the mean reversion speed parameter κ are comparatively small, while the response to ξ depends on maturity and moneyness. Increasing ξ raises the short-maturity OTM put price but reduces the OTM call price. At longer maturities, its average effect is negative for both contracts. Overall, perturbing ξ or κ parameters does not produce a significant change in the option prices. A distinct generalised effect produced by changing these parameters cannot be assigned.

The response to short maturity option prices also displays some non-linearity. For example, at $T = 0.05$, a positive 10% perturbation to v_0 raises the OTM call price by 47.59%, whereas a negative 10% perturbation reduces it by only 36.06%. The corresponding changes for the OTM put are 32.46% and -26.74%. At medium and long maturities, the responses to perturbations of the same magnitude are substantially closer to symmetric.

It should be noted that price impact on the deep OTM options are not interpreted. At $T = 0.05$, four baseline case-contract combinations for the $K/S_0 = 0.8$ put or $K/S_0 = 1.2$ call produced an estimated price of exactly zero because no simulated path generated a positive payoff.

5.7 Results of the Sobol Index Sensitivity Analysis

Case	κ	ρ	θ	v_0	ξ
A	0.051	0.100	0.462	0.268	0.199
B	0.040	0.090	0.669	0.165	0.091
C	0.030	0.193	0.089	0.460	0.306
D	0.008	0.128	0.059	0.638	0.230
E	0.031	0.134	0.504	0.310	0.090
F	0.044	0.106	0.812	0.011	0.051

Table 10: Mean total-order Sobol indices across the scalar output metrics for each baseline case. Larger values indicate greater overall contribution to output variation over the tested parameter ranges.

From Table 10, it is possible to see that the dominant Sobol index depends strongly on a chosen baseline. For instance, in the moderate equity-like case A, the long-run variance θ has the largest average total-order index, followed by v_0 and ξ. In the calm low-volatility case B, θ is more impactful, with the mean total-order index of 0.669. The same pattern appears in the no-leverage case E and the empirical benchmark case F, where θ is the main driver of output variation.

The case C behaves differently. Here, the initial variance v_0 has the largest average total-order index, followed by the volatility of volatility ξ. This indicates that, in the high-volatility regime, variation in the starting variance level and in the randomness of the variance process explains more output variation than the long-run variance level. The persistent-variance baseline D is impacted more by changes in v_0, with mean total-order Sobol index $S_T = 0.638$. This is consistent with slow mean reversion as when variance shocks decay slowly, the initial variance remains important over the maturity horizon.

Across all cases, κ remains the smallest average total-order Sobol index. This suggests that, with the experimental setup of this thesis, the speed of mean reversion explains less output variation than the variance level, volatility of volatility, or correlation parameters.

Case	Std. dev.	Total variation	Left tail	Right tail	Put rel. price	Call rel. price
A	θ (0.607)	θ (0.727)	ξ (0.348)	θ (0.416)	θ (0.618)	θ (0.565)
B	θ (0.815)	θ (0.952)	θ (0.603)	θ (0.723)	θ (0.802)	θ (0.790)
C	v_0 (0.782)	ξ (0.732)	ξ (0.557)	ρ (0.711)	v_0 (0.731)	v_0 (0.432)
D	v_0 (0.924)	v_0 (0.756)	v_0 (0.508)	v_0 (0.490)	v_0 (0.923)	v_0 (0.725)
E	θ (0.624)	θ (0.851)	θ (0.384)	θ (0.445)	θ (0.621)	θ (0.612)
F	θ (0.990)	θ (0.988)	θ (0.838)	θ (0.871)	θ (0.977)	θ (0.974)

Table 11: Most significant parameters by total-order Sobol index for selected output metrics. Each cell reports the parameter with the largest S_T, with the corresponding S_T value in parentheses. The put and call columns refer to relative price changes for the $K/S_0 = 0.90$ put and $K/S_0 = 1.10$ call.

Table 11 shows that long-run variance θ dominates most metrics in baselines B, E, and F. In the case F, this is particularly visible as θ has total-order indices above 0.83 for all reported metrics and above 0.97 for standard deviation, total variation, and both option price metrics. This can additionally reflect the significance of fast mean reversion for that baseline.

