Efficient Calibration and Greeks Estimation for a Panel of American Options Using a Stochastic Dynamic Program

Manal Teto

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science Mathematics and Statistics[1]

Department of Mathematics and Statistics
Faculty of Science
University of Ottawa

[1]The M.Sc. program is a joint program with Carleton University, administered by the Ottawa-Carleton Institute of Mathematics and Statistics

Abstract

American-style options are financial derivatives that offer the flexibility of early exercise opportunities. This feature poses the difficulty of solving a dynamic optimization problem to determine the optimal exercise strategy.

One significant advantage of the Stochastic Dynamic Program (SDP) is that its solution yields numerical approximations of option prices and sensitivities across the entire state-space partition. Through leveraging the homogeneity property, we efficiently value options with varying moneyness and maturity levels and conduct a calibration to market data. The methodology is versatile and applicable to various market dynamics, showcasing the substantial benefits of our SDP-based valuation method.

This research proposes an efficient method for pricing a panel of American-style options integrating SDP to overcome computational challenges associated with the estimation of Greeks and efficient calibration to market data. A key contribution of this study is the analysis of the convergence of Greek estimates within the SDP framework, ensuring stability and accuracy. Additionally, we introduce a novel calibration method based on the Feynman-Kac Theorem, where the objective is to find a set of Greeks that satisfy the partial differential equation. This calibration method provides an alternative approach to standard calibration techniques.

Acknowledgement

First and foremost, I would like to express my deepest gratitude to Professor François-Michel Boire for his invaluable guidance throughout this thesis. From conceptualizing an applied thesis aligned with my interests to providing constant support in navigating the complexities of coding and implementing rigorous methodologies, your dedication to my success has been truly appreciated. Your mentorship has been instrumental, and I am deeply thankful for it.

I would also like to extend my sincere thanks to Professor Rafal Kulik and Professor Yunran Wei for accepting to serve as examiners for this thesis. I am grateful for the time and effort you devoted to reviewing my work, and I truly appreciate your thoughtful feedback and suggestions.

A special thank you to the Alliance Canada team for your quick responses and for providing the computational resources that were crucial in conducting the experiments for this thesis.

I am immensely grateful to my family for their unwavering support throughout my academic journey, particularly during my Master's degree. Your belief in me and constant encouragement has been a pillar of strength. Thank you for always being there and for listening, no matter the circumstances.

Contents

List of Figures

List of Tables

List of Symbols

α_i^m, β_i^m	Local coefficients - SDP	16
$\bar{v}_m(s)$	Optimal exercise process	17
Δt	Finite time intervals	10
Δ	Delta of an option	2
ϵ	Residuals	81
Γ	Gamma of an option	2
$\hat{\sigma}$	Estimated volatility	81
$\mathbb{I}$	Indicator function	16
$\mathbb{Q}$	Risk-neutral measure	2
$\mathcal{V}$	Vega of an option	2
$\nabla f(\sigma)$	Gradient of the objective function $f(\sigma)$	63
Φ	Cumulative density function	15
ρ	Discount factor	17
σ	Volatility of the underlying asset	10
τ	Optimal stopping time	10
Θ	Theta of an option	2
$\tilde{V}$	American option market price	63
$\tilde{v}$	Approximated option value	16
$\tilde{v}_h^{\sup}$	Convex upper envelope approximation of the option	13
a_i	Asset price points in $\mathcal{S}$	14
A_{ki}, B_{ki}	Transition matrices	16
dt	Infinitesimal increment	10
$dW^{\mathbb{Q}}(t)$	Increment of the Brownian motion	10
$f(\sigma)$	Objective function - L-BFGS-B	62
G	Finite grid - SDP	13
$h(S_t)$	Payoff function	10
I	Horizontal intervals in $\overline{G}$	94
J	Vertical intervals in $\overline{G}$	94
K	Strike price of the option	10
N_t	Total number of time steps in finite grid G	81
N_{ex}	Number of exercise opportunities	81
q	Dividend yield	10

r	Risk-free rate	10
S_0	Initial asset price	16
S_0/K	Moneyness level	81
S_t	Price process of an underlying asset	9
T	Maturity date of the option	10
$T_{\max}$	Longest maturity date in panel	14
V	Value of the American option	10
v^e	Exercise value of the option	17
V_{eur}	European call option	40
w	Weights	81
$W^{\mathbb{Q}}(t)$	Randomness in price movement	10
X	Independent variable in regression	81
Y	Dependent variable in regression	81
Z	Standard normal random variable	15
$\hat{\mathbb{V}}$	Estimated variance	81
$\overline{G}$	FD Grid	94
$\overline{V}$	Derivative product	94
σ_0	Volatility initial guess	61
BT	Binomial tree	5
FD	Finite Difference	5
L-BFGS-B	Limited-memory Broyden–Fletcher–Goldfarb–Shanno	63
LSM	Least Squares Monte Carlo	7
PDE	Partial differential Equation	5
RBC	Royal Bank of Canada	59
SDP	Stochastic Dynamic Program	1
SOR	Successive Over-Relaxation	19
ATM	At-The-Money	25
BP	Breusch-Pagan	81
DTM	Days-To-Maturity	60
GBM	Geometric Brownian motion	81
GLS	Generalized Least Squares	81
ITM	In-The-Money	96
IV	Implied volatility	81
O	Order of convergence of the approximation	20
OTM	Out-of-the-money	25
RMSE	Root Mean Squared Error	64
SSE	Sum of Squared Errors	59
US	American	81
WLS	Weighted Least Squares	81
WRDS	Wharton Research Data Services	59

Introduction

Options are financial derivatives that provide their holder with a future payoff that is linked to the price of a specific underlying asset or a group of assets that can include stocks, bonds, currencies, interest rates, market indexes or commodities. These financial instruments are mainly used for hedging risk or speculative purposes. In the frame of frictionless markets, the no-arbitrage principle allows to derive the price of a derivative as the mathematical expectation of all its possible discounted future payments with respect to a so-called risk-neutral probability measure $\mathbb{Q}$.

European-style options can only be exercised at maturity, whereas American-style options provide the holder with the flexibility to exercise at any time up to and including the maturity date, whenever it is advantageous to do so. This fundamental difference significantly influences their valuation methods. Under the Black and Scholes (1973) and Merton (1973) model, standard European options can be valued using a closed-form solution. Similarly, exact Greeks for these options are easily calculated.

In contrast, American options do not admit such explicit analytical solutions due to the optionality embedded in the right to exercise the option at any point before or at expiration. This feature introduces a significant complexity to the valuation, as one must consider a multitude of possible exercise times and scenarios. The optimal exercise strategy depends on the evolving value of the underlying asset and requires a continual comparison between the immediate exercise value and the expected discounted value of holding the option for future exercise. Thus, this complexity makes deriving a straightforward pricing formula unfeasible, unlike the closed-form solution for European options.

In this thesis, I study a pricing methodology for a panel of American-style options with varying moneyness and maturities using a Stochastic Dynamic Program (SDP) and examine the convergence of different methods for estimating the Greeks. The contributions of this work are twofold. First, it provides a detailed convergence study of Greek estimates within the SDP framework, ensuring their numerical stability across different option characteristics. Second, it introduces an alternative calibration approach based on an iterated regression of the Feynman-Kac (FK) partial differential

equation, which provides a theoretically consistent way to infer volatility without directly minimizing pricing errors.

I concentrate on American put options and assume that the exercise opportunities are restricted to a finite set of equally spaced time points. These options are also known as *Bermudan options*. I address the determination of the optimal exercise strategy with SDP, where the Bellman value function is approximated using selected piecewise linear interpolations at each decision date focusing on the transition parameters of the underlying asset.

One significant advantage of this program is that it naturally provides the Delta (Δ) of an option, or the first derivative with respect to the underlying asset price. Other Greeks, such as Gamma (Γ), Theta (Θ), and Vega ($\mathcal{V}$), can be extracted using finite difference methods based on the computed option prices for a panel of options. The sensitivities of these Greeks are discussed in Section (1.4). In a single computation, SDP determines the prices for a panel of options along with the corresponding Greeks.

A key takeaway from these computations is the high accuracy obtained when pricing a panel of American options within less than 5 minutes. The largest error in option prices remains below 0.3%, while the errors in the Greeks are also very small, with a maximum of 0.1% for Delta, 1% for Gamma, 0.3% for Theta, and 0.1% for Vega. These results demonstrate the robustness of the method in efficiently computing option prices and their sensitivities.

To implement the SDP into pricing a panel of options, we use option price homogeneity, a property shared by most of the option pricing models used nowadays. As demonstrated by Joshi (2001), this characteristic generally holds for models where the asset price follows a stochastic process with log-normal dynamics, meaning that increments of the logarithm of the asset price are independent of its current level under the risk-neutral measure $\mathbb{Q}$. This property is true for a wide range of models, including the geometric Brownian motion of Black and Scholes (1973), the compound jump-diffusion process of Merton (1976), the Variance Gamma process of Madan et al. (1998), and the stochastic volatility model of Heston (1993).