In case D, by contrast, the parameter v_0 explains the largest share of variance for all outputs. This is consistent with slow mean reversion: since variance reverts slowly, the initial variance level remains influential throughout the time horizon and affects both distributional and option pricing outputs. A different situation can be observed for baseline C where no single parameter explains the largest share of variance across all outputs. For this case, v_0 explains most of the variance in the standard deviation and option price responses, ξ is a more impactful parameter for total variation and the left tail probability ratio, while ρ explains largest share of changes in extreme outcomes on the right tail of the distribution.

6 Limitations of the Methodology

Firstly, the results provided in the section 5 do not represent an empirical estimate of parameter uncertainty. The six baseline specifications were selected to model distinct volatility regimes. However, these do not account for all conceivable market conditions. Further, the perturbation range of $\pm 20\%$ is an experimental design choice rather than a confidence interval for parameter estimation. Consequently, the conclusions should be interpreted conditional on the selected baselines, perturbation range, maturities and selected option types.

Perturbing only one parameter at a time provides additional interpretability as a response to predefined variables can be isolated. Using a different methodology design and perturbing multiple parameters at a time can lead to differing results. Sobol indices address joint variation in part, but their values still remain dependent on the selected parameter ranges, input distributions and sampling assumptions. Consequently, the findings from the variance decomposition study should be interpreted conditional to the experimental design of this thesis. Further, perturbations to the positive parameters are expressed as relative changes, whereas changes in ρ are additive. The resulting response magnitudes are therefore measured in different units and only serve to indicate the largest output response.

The number of Euler updates is scaled with maturity as

$$N(T) = \mathrm{round}(252T),$$

so that the time increment remains approximately $\Delta t = 1/252$ years. This corresponds to approximately one update per trading day. As a result, short-maturity simulations contain fewer updates. For instance, there are 13 updates at $T = 0.05$.

The use of a daily grid is primarily a computational design choice. Since the experiment was not repeated with higher amount of time updates, it was not established whether the stated results will differ when the amount of time updates is increased.

The Fokker–Planck comparison at $T = 1$ uses $N = 504$ time steps, while the associated convergence exercise varies the number of Monte Carlo paths with N held fixed. The comparison can be considered as evidence for sampling convergence and agreement under one representative configuration, but it does not constitute a systematic time-step convergence study across all possible baselines and maturities.

Truncating the negative variance at 0 prevents negative raw variance states from entering the drift and diffusion calculations, but may introduce additional bias if the amount of negative variance updates before the truncation is applied is high. This is particularly relevant for baseline C, for which the raw Euler variance state is negative at maturity in approximately 14.08% of paths at

$T = 1$ and 15.28% at $T = 2$. This means that additional caution is needed when interpreting results for Baseline C. A targeted experiment using another truncation scheme would be necessary to identify the impact of the scheme used in the current version of the methodology.

Although 10^6 paths substantially reduce ordinary Monte Carlo error, very extreme terminal asset prices can still be absent from produced data. For example, under the largest negative perturbation to ρ in baseline C, no right-tail exceedances are observed at $T \in \{0.5, 1, 2\}$. Consequently, zero occurrences of the terminal asset prices above a threshold is more indicative of the resolution of the simulation rather than a proof that the event is impossible. Moreover, the small numerical constant used when calculating logarithmic probability ratios converts such zero counts into finite but potentially very large sensitivity estimates. These observations should be treated as numerically unstable, with the underlying probability levels reported alongside the transformed sensitivities.

An analogous problem occurs for option price ratios when the baseline value is equal or close to zero, particularly for deep OTM contracts at short maturities. Percentage and logarithmic price changes can become undefined or disproportionately large even when the absolute price change is negligible. Under the current methodology, interpretation of the impact of parameter uncertainty for contracts with non-negligible baseline values and on maturities of at least six months should carry greater value. Absolute price changes and paired Monte Carlo uncertainty should be considered alongside relative or logarithmic changes. Conclusions concerning linearity are also conditional on the reporting scale since a response that is approximately linear in the option price level need not be linear after a logarithmic transformation.