However, price homogeneity can fail when there are market imperfections such as transaction costs, stochastic interest rates, or liquidity constraints, which introduce non-linearities that disrupt the predictable scaling relationship between the underlying asset price, strike price, and the resulting option price. The benefit of this homogeneity property is that, since early exercise strategies are independent of the initial value of the state variables and only depend on current values and strike prices, the same optimal strategy can be used for all options in the panel.

Solving the SDP approximates both the option value and the optimal exercise

strategy, resulting in a standardized grid that already contains option prices for various maturities and asset prices. By leveraging the homogeneity property of option pricing models, we transform the problem of pricing options with different strike prices into an equivalent problem where all options share the same strike price. This transformation allows us to price all options simultaneously within the same computational framework. The pre-computed grid enables efficient interpolation to determine the prices of different options. Thus, the SDP approach not only provides a robust and efficient solution for pricing American-style options but also facilitates the efficient pricing of a panel of such options.

In addition, I study the sensitivities of the options which are critical for risk management and hedging strategies. We test several numerical approximation methods across a wide range of option characteristics to determine which approach offers the ideal balance between numerical accuracy and computational cost.

Finally, a calibration procedure was performed using the SDP framework to estimate the volatility parameter σ in two distinct ways. The first calibration method directly minimizes the Sum of Squared Errors (SSE) between market-observed and SDP-predicted option prices to ensure that the model prices closely match market prices. The second calibration method leverages the Feynman-Kac Theorem, using the SDP-computed Greeks to construct an estimable relationship between option prices and sensitivities. This approach does not seek to minimize pricing discrepancies but rather derives an implied volatility that ensures consistency with the theoretical pricing PDE. The resulting estimate can be particularly relevant in contexts where accurate sensitivities are required, such as hedging applications.

Both calibration procedures are applied to 1354 American-style call and put options on Royal Bank of Canada (RBC) stock, traded in 2017. The dataset is sourced from OptionMetrics via Wharton Research Data Services (WRDS), filtering for options with at least 30 days to maturity and a moneyness range between 0.9 and 1.1 to ensure stability and liquidity. Market prices are derived from the midpoint of bid-ask quotes, and risk-free rates are interpolated from zero-coupon bond yields.

The rest of this thesis is organized as follows. We first review the core concepts and literature that form the foundation of American option pricing, situating this work within research. Chapter 1 establishes the theoretical framework for American option valuation, formulates the SDP and provides details on estimation methods to compute the Greeks alongside considerations for computational efficiency. Chapter 2 extends the methodology developed in Chapter 1 to price a panel of options and includes a study on the convergence of each Greek with different methods, providing insights into the stability and accuracy of the model as it is applied to a broader

set of options. Chapter 3 focuses on calibrating the proposed methodology to market data, discussing both the calibration procedures and the resulting accuracy of the fitted models. Finally, we summarize the main contributions, acknowledge the study's limitations, and suggest avenues for further research. The Appendices provide additional computational details and numerical results.

Background and Literature Review

Pricing options with early-exercise features has always been challenging due to the complexity of determining the optimal timing for exercise and accurately forecasting future price movements of the underlying asset.

Over the last decades, various sophisticated techniques have been implemented to value American options. The earliest methods primarily relied on numerical approaches. These techniques include the Finite Difference (FD) method, first introduced by Brennan and Schwartz (1978), which involves solving a discretized partial differential Equation (PDE). In FD, the continuous time and asset price domains are divided into small intervals, to create a mesh where each node represents a specific time and asset price level. The algorithm iteratively computes the option value at each time to maturity on this mesh, starting from the known terminal payoff at maturity and working backward to the present time. While FD can provide accurate solutions for option prices and Greek estimates, it requires careful treatment of stability conditions: explicit schemes are conditionally stable, meaning they require extremely fine time steps. In contrast, the SDP offers a more efficient and stable alternative.

Beyond the role of FD in solving partial differential equations for option pricing, numerical techniques also play a crucial role in estimating Greeks, which measure the sensitivities of option prices to different parameters. The evolution of numerical differentiation techniques in the literature reflects the broader development of calculus and numerical analysis over the past several centuries. As the field of numerical analysis expanded, so too did the sophistication of numerical differentiation methods. By the mid-20^{th} century, these methods had become essential for calculating derivatives that were not analytically tractable (Hull et al., 2009).

Binomial option pricing models, also known as Binomial Trees (BT) were then structured and refined through the seminal work of Cox et al. (1979). BT discretizes the evolution of the asset price into a tree where at each time step the price either moves up or down by fixed factors. It requires a very fine time discretization to achieve accuracy, which leads to an exponential increase in computational cost. This limits its scalability compared to SDP, which operates efficiently without requiring

fine-time discretization and remains computationally stable regardless of the number of decision points.

Although these numerical methodologies offer high precision in pricing American options, they tend to become impractical in high-dimensional scenarios or in situations where the option's payoff is dependent on the historical path of the asset. This obstacle causes the computational complexity to increase exponentially, making these methods very time-consuming and potentially less accurate (Hull et al., 2009).

Other procedures that combined analytical techniques with numerical methods were subsequently suggested but still involved trade-offs between accuracy, complexity, and computational efficiency. Geske and Johnson (1984) derived an analytic solution for American options by treating early exercise as an infinite series of discrete decisions, using polynomial extrapolation to achieve a higher accuracy. This method improves computational efficiency compared to fully numerical methods, but still requires significant computational effort to achieve very high precision. MacMillan (1986) and Barone-Adesi and Whaley (1987) proposed efficient analytic approximations that reduce complexity compared to full numerical methods. These approximations involve simplifying assumptions to create semi-closed-form solutions to price American options, which significantly reduces the computational time needed for valuation. However, these approximations often sacrifice some accuracy because they do not capture the full dynamics of the option's early exercise feature as precisely as full numerical methods.

SDP offers an alternative for low-dimensional option pricing, differing from the FD method as it does not require a fine time discretization. SDP can operate efficiently with a limited number of decision dates, which reduces the computational cost. Originally introduced by Bellman (1957), Stochastic Dynamic Programming is a technique for modeling and solving decision-making problems under uncertainty. It combines elements of stochastic programming and dynamic programming and formulates the problem as a Bellman equation, which recursively determines the optimal policy at each stage. The objective is to determine an optimal decision-making strategy that dynamically adapts to uncertainty, ensuring the best possible outcome at each step.

Pricing American options with SDP dates back to Chen (1970), who managed to directly generate theoretical prices for a limited number of decision dates. However, Chen's work predates the influential work of Black and Scholes (1973) and Merton (1973), thus it does not incorporate risk-neutral pricing. In this thesis, we consider the SDP presented in Ben-Ameur et al. (2002) which introduces the combination of SDP with finite element methods to price American-style Asian options. Finite

elements are employed to discretize the state space and benefit from an increased numerical stability and precision in the pricing process.

Beyond these numerical methods, simulation-based techniques have also been introduced to price American options. These techniques are easy to use and work well because they rely on simple assumptions, where averages of random observations usually lead to accurate results. For instance, the Monte Carlo simulation approach which was applied by Boyle (1977) to European options and later adapted to handle the complexities of options with early exercise features (Longstaff and Schwartz, 2001). Additionally, Tilley (1993) successfully merged a dynamic programming algorithm with a simulation framework, allowing research to develop effective simulation based algorithms for American option pricing.

A widely recognized procedure within the Monte Carlo simulation framework is the Least Squares Monte Carlo (LSM) method developed by Longstaff and Schwartz (2001). This method was groundbreaking and contributed significantly to the field by providing a practical solution for pricing American-style options using Monte Carlo simulation, which was previously very challenging.

The procedure is simple and powerful; it combines Monte Carlo simulation with a least squares regression approach to determine the optimal exercise strategy and thus to estimate the option's value. A key advantage of this method lies in the scalability of regressions: they can be easily adapted to multivariate settings and are suitably adapted to a wide range of option payoff structures.

The Monte Carlo simulation generates random future price paths for the underlying asset based on a stochastic process, such as Geometric Brownian Motion (GBM). At each step, the option's payoff is computed, and least squares regression is used to estimate the continuation value by regressing payoffs against basis functions of the simulated asset prices.

Working backward from maturity, the continuation value is compared to the immediate payoff to determine whether early exercise is optimal. The final option value is obtained by averaging across all simulated paths, yielding an estimate at the initial time.

Prior to this approach, there was no efficient and accurate method for pricing such options using Monte Carlo simulations. However, while the LSM was innovative and effective, it holds certain limitations and drawbacks, such as the approximation bias, which arises because the continuation value is estimated using a finite set of basis functions, which may not fully capture its true underlying structure, leading to systematic errors.

Additionally, the error bias occurs due to the use of Monte Carlo simulations, which are inherently subject to random sampling variability. Since a finite number of simulated paths estimate expected payoffs, different path sets can lead to an

over- or under-estimation of the option value. In practice, increasing the number of simulations can help reduce variability and make the estimate more accurate, but it will never fully eliminate the randomness component unless an infinite number of simulations are used, at the cost of increasing the computational effort required.

Traditional numerical methods, as well as initial simulation-based methods, each have distinct advantages depending on the dimensionality and complexity of the problem.