Furthermore, total variation distance is estimated from discretised histograms and is affected by the selected bins and support. The tail measures depend on the baseline-defined 1% thresholds and describe changes relative to those thresholds rather than the complete shape of the perturbed distribution.

Finally, the experiment examines parameter uncertainty only within the standard Heston model. It does not quantify broader model risk arising from jumps, stochastic interest rates, time-varying parameters, transaction costs, early exercise or path-dependent payoffs. The assumption of $r = q = 0$ also separates parameter sensitivity from discounting and dividend effects. Only aggregate results were retained from the full experiment rather than the complete terminal path samples.

7 Conclusions and Further Research

This thesis investigated how perturbations to the Heston model parameters affect the terminal log return distribution, fixed-threshold tail probabilities, and European option prices. The central conclusion is that parameter uncertainty can substantially affect each of these outputs. Moreover, the effects of parameter uncertainty vary across the five Heston parameters, with some exerting a substantially greater influence than others. Overall, parameter importance depends on maturity, the baseline configuration, the output metric under consideration, and, for option prices, the moneyness of the contract.

The Monte Carlo implementation was first validated against the Fokker–Planck benchmark density. It was shown that the discrepancy between the two distributions decreased as the number of simulated paths increased, which provided support for the use of the simulation method in the subsequent sensitivity experiments. Each of the posed hypotheses posed in Section 4.8 can be evaluated as follows.

H1: Effects of the variance-level parameters. Hypothesis 1 is generally supported. Positive perturbations to v_0 and θ broadly increased the dispersion of terminal returns and raised both fixed threshold tail probability ratios. They also increased the prices of OTM puts and calls that could reliably be estimated in all tested baseline-maturity combinations. This is consistent with the increase in accumulated variance and the convexity of European option payoffs.

H2: Dependence on maturity and mean reversion. There is a strong support for Hypothesis 2. The influence of v_0 was greatest at short maturities and declined as the maturity increased. Conversely, the effect of θ strengthened with maturity. This effect was observed in the distributional distances, tail probabilities, and option price results. The baseline comparison produced the same pattern as initial variance and remained particularly important in the persistent variance scenario. On the other hand, long-run variance was most impactful in driving model outputs for the baseline F adopted from literature. Perturbations to κ generally produced the smallest effects, particularly when the initial variance v_0 was equal to the long-run variance θ.

H3: Effect of volatility of variance. Hypothesis 3 is supported with an important distinction. Perturbations to ξ affected skewness, excess kurtosis and tail probabilities more clearly than they affected the mean of the distribution. Increasing ξ consistently increased the left-tail probability ratio, however, its effect on the right tail depended on the baseline correlation. In the negatively correlated cases A-D, increasing ξ increased the left-tail ratio while reducing the right-tail ratio, which indicated additional downside risk. In the zero-correlation cases E and F, increasing ξ raised both tail ratios, which indicates a higher probability of observing extreme outcomes at both ends of the distribution. It can be therefore stated that the effect of ξ on tail behaviour is affected by ρ in a large part and does not have a uniform effect on both tails of the terminal returns distribution.

H4: Effect of correlation parameter. Hypothesis 4 is strongly supported. The comparison between cases A and E showed that removing negative correlation substantially reduced negative skewness. More generally, increasing ρ reduced the left-tail ratio and increased the right-tail ratio for every tested baseline, maturity, and feasible perturbation magnitude. Decreasing ρ produced the opposite response. The option price results provide a different view as increasing ρ raised OTM call prices while reducing OTM put prices.

H5: Local linearity of parameter responses. Hypothesis 5 is broadly supported, although the degree of linearity of the response to parameter perturbations depends on the output and option type. Most smooth distributional and reliably estimated option price responses were approximately linear over small perturbations. For OTM options with short maturity the response was disproportional, especially when the parameter v_0 was changed. The same could be observed for selected tail-related outputs. The results do not support a more general claim that options across all strikes and maturities will respond non-linearly to parameter perturbations.