Numerical techniques are generally preferred for low-dimensional problems, such as those involving one or two underlying assets. This is because they use structured approaches and have deterministic paths as they directly capture the dynamics of early exercise through backward induction without the need for extensive random sampling. However, as the number of dimensions increases, such as when there are multiple stochastic factors affecting the underlying asset, numerical methods become computationally impractical as it leads to a rapid increase in memory requirements and computational time, often referred to as *the curse of dimensionality*.

With recent advances in computational techniques, the curse of dimensionality in option pricing has been alleviated to some extent. Techniques such as hierarchical tensor formats and low-rank tensor representations have proven effective in reducing the computational burden in high-dimensional settings (Bayer et al., 2023). It is worth noting that SDP performs efficiently in low to moderate-dimensional settings compared to other numerical methods such as FD (Ben-Abdellatif et al., 2024), especially when parallel computing is leveraged to accelerate the recursive structure of the algorithm. In higher-dimensional problems, SDP can be combined with Monte Carlo simulations. However, this introduces statistical bias and increases variance. An alternative is to apply dimensionality reduction techniques, such as principal component analysis, to reduce the nominal dimension of the state space to a lower effective dimension. This is often justified in financial applications, where asset prices tend to be highly correlated and exhibit co-movements, allowing for more efficient modeling without significant loss of information.

Among the various numerical and simulation-based methods for pricing American options, SDP stands out as a particularly powerful and efficient approach. Unlike regression-based methods such as LSM, which introduce approximation bias due to finite basis function representation, SDP directly constructs an upper envelope of continuation values, ensuring a more precise estimation of the option price. In addition to its accuracy, SDP remains computationally robust and unconditionally stable regardless of the number of decision points. It also eliminates sampling variability by systematically optimizing exercise decisions through backward induction in a structured and deterministic framework.

Chapter 1

American Option Pricing

This chapter formulates the pricing of American options using the SDP approach and presents the benchmark method used for comparison. We begin by introducing the theoretical model that frames American option pricing as an optimal stopping problem and establish its valuation as an expectation under optimal exercise. We then detail the implementation of the SDP approach and introduce the FD method as a benchmark, which serves as a reference for assessing the accuracy of our pricing and Greeks estimation. The latter sections focus on computing the Greeks using various numerical techniques and evaluate the computational efficiency of the proposed methods.

1.1 American Option Pricing as an Optimal Stopping Problem

We consider $\{S_t : t \in [0, T]\}$ as the price process of an underlying asset, interpreted as the state process defined on the probability space $(\Omega, \mathcal{F}, \mathbb{Q})$, where Ω is the set of all possible asset price paths, $\mathcal{F}$ is the filtration representing the information available up to a given time, and $\mathbb{Q}$ is the risk-neutral measure under which it is assumed that the asset's expected return is equal to the risk-free rate r.

As discussed in Karatzas and Shreve (1998), the *risk-neutral* probability measure $\mathbb{Q}$ ensures that the discounted price process of a tradable asset is a martingale, eliminating arbitrage opportunities. It can be defined under which the price movements of the asset adhere to a specific stochastic differential equation in continuous time as

$$dS_t = (r - q)S_t dt + \sigma S_t dW^{\mathbb{Q}}(t), \quad \text{for} \quad 0 \leq t \leq T. \tag{1.1.1}$$

This equation captures the risk-free rate of return r as well as the volatility of asset returns σ, T the maturity date, and the randomness in price movement represented by the Brownian motion $\{W^{\mathbb{Q}}(t) : t \in [0,T]\}$. Additionally, $dW^{\mathbb{Q}}(t)$ represents the increment of the Brownian motion and dt is an infinitesimal increment.

In our model, we incorporate the dividend yield q, which accounts for the impact of dividend payments on the asset price. The dividend yield represents the continuous rate at which the asset distributes dividends as a fraction of its price. This affects the pricing dynamics by reducing the expected growth rate of the asset under the risk-neutral measure $\mathbb{Q}$, as the asset's total return must compensate for both price appreciation and dividend payments.

The solution for (1.1.1) is given as

$$S_{t+\Delta t} = S_t e^{(r-q-\frac{\sigma^2}{2})\Delta t+\sigma\Delta W_t^{\mathbb{Q}}}, \quad \text{for} \quad 0 \le t \le t+\Delta t \le T, \tag{1.1.2}$$

where the term $\Delta W_t^{\mathbb{Q}} = W^{\mathbb{Q}}(t+\Delta t) - W^{\mathbb{Q}}(t)$ represents the increment of a Brownian motion under the risk-neutral measure $\mathbb{Q}$, satisfying $\Delta W_t^{\mathbb{Q}} \sim \mathcal{N}(0,\Delta t)$. In practice, we rely on the discrete setting to approximate changes over finite intervals of time Δt rather than infinitesimal ones dt.

Equation (1.1.1) describes the evolution of the underlying asset price under the risk-neutral measure. However, pricing an American option requires not only modeling the asset's price dynamics but also determining the optimal exercise time. This naturally leads to an optimal stopping problem, where the holder must decide at each admissible exercise date whether to continue holding the option or to exercise it. We formulate this problem rigorously within the risk-neutral framework.

Let S_t represent the price of the underlying asset at time t, and consider an American put option with a strike price K where the payoff function at exercise is denoted by $h(S_t) = \max(K - S_t, 0)$, which represents the exercise value at time t and remains the same in both continuous and discrete settings. To determine the value $V(S_t,t)$ of the American option at time t, we consider the optimal stopping problem under the *risk-neutral* measure $\mathbb{Q}$. The theoretical value of the American option in continuous time can be expressed as

$$V(S_t,t) = \sup_{\tau\in[t,T]} \mathbb{E}^{\mathbb{Q}}\left[e^{-r(\tau-t)}h(S_\tau) \mid S_t\right], \tag{1.1.3}$$

where τ is a stopping time representing the exercise time, chosen from the interval $[t,T]$, and $\mathbb{E}^{\mathbb{Q}}$ denotes the expectation under the risk-neutral measure $\mathbb{Q}$, which is used to evaluate discounted payoffs in the absence of arbitrage. The risk-free interest rate is r, $e^{-r(\tau-t)}$ is the discount factor applied to the payoff to account for the

time value of money, and $h(S_\tau)$ represents the payoff function of the option at τ.

The value $V(S_t, t)$ represents the expected discounted payoff at the optimal time $\tau \in [t, T]$. In practical applications, this continuous-time problem is approximated using numerical methods by discretizing time into small intervals $[t, t + \Delta t]$. The value function $V(S_t, t)$ is then estimated at each discrete time step. To determine the option's value numerically, it is necessary to compare the continuation value to the exercise value over a sequence of discrete steps. Specifically, the continuation value represents the discounted expected value of holding the option until a potentially more favourable future exercise time, while the exercise value is simply the payoff function $h(S_t)$. At each admissible exercise date, the holder will choose to exercise the option if the immediate exercise value exceeds the continuation value. The Bermudan option value thus satisfies

$$V(S_t, t) = \max\left(h(S_t), \mathbb{E}^{\mathbb{Q}}\left[e^{-r(\Delta t)} V(S_{t+\Delta t}, t + \Delta t) \mid S_t\right]\right). \tag{1.1.4}$$

Equation (1.1.4) provides a discrete-time approximation of the value function $V(S_t, t)$ at any discrete time step t with the corresponding asset price S_t where Δt represents the time step in the numerical scheme. Although the left-hand side does not explicitly depend on Δt, it is understood that $V(S_t, t)$ is computed on a discrete time grid and therefore implicitly depends on the choice of Δt. This formulation is fundamental for numerical methods that approximate the option price through backward recursion.

Equation (1.1.4) implies that the value function $V(S_t, t)$ must satisfy an *optimal exercise boundary* that determines whether it is best to exercise or hold the option at each point in time and asset price. As the time interval $\Delta t \to 0^+$, the discrete approximation of $V(S_t, t)$ converges to the continuous-time solution, meaning that the numerical method more closely approximates the true solution as the time step Δt becomes finer. This convergence occurs because smaller intervals provide a more precise representation of the underlying stochastic process.

The discrete formulation of the American option pricing problem in Equation (1.1.4) sets the foundation for numerical approximation techniques. The next Section introduces the SDP framework and its implementation in the context of American option pricing.

1.2 Stochastic Dynamic Program

This Section will outline the SDP configuration, emphasizing the use of transition matrices, that serve as a foundation for our pricing algorithm. The

partition consists of a discretized version of both time and state space, where the time interval $[0, T]$ is divided into a set of discrete time steps $\{t_0, t_1, \cdots, t_N\}$ and the asset price is discretized into a finite set. Additionally, the transition matrices represent the transition moments that describe the evolution of the asset price within the constructed grid. They are calculated only once as long as the time steps are uniform, and depend on fixed parameters determined by the assumptions in the Black and Scholes (1973) model. SDP divides the valuation problem into two distinct components; the dynamics of the underlying asset and the payoff structure of the derivative being priced. The dynamics of the underlying asset are characterized by the defined transition parameters, which describe how the asset's price evolves.

Pricing American options involves identifying and calculating the early-exercise boundary. This indicates that the boundary itself is part of the solution and must be determined as part of the valuation procedure. SDP provides a powerful approach to address this challenge by breaking down the decision-making process into a sequence of smaller and manageable segments, treating the pricing problem as a series of decisions, where each is based on the present condition of the underlying asset.

Definition 1.2.1. Let $\mathbb{F}$ be a family of real-valued functions defined on a domain $D \subseteq \mathbb{R}^n$. The *upper envelope* of $\mathbb{F}$ is the pointwise supremum

$$\tilde{f}(x) = \sup_{f \in \mathbb{F}} f(x), \quad \forall x \in D.$$

More generally, a function $\tilde{f}$ is called an upper envelope of a function f on D if $\tilde{f}(x) \geq f(x)$ for all $x \in D$, and $\tilde{f}$ belongs to a specific class of functions of interest.

Theorem 1.2.1. *SDP is an upper envelope (Ben-Ameur et al., 2013).*

Let $v_h(s, t_n)$ denote the value function of an American option at time t_n with underlying asset price s. A function $\tilde{v}_h^{\sup}(s, t_n)$ is an upper envelope of $v_h(s, t_n)$ if it satisfies

$$\tilde{v}_h^{\sup}(s, t_n) \geq v_h(s, t_n), \quad \forall (s, t_n) \in [0, T] \times \mathbb{R}^+,$$

where $\tilde{v}_h^{\sup}$ is convex in s.

SDP constructs an upper envelope by backward induction, ensuring that at each discrete time step the approximation satisfies

$$\tilde{v}_h^{\sup}(s, t_n) = \max\left(v_e(s, t_n), \mathbb{E}^{\mathbb{Q}}\left[e^{-r\Delta t}\tilde{v}_h^{\sup}(S_{t_{n+1}}, t_{n+1}) \mid S_{t_n} = s\right]\right),$$

where $v_e(s, t_n)$ is the early-exercise value.

Proof. To demonstrate that the SDP approach constructs an upper envelope for the value function of a derivative contract, we present the proof outlined in Ben-Ameur et al. (2013). Let s represent the price of the underlying asset, $n = 0, 1, ..., N$

denote the discrete time step index, with t_n representing the specific time step where $t_0 = 0, t_N = T$ and $t_n < t_{n+1}$. For simplicity, we assume that the time steps are equally spaced. Suppose that a convex upper envelope approximation $\tilde{v}_h^{\sup}(\cdot, t_{n+1})$ of $v_h(\cdot, t_{n+1})$ that overestimates the value function at a future date is available at a certain future date t_{n+1} for all s such that

$$\tilde{v}_{t_{n+1}}^{\sup}(s) \geq v_h(s, t_{n+1}).$$

This assumption holds at maturity $t_N = T$ since

$$\tilde{v}_h^{\sup}(s, t_N) = v_h(s, t_N) = 0.$$

Starting from $t_N = T$

$$\tilde{v}^{\sup}(s, t_N) = v(s, t_N) = v_e(s, t_N).$$

To maintain the upper-envelope property, the value function at t_{n+1} is defined as

$$\tilde{v}_{t_{n+1}}^{\sup}(s) = \max(v_e(s, t_{n+1}), \tilde{v}_h^{\sup}(s, t_{n+1})).$$

This result holds under the assumption that the payoff function $h(s)$ is convex, which implies that the early-exercise value $v_e(s, t_n)$ is convex. Since the maximum of convex functions is convex, and the expectation of a convex function under a probability measure preserves convexity, the recursive construction maintains convexity at each time step. This ensures that $\tilde{v}_{t_{n+1}}^{\sup}(s)$ is convex and overestimates $v(s, t_{n+1})$ since the maximum of two convex functions is convex. By recursively applying this method through time steps $t_n, t_{n-1}, \ldots, t_0$, the convexity and upper-bound properties are preserved. Therefore, the SDP approach constructs a sequence of upper-envelope approximations that provide an overestimate of the value function at each time step. □

1.2.1 Initializing the Grid

As shown in Figure 1.1, we present a finite grid G that illustrates a discretized state space for the possible prices of the underlying asset over the life of the option. The horizontal axis denotes discrete time increments, $\Delta(t) = t_{n+1} - t_n$ which describes the length of each time step. The discrete time points $t_n \in \mathcal{P}$, $n = 0, 1, ..., N$ correspond to different points in the life of the option up to its maturity T. The vertical axis represents a set of potential prices for the underlying

asset at these times, segmented into a mesh of grid points $a_i \in \mathcal{S}$, $i = 0, ..., p$.

Each node on the grid represents a specific state of the option. Given the holder's flexibility to exercise the option at any of these points, our model must calculate the value of the option at each node in the grid $G = \mathcal{S} \times \mathcal{P}$ to determine the most advantageous exercise policy. In this study, the option holder is permitted to exercise only at certain predetermined points on the grid, since we are working with Bermudan options, which means that the optimal exercise strategy is limited to specific dates in $\mathcal{P}$. However, as the number of predetermined exercise points increases, the Bermudan option increasingly mimics the American option. This is precisely what we do in our approach, which allows us to approximate the American option by refining the exercise grid. Notably, unlike the FD algorithm, the SDP approach does not require additional discretization between these exercise dates.

The structured grid G of Figure 1.1 is particularly efficient for pricing a panel of options because once it is established, it can be used to price any option with the same payoff structure, strike price, asset price distribution and exercise dates, whose parameters fall within the maturity range covered by the grid. This is because it is strategically sized to account for the longest maturity date $T_{\max}$ in the options panel so that each option, irrespective of its individual maturity, is encompassed within the grid, eliminating the need for multiple grid configurations. If an option's initial price falls between two grid points, interpolation can be used to estimate its exact value. By knowing the values at the surrounding grid points, we can interpolate to find the value at any point within the grid. By accommodating the longest maturity, we will need to compute the transition probabilities calculation only once, and these can be reused for all other options with shorter maturities, saving significant computational time.

Initializing the grid G involves dividing the positive real axis into a set of intervals and subsequently estimating the option value using a piecewise linear interpolation method. For that, let $\mathcal{S} = \{a_0, \ldots, a_p\}$ be a set of grid points that covers the state space into $p+1$ intervals as $(a_i, a_{i+1}]$ for $i = 0, ..., p$. We choose to construct the grid points as quantiles of the underlying asset at $t_N = T$. We consider the asset price distribution at time t_N, given the initial asset price at time t_0 is s_0, as selected in Ben-Ameur et al. (2013), under the Black and Scholes (1973) model where $a_0 = 0 < a_1 < \cdots < a_p < a_{p+1} = +\infty$, as

$$
\begin{aligned}
a_1 &= s_0 e^{\left(-\frac{r-\sigma^2}{2}\right)T - 7\sigma\sqrt{T}}, \\
a_2 &= s_0 e^{\left(-\frac{r-\sigma^2}{2}\right)T - 5\sigma\sqrt{T}}, \\
a_{p-1} &= s_0 e^{\left(-\frac{r-\sigma^2}{2}\right)T + 5\sigma\sqrt{T}},
\end{aligned}
$$

$$a_p = s_0 e^{\left(-\frac{r-\sigma^2}{2}\right)T+7\sigma\sqrt{T}},$$

$$a_i = s_0 e^{\left(-\frac{r-\sigma^2}{2}\right)T+\sigma\sqrt{T}z_i}, \tag{1.2.1}$$

where z_i is the quantile of the standard normal distribution associated with fraction $\frac{i}{p}$, for $i = 3, \ldots, p-2$. This method features the characteristic of placing more evaluation nodes in areas most frequently visited, as detailed in Equation (1.2.1). Additionally, the grid points remain constant over time.

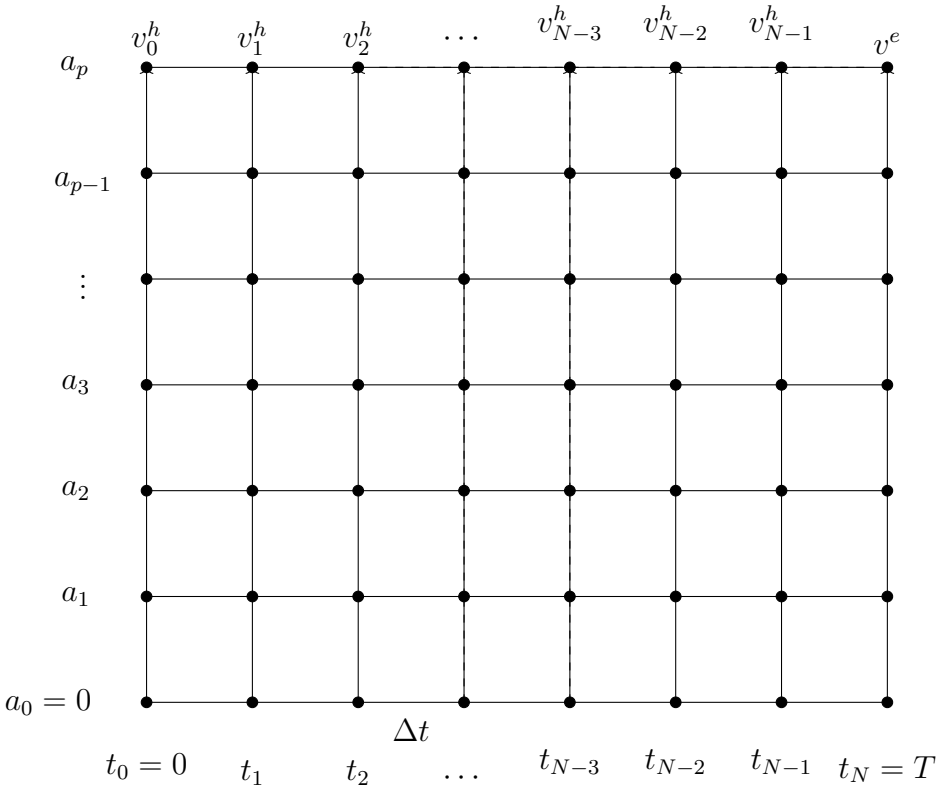

Figure 1.1: Finite grid G illustration.