H6: Regime-dependent parameter importance. The regime dependence component of Hypothesis 6 is supported. No parameter was most significant in driving changes for every output, maturity, and baseline configuration considered in this thesis. Initial variance was most important at short maturities and under slow mean reversion, long-run variance became more influential at longer maturities and under fast mean reversion. Correlation and volatility-of-volatility were more impactful for tail-related outputs. The Sobol results reinforced this conclusion, as θ was most significant for many outputs in baselines B, E, and F. Initial variance v_0 represented largest share of the variance for most outputs in the baseline D, and different parameters dominated different outputs in the high volatility case C.

Taken together, the presented findings answer the research question by showing that parameter uncertainty has different impact on the Heston model depending on the time horizon and the feature of the distribution being considered. Parameters related to variance primarily determine the overall dispersion of asset returns, while ρ and ξ mostly determine the skewness of the return distribution. These changes can produce a significant impact on the European option prices produced by the Heston model, particularly for OTM contracts at short maturities.

References

[1] Yacine Ait-Sahalia and Robert Kimmel. Maximum likelihood estimation of stochastic volatility models. *Journal of Financial and Quantitative Analysis*, 42(4):951–972, 2007.

[2] Robert Azencott, Yutheeka Gadhyan, and Roland Glowinski. Option pricing accuracy for estimated heston models. *arXiv: Mathematical Finance*, 2014.

[3] David S. Bates. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies*, 9(1):69–107, 1996.

[4] Eric Benhamou, Emmanuel Gobet, and Mohammed Miri. Time dependent Heston model. *SIAM Journal on Financial Mathematics*, 1(1):289–325, 2010.

[5] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

[6] Samuel N. Cohen and Martin Tegnér. European option pricing with stochastic volatility models under parameter uncertainty. In Samuel N. Cohen, István Gyöngy, Gonçalo dos Reis, David Siska, and Łukasz Szpruch, editors, *Frontiers in Stochastic Analysis—BSDEs, SPDEs and Their Applications*, volume 289 of *Springer Proceedings in Mathematics & Statistics*, pages 123–167. Springer, Cham, 2019.

[7] Rama Cont. Model uncertainty and its impact on the pricing of derivative instruments. Technical Report FC-04-02, Finance Concepts Working Paper, RiskLab/ETH Zurich, 2006.

[8] Michel Crouhy, Dan Galai, and Robert Mark. *Model Risk*. John Wiley & Sons, Ltd, 2008.

[9] Yiran Cui, Sebastián del Baño Rollin, and Guido Germano. Full and fast calibration of the heston stochastic volatility model. *European Journal of Operational Research*, 263(2):625–638, 2017.

[10] Emanuel Derman. Model risk. Quantitative strategies research notes, Goldman Sachs, April 1996.

[11] Adrian A. Dragulescu and Victor M. Yakovenko. Probability distribution of returns in the heston model with stochastic volatility. *Quantitative Finance*, 2(6):443–453, 2002.

[12] Alberto Elices. Models with time-dependent parameters using transform methods: Application to Heston's model. arXiv preprint arXiv:0708.2020, 2007. Revised October 2008.

[13] Futures Industry Association. ETD Volume – December 2025.

[14] Manfred Gilli and Enrico Schumann. Heuristic optimisation in financial modelling. *Annals of Operations Research*, 193(1):129–158, 2012.

[15] F.M.Y. Guillaume and W. Schoutens. Calibration risk: Illustrating the impact of calibration risk under the heston model. *Review of Derivatives Research*, 15(1):57–79, 2012.

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

[17] Yan Cai Jin Zheng. Modeling the implied volatility surface-: A study for sp 500 index option. *Research Journal of Applied Sciences, Engineering and Technology*, 5:1973–1977, 2013.

[18] Roger Lord, Remmert Koekkoek, and Dick van Dijk. A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2):177–194, 2010.

[19] Jiayi Luo and Cindy Long Yu. The application of symbolic regression on identifying implied volatility surface. *Mathematics*, 11(9):2108, 2023.

[20] Stephen Slivinski. Too interconnected to fail? the rescue of long-term capital management. *Region Focus*, (Summer):34–36, 2009.

[21] I. M. Sobol'. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, 55(1–3):271–280, 2001.

[22] María Suárez-Taboada, Jeroen A.S. Witteveen, Lech Grzelak, and Cornelis W. Oosterlee. Uncertainty quantification and heston model. *Journal of Mathematics in Industry*, 8(1):5, 2018.