1.2.2 Transition Matrices

In our algorithm, transition probabilities are used to model the dynamics of the underlying asset's price movements. Considering that Z is a standard normal random variable independent of the past of the underlying asset prices s_t and the cumulative density function Φ of Z

$$A_{ki} = \mathbb{E}^{\mathbb{Q}}\left[\mathbb{I}\left(\frac{a_i}{a_k} < e^{\mu\Delta t+\sigma\sqrt{\Delta t}Z} \leq \frac{a_{i+1}}{a_k}\right)\right]$$
$$= \begin{cases} \Phi(c_{k,1}) & \text{for } i = 0, \\ \Phi(c_{k,i+1}) - \Phi(c_{k,i}) & \text{for } 1 \leq i \leq p-1, \\ 1 - \Phi(c_{k,p}) & \text{for } i = p. \end{cases}$$

$$B_{ki} = \mathbb{E}^{\mathbb{Q}}\left[a_k e^{\mu\Delta t+\sigma\sqrt{\Delta t}Z}\mathbb{I}\left(\frac{a_i}{a_k} < e^{\mu\Delta t+\sigma\sqrt{\Delta t}Z} \leq \frac{a_{i+1}}{a_k}\right)\right]$$
$$= \begin{cases} a_k\Phi(c_{k,1} - \sigma\sqrt{\Delta t})e^{-(r-q)\Delta t} & \text{for } i = 0, \\ a_k[\Phi(c_{k,i+1} - \sigma\sqrt{\Delta t}) - \Phi(c_{k,i} - \sigma\sqrt{\Delta t})]e^{-(r-q)\Delta t} & \text{for } 1 \leq i \leq p-1, \\ a_k[1 - \Phi(c_{k,p} - \sigma\sqrt{\Delta t})]e^{-(r-q)\Delta t} & \text{for } i = p, \end{cases}$$

where $\mu = r - q - \frac{\sigma^2}{2}$ and $c_{k,i} = [\log(a_i/a_k) - \mu\Delta t]/(\sigma\sqrt{\Delta t})$ for $k = 1, .., p$ and $i = 0, ..., p$ represents the normalized quantiles that are used to calculate transition probabilities in matrices A_{ki} and $B_{k,i}$. Matrix A_{ki} illustrates the probability that the next asset price transitions from the current state a_k to the next one a_{k+1} in the following time step a_{i+1}. Matrix B_{ki} considers not only the probability of transitioning between the price levels but also weights this probability by the future price ($a_k e^{\mu\Delta t+\sigma\sqrt{\Delta t}Z}$). This weight accounts for the value of the asset at the next step a_{i+1}, conditional on falling within the required interval.

1.2.3 Approximating the Option Value via Piecewise Linear Interpolation

To find the option value, we utilize a piecewise linear interpolation based on the initial asset price $s \in \mathcal{S}$ and the local coefficients α_i^m and β_i^m corresponding to the optimal exercise strategy at node i and time m. Given an approximation $\tilde{v}_m$ of the option value V_m at the discrete points a_i and steps $m = \{0, \ldots, n-1\}$, this function is interpolated

$$\tilde{v}_m(s) = \sum_{i=0}^{p} (\alpha_i^m + \beta_i^m s)\, \mathbb{I}(a_i < s \leq a_{i+1}), \quad t_m \in \{t_0, \ldots, t_{n-1}\} \tag{1.2.2}$$

where $\mathbb{I}$ is the indicator function and the local coefficients of this interpolation at step m, α_i^m and β_i^m, are obtained by solving the linear equations

$$\alpha_i^m = \frac{\tilde{v}_m(a_i)a_{i+1} - \tilde{v}_m(a_{i+1})a_i}{a_{i+1} - a_i}, \tag{1.2.3}$$

$$\beta_i^m = \frac{\tilde{v}_m(a_{i+1}) - \tilde{v}_m(a_i)}{a_{i+1} - a_i}, \quad i \in \{0, \ldots, p+1\}. \tag{1.2.4}$$

For $i = p$ we take $\alpha_p^m = \alpha_{p-1}^m$ and $\beta_p^m = \beta_{p-1}^m$.

The exercise value of an American option at any decision date $t = \{t_1, \ldots, t_n\}$, as stated in the contract is given by

$$v^e(s) = \begin{cases} (s-K)^+ & \text{for the American call option,} \\ (K-s)^+ & \text{for the American put option,} \end{cases} \tag{1.2.5}$$

where $(x)^+ \equiv \max\{0, x\}$. Note that $v^e(s)$ is equal to the payoff function $h(s)$. However, we use the distinct notation v^e to emphasize its role in the early-exercise decision rule, where it is compared against the continuation value.

Accounting for the risk-neutral principle, the continuation value represents the expected value of keeping the option until the next time step rather than exercising it immediately and can be mathematically written at any $t_m \in \{0, \ldots, n-1\}$ as

$$\begin{aligned} \tilde{v}_m^h(s) &= \mathbb{E}^{\mathbb{Q}}[\rho \tilde{v}_{m+1}(a_i+1)(S_{m+1}) | S_m = s] \\ &= \rho \sum_{i=0}^{p} \alpha_i^{m+1} \mathbb{E}^{\mathbb{Q}} \left[\mathbb{I}(a_i < S_{m+1} \leq a_{i+1}) \mid S_m = a_k\right] \\ &+ \beta_i^{m+1} \mathbb{E}^{\mathbb{Q}} \left[S_{m+1} \mathbb{I}(a_i < S_{m+1} \leq a_{i+1}) \mid S_m = a_k\right], \end{aligned} \tag{1.2.6}$$

where $\rho = e^{-\Delta t(r-q)}$ is the discount factor and $\tilde{v}_m^h(s)$ denotes the approximate holding value of the option. For $k = 1, \ldots, p$ and $i = 0, \ldots, p$, we have

$$\begin{aligned} A_{ki} &= \mathbb{E}^{\mathbb{Q}} \left[\mathbb{I}(a_i < S_{m+1} \leq a_{i+1}) \mid S_m = a_k\right], \\ B_{ki} &= \mathbb{E}^{\mathbb{Q}} \left[S_{m+1} \mathbb{I}(a_i < S_{m+1} \leq a_{i+1}) \mid S_m = a_k\right]. \end{aligned} \tag{1.2.7}$$

Combining Equations (1.2.6) and (1.2.7), we obtain

$$\tilde{v}_m^h(s) = \rho \sum_{i=0}^{p} \left(\alpha_i^{m+1} A_{ki} + \beta_i^{m+1} B_{ki}\right), \quad \text{for } m = 0, \ldots, n-1. \tag{1.2.8}$$

Finding the optimal exercise then involves comparing the exercise value $v^e(s)$ with the holding value $\tilde{v}_m^h(s)$ at each exercise step, which marks the choice between conditions where the early exercise of the option is advantageous and where holding onto it is preferred. We can express the optimal exercise process $\bar{v}_m(s)$ as follows

$$\bar{v}_m(s) = \begin{cases} \tilde{v}_0^h(s) & \text{for } m = 0, \\ \max(v^e(s), \tilde{v}^h(s)) & \text{for } m = 1, \ldots, n-1, \\ v^e(s) & \text{for } m = n. \end{cases} \tag{1.2.9}$$

Equation (1.2.9) reflects the decision-making process for the option holder: the option will be maintained as long as its holding value exceeds the exercise value. Depending on the exercise value $v^e(s)$, the holder will exercise the option if it yields a positive value, or abandon the contract if the exercise value is zero. The approach to determine the value of the option involves using backward recursion through Equations (1.2.5)-(1.2.9), starting from the established function, or boundary condition $\tilde{v}_n(s) = v^e(s)$. For $m \in \{0, \dots, n-1\}$, the value function $\tilde{v}_m(s)$ is not analytically explicit and requires an approximation over the finite grid defined in Section (1.2.1).

1.2.4 Summary of the SDP Algorithm

The stochastic dynamic program for option pricing involves a backward recursion process that starts at the option's maturity T and iteratively works backward. The SDP approach works as follows

1. *Compute the transition matrices A_{ki} and B_{ki}:* These transition parameters are computed only once and reused for the entire partition G.

2. *Calculate the local coefficients at t_n:* At maturity T, the option's value is known and corresponds to its intrinsic value $v^e(s)$ from Equation (1.2.9). At this stage, we can compute α_i^m and β_i^m at $m = T$ with Equations (1.2.3) and (1.2.4) respectively.

3. *Compute the continuation value at t_{n-1}:* Use the computed α_i^m and β_i^m and the transition matrices A_{ki} and B_{ki} to calculate $\tilde{v}_m^h(s)$ at $t = n-1$ from Equation (1.2.8).

4. *Maximize to find the approximated option value:* Maximize between $\tilde{v}_m^h(s)$ and $v^e(s)$ as in (1.2.9).

Steps two to four are repeated recursively for each time step until iterating back to the initial time step t_0, and then apply Equation (1.2.2) to interpolate for the current option value $\tilde{v}_m$ using the corresponding local coefficients α_i^m and β_i^m of the closest grid points a_i to the current asset price S_0.

Through this SDP algorithm, we obtain not only an approximation of the option's current value but also an approximation of the Delta Δ, which is the sensitivity of the option's price to a small change in the price of the underlying asset $\frac{\partial V}{\partial S}$. A more detailed discussion on this topic, as well as the extraction of other Greeks, will be presented in Section (1.4). Gamma, Theta and Vega can also be easily extracted from the grid $G = \mathcal{S} \times \mathcal{P}$ with finite-difference approximations.

1.3 Finite-Difference Method as a Benchmark

In this Section, we briefly present the FD method that will serve as a benchmark to verify our option price and Greeks estimates. This method allows for the resolution of partial differential Equations (PDEs) satisfied by derivatives when they do not offer an explicit solution. The FD method was introduced to finance by Brennan and Schwartz (1978) and has since been suited for valuing derivative products with anticipated exercise features, as further demonstrated by the general application of the explicit FD approach by Hull and White (1990).

The chosen benchmark is computed with the implicit FD method derived from Wilmott et al. (1995) and Hull et al. (2009) which uses the Crank-Nicolson method and aims to ameliorate the performance of the basic explicit FD in terms of precision and stability.

The FD approach discretizes the Black and Scholes (1973) PDE over a structured grid representing different time and asset price levels. Using backward induction, it iteratively determines option values and incorporates early exercise conditions for American options. The solution is refined using the Successive Over-Relaxation (SOR) method to accelerate convergence. This approach enables a direct estimation of option Greeks through finite difference approximations.

For consistency, we compute the FD benchmark with a grid size of $(N_s, N_t) = (4000, 4000)$, ensuring numerical stability since values remain consistent, with variations appearing only at the sixth decimal place. The detailed formulation, discretization steps, and solution procedure are presented in Appendix B.

1.4 Estimation of Greeks for American Options

The Greeks are fundamental tools in risk management and option pricing, they provide a mathematical way to measure the sensitivity of an option's price to various factors and allow the strategic structuring of portfolios.

Each Greek quantifies the sensitivity of an option's value to a specific parameter. We will estimate Delta, Gamma, Theta, and Vega. Although our analysis primarily relies on the Black and Scholes (1973) model, which assumes constant volatility of returns, we estimate Vega with small variations in volatility to assess the sensitivity of the option's value to these changes. First, we will introduce numerical differentiation methods, followed by Richardson Extrapolation, as both will be used for estimation purposes. We then provide a detailed presentation focused on the efficient computation of the Greeks.

1.4.1 Numerical Differentiation for Greeks Estimation

All numerical differentiation methods approximate Greeks by utilizing values in the SDP grid G of Section (1.2.1). Using well-established mathematical results from Abramowitz and Stegun (1948), we present the following versions for each approximation.

For the first-order partial derivative, the central approximation can be expressed as

$$\frac{\partial f(x_i)}{\partial x} = \frac{f(x_{i+1}) - f(x_{i-1})}{2\overline{h}} + O(\overline{h}^2) \quad \text{for a small selected } \overline{h}, \tag{1.4.1}$$

where f is a given function and $\overline{h}$ is a small step size approaching zero as a limit. This approximation has an order of convergence of $O(\overline{h}^2)$. The error term indicates that the error decreases quadratically as the step size h tends to zero. This result is derived from the Taylor series expansion, where the first-order error terms cancel out, resulting in a higher-order error of $O(\overline{h}^2)$.

For the second-order partial derivative, the central approximation is presented as

$$\frac{\partial^2 f(x_i)}{\partial x^2} = \frac{f(x_{i+1}) - 2f(x_i) + f(x_{i-1})}{\overline{h}^2} + O(\overline{h}^2) \quad \text{for a small selected } \overline{h}, \tag{1.4.2}$$

or as

$$\frac{\partial^2 f(x_i)}{\partial x^2} = \frac{-f(x_{i+2}) + 16f(x_{i+1}) - 30f(x_i) + 16f(x_{i-1}) - f(x_{i-2})}{12\overline{h}^2} + O(\overline{h}^4), \tag{1.4.3}$$

for a small selected $\overline{h}$ approaching zero as a limit and f is a given function. Equation (1.4.2) has an order of convergence of $O(\overline{h}^2)$, whereas Equation (1.4.3) has an order of convergence of $O(\overline{h}^4)$. This means that the approximation of (1.4.3) has an error term that decreases faster compared to the approximation of (1.4.2). This is because Equation (1.4.3) allows more terms in the Taylor series expansion to cancel out since it uses more points.

We also introduce the forward approximation which uses the function values at discrete points that follow the point of interest x_i at which we want to approximate the derivative of the function

$$\frac{\partial f(x_i)}{\partial x} = \frac{f(x_{i+1}) - f(x_i)}{\overline{h}} + O(\overline{h}) \quad \text{for a small selected } \overline{h}, \tag{1.4.4}$$

where f is a given function and $\overline{h}$ is a small step size approaching zero as a limit. The error term for the forward difference approximation is $O(\overline{h})$ because the method

uses only the current point x_i and the next point x_{i+1} to approximate the derivative. When applying the Taylor series expansion to the function $f(x)$ around the point x_i, the leading truncation error term is proportional to the first power of the step size $\overline{h}$. This means that as $\overline{h}$ becomes smaller, the error decreases linearly with respect to $\overline{h}$, leading to a first-order convergence of $O(\overline{h})$.

1.4.2 Richardson Extrapolation

Richardson (1911) acknowledged the power of the FD method as a general technique for solving differential equations numerically and noted that by systematically refining a mesh, it was possible to obtain more accurate results and that the errors in the numerical solutions would tend to decrease predictably. He also examined the effects of different mesh sizes on the accuracy of the results, an analysis that led to his ideas about extrapolating from multiple approximations to improve accuracy.

The essence of his findings lay in improving the rate of convergence of a sequence of estimates and increasing the order of accuracy of an approximation by using a combination of approximations at different step sizes or, in other words, combining two derivative estimates to compute a third one. The idea is to start with an approximation using a step size $\overline{h}$ that has a known error term expressed as a series expansion in terms of $\overline{h}$. Richardson extrapolation takes advantage of this by calculating approximations with different step sizes, for example, $\overline{h}$ and $\overline{h}/2$, and combining them with a weighted combination of both to eliminate the leading error term.

We consider a numerical approximation $R_0(\overline{h})$ to a quantity R^*. The true value R^* and the approximation $R_0(\overline{h})$ are related by an error term that can be expanded as a power series in $\overline{h}$

$$R^* = R_0(\overline{h}) + C_0\overline{h}^{k_0} + C_1\overline{h}^{k_1} + C_2\overline{h}^{k_2} + \sum_{n=3}^{\infty} C_n\overline{h}^{k_n},$$

where C_i are unknown constants and k_i are known constants such that $\overline{h}^{k_i} > \overline{h}^{k_{i+1}}$, and it can be rewritten as

$$R^* = R_i(\overline{h}) + O(\overline{h}^{k_i}), \tag{1.4.5}$$

where the last term $O(\overline{h}^{k_i})$ represents the leading-order term of the truncation error of the $R_i(\overline{h})$ approximation. This means that the error in the approximation $R_i(\overline{h})$ is bounded above by a constant times $\overline{h}^{k_i}$ as $\overline{h} \to 0$. Formally, there exists a constant $M > 0$ such that $\left|R^* - R_i(\overline{h})\right| \leq M\overline{h}^{k_i}$.