[23] Arman Zadgar, Somayeh Fallah, and Farshid Mehrdoust. Deep learning-enhanced calibration of the heston model: A unified framework. *ArXiv*, 2025. preprint.

A Derivations

A.1 Log-return dynamics

Under the risk-neutral measure, the asset-price process in the Heston model is

$$dS_t = (r - q)S_t\,dt + \sqrt{v_t}S_t\,dW_t^S.$$

We define log-return as

$$X_t = \log\left(\frac{S_t}{S_0}\right).$$

Let

$$f(S_t) = \log\left(\frac{S_t}{S_0}\right).$$

The first and second derivatives of f with respect to S_t are

$$f'(S_t) = \frac{1}{S_t}, \qquad f''(S_t) = -\frac{1}{S_t^2}.$$

Applying Itô's lemma gives

$$dX_t = f'(S_t)dS_t + \frac{1}{2}f''(S_t)(dS_t)^2.$$

The quadratic-variation term is

$$(dS_t)^2 = \left(\sqrt{v_t}S_t\,dW_t^S\right)^2 = v_tS_t^2\,dt,$$

where $(dW_t^S)^2 = dt$, while terms containing dt^2 or $dt\,dW_t^S$ vanish. Therefore,

$$\begin{aligned}
dX_t &= \frac{1}{S_t}\left[(r-q)S_t\,dt + \sqrt{v_t}S_t\,dW_t^S\right] - \frac{1}{2S_t^2}\left(v_tS_t^2\,dt\right) \\
&= (r-q)\,dt + \sqrt{v_t}\,dW_t^S - \frac{1}{2}v_t\,dt.
\end{aligned}$$

Hence, the log-return dynamics are

$$\boxed{dX_t = \left(r - q - \frac{1}{2}v_t\right)dt + \sqrt{v_t}\,dW_t^S}.$$

A.2 Expected variance, accumulated variance, and expected log-return

The variance process in the Heston model is

$$dv_t = \kappa(\theta - v_t)\,dt + \xi\sqrt{v_t}\,dW_t^v.$$

Let

$$m(t) = \mathbb{E}^{\mathbb{Q}}[v_t].$$

Taking expectations and using the fact that the Itô integral has expectation zero gives

$$\frac{dm(t)}{dt} = \kappa\left(\theta - m(t)\right), \qquad m(0) = v_0.$$

Equivalently,

$$\frac{dm(t)}{dt} + \kappa m(t) = \kappa\theta.$$

Multiplying both sides by the integrating factor $e^{\kappa t}$ yields

$$\frac{d}{dt}\left(e^{\kappa t}m(t)\right) = \kappa\theta e^{\kappa t}.$$

Integrating from 0 to t gives

$$e^{\kappa t}m(t) - v_0 = \theta\left(e^{\kappa t} - 1\right).$$

Solving for $m(t)$ produces

$$\boxed{\mathbb{E}^{\mathbb{Q}}[v_t] = \theta + (v_0 - \theta)e^{-\kappa t}}.$$

The expected accumulated variance over the interval $[0, T]$ is defined as

$$A_T = \mathbb{E}^{\mathbb{Q}}\left[\int_0^T v_t\,dt\right].$$

Interchanging expectation and integration gives

$$\begin{aligned}
A_T &= \int_0^T \mathbb{E}^{\mathbb{Q}}[v_t]\,dt \\
&= \int_0^T \left[\theta + (v_0 - \theta)e^{-\kappa t}\right] dt \\
&= \theta T + (v_0 - \theta)\frac{1 - e^{-\kappa T}}{\kappa}.
\end{aligned}$$

Thus,

$$\boxed{A_T = \theta T + (v_0 - \theta)\frac{1 - e^{-\kappa T}}{\kappa}}.$$

Integrating the log-return dynamics from 0 to T, and using $X_0 = 0$, gives

$$X_T = (r - q)T - \frac{1}{2}\int_0^T v_t\,dt + \int_0^T \sqrt{v_t}\,dW_t^S.$$

Taking expectations and again using the zero expectation of the Itô integral gives