Similarly, in Equation (1.4.5) the approximation $R_i(\overline{h})$ is said to be an $O(\overline{h}^{k_i})$ approximation. Note that by simplifying big O notation, the following formulae are equivalent

$$R^* = R_0(\overline{h}) + C_0\overline{h}^{k_0} + C_1\overline{h}^{k_1} + C_2\overline{h}^{k_2} + \sum_{n=3}^{\infty} C_n\overline{h}^{k_n},$$
$$R^* = R_0(\overline{h}) + C_0\overline{h}^{k_0} + O(\overline{h}^{k_1}),$$
$$R^* = R_0(\overline{h}) + O(\overline{h}^{k_0}).$$

Richardson extrapolation is a process that finds a better approximation of R^* by changing the error formula from $R^* = R_0(\overline{h}) + O(\overline{h}^{k_0})$ to $R^* = R_1(\overline{h}) + O(\overline{h}^{k_1})$. Therefore, by replacing $R_0(\overline{h})$ with $R_1(\overline{h})$ the truncation error has reduced from $O(\overline{h}^{k_0})$ to $O(\overline{h}^{k_1})$ for the same step size $\overline{h}$. The general pattern occurs in which $R_i(\overline{h})$ is a more accurate estimate than $R_j(\overline{h})$ when $i > j$. By this process, we have achieved a better approximation of R^* by subtracting the largest term in the error which was $O(\overline{h}^{k_0})$. This process can be repeated to remove more error terms to get even better approximations.

We now demonstrate how we can eliminate the leading error term $O(\overline{h}^{k_i})$ (Abramowitz and Stegun, 1948). We use two approximations with different step sizes $\overline{h}$ and $\frac{\overline{h}}{t}$

$$R^* = R_0(\overline{h}) + C_0\overline{h}^{k_0} + C_1\overline{h}^{k_1} + C_2\overline{h}^{k_2} + O(\overline{h}^{k_3}), \tag{1.4.6}$$
$$R^* = R_0\left(\frac{\overline{h}}{t}\right) + C_0\left(\frac{\overline{h}}{t}\right)^{k_0} + C_1\left(\frac{\overline{h}}{t}\right)^{k_1} + C_2\left(\frac{\overline{h}}{t}\right)^{k_2} + O\left(\frac{\overline{h}}{t}\right)^{k_3}. \tag{1.4.7}$$

To improve our approximation from $O(h^{k_0})$ to $O(h^{k_1})$, we multiply Equation (1.4.7) by a factor t^{k_0} and subtract Equation (1.4.6) from it to reduce the first error term

$$(t^{k_0}-1)R^* = t^{k_0}R_0\left(\frac{\overline{h}}{t}\right) - R_0(\overline{h}) + \left(t^{k_0}C_1\left(\frac{\overline{h}}{t}\right)^{k_1} - C_1\overline{h}^{k_1}\right) + \left(t^{k_0}C_2\left(\frac{\overline{h}}{t}\right)^{k_2} - C_2\overline{h}^{k_2}\right) + O(h^{k_3}).$$

This multiplication and subtraction was performed because $t^{k_0}R_0(\frac{\overline{h}}{t}) - R_0(\overline{h})$ is an $O(\overline{h}^{k_1})$ approximation of $(t^{k_0}-1)R^*$.

The term $C_0\left(\frac{\overline{h}}{t}\right)^{k_0} t^{k_0} - C_0\overline{h}^{k_0}$ does not appear in the final equation because it simplifies to zero. We solve for R^* and get

$$R^* = \frac{\left[t^{k_0}R_0\left(\frac{\overline{h}}{t}\right) - R_0(\overline{h})\right] + \left(t^{k_0}C_1\left(\frac{\overline{h}}{t}\right)^{k_1} - C_1\overline{h}^{k_1}\right)}{t^{k_0}-1} + \frac{\left(t^{k_0}C_2\left(\frac{\overline{h}}{t}\right)^{k_2} - C_2h^{k_2}\right)}{t^{k_0}-1} + O(\overline{h}^{k_3}),$$

and we can rewrite it as

$$R^* = R_1(\overline{h}) + O(h^{k_1}),$$

where :

$$R_1(\overline{h}) = R_{\text{Rich}} = \frac{t^{k_0} R_0(\frac{\overline{h}}{t}) - R_0(\overline{h})}{t^{k_0} - 1} \tag{1.4.8}$$

is the estimate of the true value R^* and k_0 is the order of the leading error term.

1.4.3 Delta Estimation

Denoted by Δ, it is the first derivative of the option's price V, with respect to the price of the underlying asset S and is defined as $\Delta = \frac{\partial V}{\partial S}$. Delta is often referred to as the *hedge ratio* because it provides the proportion of shares needed to hedge the option's price movement, which allows it to set up a delta-neutral position.

In the SDP algorithm, since the approximated value of the option $\tilde{v}$ is known for all periods, the estimated Delta is simply the beta coefficient β since it is the first derivative of Equation (1.2.2) with respect to the stock price S. We refer to this as Δ_{SDP}.

The approximation of the Delta based on the first-order central finite-difference approximation of Equation (1.4.1) is

$$\Delta_{\text{CD}} = \frac{\tilde{v}(S_0 + \delta S, t_0) - \tilde{v}(S_0 - \delta S, t_0)}{2\delta S}, \quad \text{for a selected small } \delta S. \tag{1.4.9}$$

This method is generally recognized for its computational efficiency, accuracy across all S values, and ease of implementation, which makes it widely adopted for estimating Delta (Hull et al., 2009). To confirm this, we will benchmark it against other methods to identify the most precise and robust approach.

Another innovative method would be to refine the calculation of Delta within the SDP framework by integrating the current market price S_0 and its surrounding points directly into the computational grid. This inclusion ensures that the points are appropriately considered within the continuum of possible stock prices over time. The method proceeds by calculating the Delta, for points immediately adjacent to S_0 at the left and the right $\Delta(S_0 - \delta S, t_0)$ and $\Delta(S_0 + \delta S, t_0)$ respectively, using the first-order central finite-difference approximation of Equation (1.4.9) and then averaging these two Deltas to interpolate the option's delta Δ_{mean} precisely at S_0, as shown in Figure 1.2.

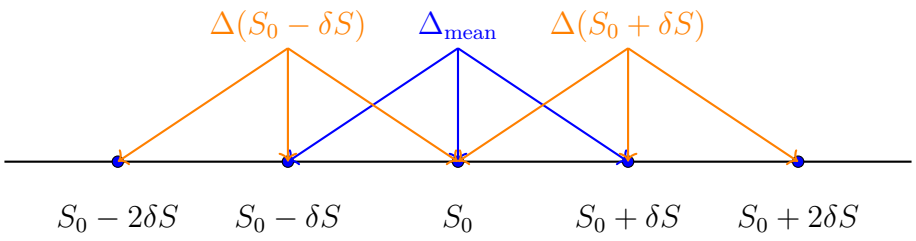

Figure 1.2: Construction of Δ_{mean}.

We include the four surrounding points of S_0 into the grid G.

We can also approximate Delta using Richardson extrapolation of Section (1.4.2) by computing it for two different step sizes δS and $\delta S/2$, in the central difference approach and then use these two approximations to extrapolate for a higher accuracy result

$$\Delta_1 = \frac{\tilde{v}(S_0 + \delta S, t_0) - \tilde{v}(S_0 - \delta S, t_0)}{2\delta S}, \tag{1.4.10}$$

$$\Delta_2 = \frac{\tilde{v}(S_0 + \delta S/2, t_0) - \tilde{v}(S_0 - \delta S/2, t_0)}{\delta S}, \tag{1.4.11}$$

$$\Delta_{\text{Rich}} = 2\Delta_2 - \Delta_1, \tag{1.4.12}$$

where Δ_1 and Δ_2 correspond to computed Deltas with respective δS and $\delta S/2$ step sizes.

We now evaluate the robustness of our approximations by comparing them to the FD benchmark. Specifically, we study the numerical accuracy of four Delta approximations: Δ_{SDP} derived from direct SDP computation as the first derivative of Equation (1.2.2), Δ_{CD} obtained through first-order central finite difference of Equation (1.4.9), Δ_{mean} integrating S_0 within the grid as presented in Figure 1.2, and Δ_{Rich} using Richardson extrapolation of Equation (1.4.10). These results are benchmarked against the FD [2] numerical approach discussed in Section (1.3). Throughout this Section and the upcoming ones, we consider a put option with 50 exercise opportunities in each year $N_{\text{ex}} = 50$ with a total time step of $N_t = 4000$, $T = 3$ years maturity, volatility of $\sigma = 10\%$, and a risk-free interest rate of $r = 5\%$. The initial stock price is fixed at $S_0 = 100$, and the strike price $K = 100$. We also set $\delta S = 0.01 S_0$, so we consider a 1% change in S_0.

Figure 1.3 illustrates the Delta Δ estimates of the studied numerical methods compared to the FD benchmark in the region where $90 \leq S_0 \leq 110$, to provide insight

[2] The code to price a single option using the FD method was obtained from Zheng (2021). Modifications were made to extend the implementation for pricing a panel of options by applying the homogeneity property. I sincerely thank `Wencai Zheng` for making the original implementation available.

into the behavior of Delta estimates in the in-the-money (ITM), at-the-money (ATM), and out-of-the-money (OTM) regions. The goal is to highlight the stability and convergence of each method relative to the benchmark, to allow for a clear distinction in performance as the partition size increases.