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}}[X_T] &= (r - q)T - \frac{1}{2}\mathbb{E}^{\mathbb{Q}}\left[\int_0^T v_t\,dt\right] \\
&= (r - q)T - \frac{1}{2}A_T.
\end{aligned}$$

Therefore,

$$\boxed{\mathbb{E}^{\mathbb{Q}}[X_T] = (r - q)T - \frac{1}{2}\left[\theta T + (v_0 - \theta)\frac{1 - e^{-\kappa T}}{\kappa}\right]}.$$

Finally, the expected deviation of variance from its long-run level satisfies

$$\mathbb{E}^{\mathbb{Q}}[v_t] - \theta = (v_0 - \theta)e^{-\kappa t}.$$

The mean-reversion half-life $\tau_{1/2}$ is the time required for the magnitude of this deviation to decrease by one half. Therefore,

$$e^{-\kappa \tau_{1/2}} = \frac{1}{2},$$

which gives

$$\boxed{\tau_{1/2} = \frac{\log 2}{\kappa}}.$$

A.3 Numerical and reproducibility details

The one-parameter perturbation experiment and the Sobol analysis used the same Euler discretisation with full truncation, but different computational designs. Table 12 records the settings used to generate the reported results. Common random numbers were used within each case-maturity block. The baseline and perturbed simulations used the same sequences of underlying standard normal shocks. This reduced Monte Carlo noise in the estimated differences.

Table 12: Numerical implementation and reproducibility settings.

Component	Setting
Python packages and precision	The fine-grid experiment was implemented in Python 3.12.3 using NumPy 2.1.2, pandas 3.0.3, and SciPy 1.18.0. The Sobol notebook used Python 3.14.4 and the SALib Sobol interface. Simulation states and Gaussian samples were stored as 64-bit floating-point arrays.
One-parameter design	For each of the six baselines, 10^6 paths were simulated at $T \in \{0.05, 0.10, 0.25, 0.50, 1.00, 2.00\}$, with $N_T = \text{round}(252T) \in \{13, 25, 63, 126, 252, 504\}$. For $p \in \{v_0, \kappa, \theta, \xi\}$, p was replaced by $p(1+\delta)$; for ρ, it was replaced by $\rho + \delta$. The non-zero grid was $\delta \in \{-0.20, -0.18, \ldots, 0.18, 0.20\}$. Excluding the infeasible $\rho = -1$ setting in case C left 3,594 perturbed simulations.
Random-number control	A fixed seed was used to make the experiment reproducible. Within each case and maturity, the same seed was used for the baseline and perturbed simulations, thereby reducing random variation in their comparison.
Distribution and tail threshold details	Standard deviations used one degree-of-freedom correction. Skewness and excess kurtosis used the bias-corrected pandas definitions, with Fisher excess kurtosis. Quantiles used linear interpolation. Tail thresholds were the baseline 1% and 99% quantiles, with strict exceedance events. Total variation used 600 equal-width density bins on a common, pair-specific range: the pooled 0.001 and 0.999 quantiles, padded on each side by 5% of their distance. Wasserstein distance was the empirical first Wasserstein distance. Probability ratios and logarithms used $\varepsilon = 10^{-15}$.
Sobol input domain and sample	The five inputs were independent uniforms. Positive parameters varied by $\pm 20\%$ around the applicable baseline and ρ varied additively by ± 0.10; no Feller-condition rejection was applied. A scrambled Sobol–Saltelli design with base size $N = 512$, sampling seed 123, and second-order indices disabled produced $512(5+2) = 3,584$ parameter vectors per case, or 21,504 in total.
Sobol simulation and estimation	The Sobol experiment was conducted only at $T = 1$, using 50,000 paths and 252 Euler steps per parameter vector. Within case $j = 1, \ldots, 6$, all evaluations shared seed $43 + 10^6 j + 10000$. Tail settings were $\alpha = 0.01$ and $k = 2$. Total variation used 300 baseline-defined bins spanning the baseline 0.001 to 0.999 quantiles with 10% padding and infinite outer edges. SALib first- and total-order indices used 1,000 bootstrap resamples, although the bootstrap seed was not fixed. Second-order indices and independent Monte Carlo replications were not computed.