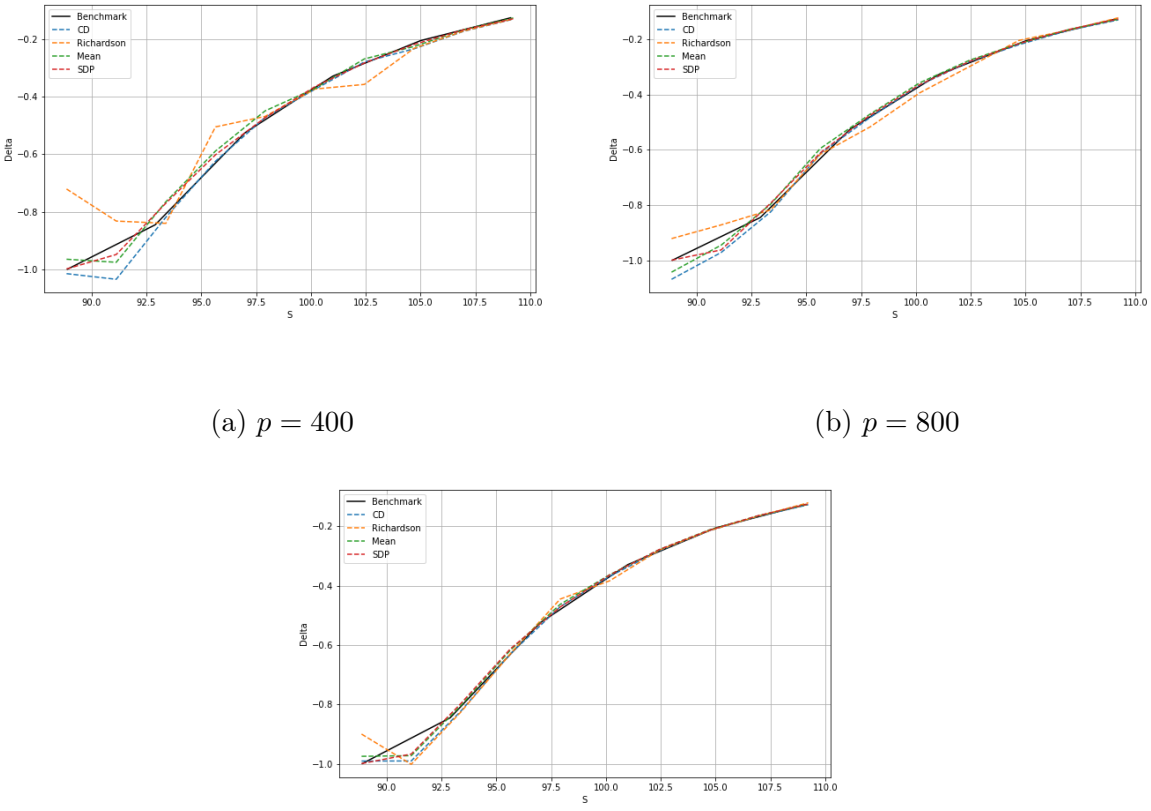

(a) $p = 400$

(b) $p = 800$

(c) $p = 1200$

Figure 1.3: Delta for $S_0 \in [90, 110]$ with different partition sizes p.

The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$. The option is an American put option. Each panel represents the convergence behavior with increasing partition sizes p.

Upon analyzing the panels of Figure 1.3, it becomes evident that certain methods demonstrate improved stability. At a low partition size of $p = 400$, we observe that the Richardson method exhibits some instability relative to the benchmark. The SDP method provides the most precise estimation at low partition sizes. However, at an increased partition size of $p = 1200$, all methods show reduced instability. They are

accurate within 1% of the benchmark value. The SDP method continues to offer the most stability compared to the benchmark with an error of 0.7%, followed by the mean approximation with 0.73%.

Table A.1 in Appendix A, provides a quantitative comparison of the Delta estimates at partition sizes $p = 400$ and $p = 1200$ relative to the benchmark. From these tables, it is apparent that the Δ_{SDP} estimates consistently align closely with the benchmark values $\Delta_{\text{Benchmark}}$.

1.4.4 Gamma Estimation

Denoted by Γ, it measures the convexity of an option's value in relation to the underlying asset's price. Gamma is defined as the second partial derivative of the option value with respect to the underlying asset price $\Gamma = \frac{\partial^2 V}{\partial S^2}$.

To estimate the Gamma of an option, we first approximate by a second-order central finite-difference of Equation (1.4.2)

$$\Gamma = \frac{\tilde{v}(S_0 + \delta S, t_0) - 2\tilde{v}(S_0, t_0) + \tilde{v}(S_0 - \delta S, t_0)}{(\delta S)^2}, \tag{1.4.13}$$

where $\tilde{v}(S_0 + \delta S, t_0)$ and $\tilde{v}(S_0 - \delta S, t_0)$ are the approximate option values at prices $S_0 + \delta S$ and $S_0 - \delta S$ respectively, and δS is a small change in the price.

For more accuracy, we test higher-order central differences since they can provide more precise approximations of derivatives by incorporating information from a larger number of points around the point of interest as explained in Section (1.4.1). From Equation (1.4.3) we get

$$\tilde{\Gamma} = \frac{-\tilde{v}(S_0 + 2\delta S, t_0) + 16\tilde{v}(S_0 + \delta S, t_0) - 30\tilde{v}(S_0, t_0) + 16\tilde{v}(S_0 - \delta S, t_0) - \tilde{v}(S_0 - 2\delta S, t_0)}{12(\delta S)^2}. \tag{1.4.14}$$

We also apply Richardson extrapolation, where we compute Γ_1 and Γ_2 corresponding to step sizes of δS and $\delta S/2$, respectively, as follows

$$\Gamma_1 = \frac{\tilde{v}(S_0 + \delta S, t_0) - 2\tilde{v}(S_0, t_0) + \tilde{v}(S_0 - \delta S, t_0)}{(\delta S)^2}, \tag{1.4.15}$$

$$\Gamma_2 = \frac{\tilde{v}(S_0 + \delta S/2, t_0) - 2\tilde{v}(S_0, t_0) + \tilde{v}(S_0 - \delta S/2, t_0)}{(\delta S/2)^2},$$

$$\Gamma_{\text{Rich}} = (4\Gamma_2 - \Gamma_1)/3.$$

To compare the accuracy of our three approximation methods, we again benchmark them with the same FD solution of Section (1.3). We compare the Gamma estimates

to the benchmark, and examine their performance as partition sizes p increase for different values of S_0 in the grid G.

Figure 1.4 illustrates the Gamma Γ estimates of the studied numerical methods compared to the FD benchmark in the region where $90 \leq S_0 \leq 110$.

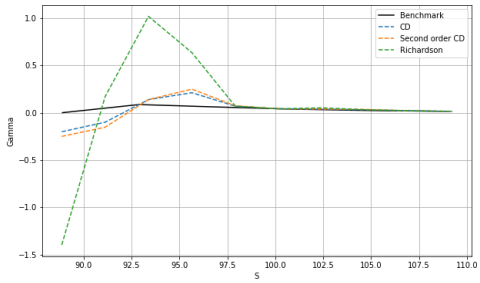

(a) $p = 400$

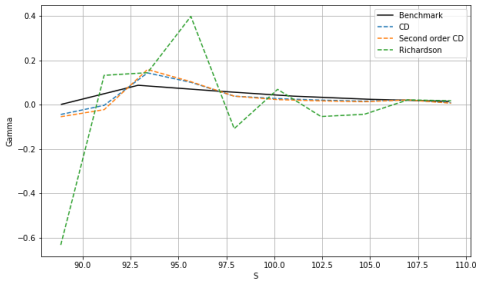

(b) $p = 800$

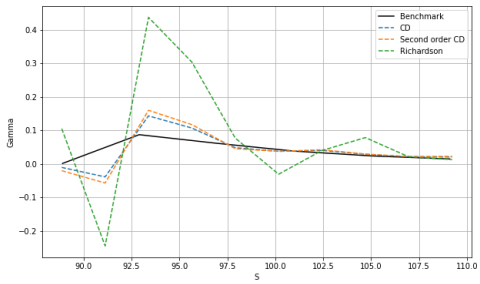

(c) $p = 1200$

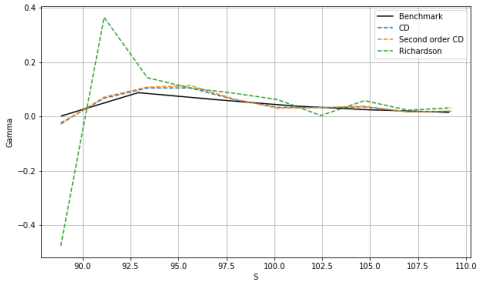

(d) $p = 2000$

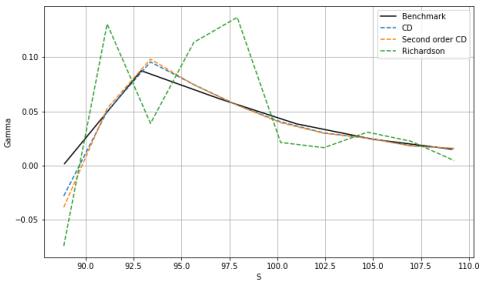

(e) $p = 3000$

Figure 1.4: Gamma for $S_0 \in [90, 110]$ with different partition sizes p.

The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$. The option is an American put option. Each panel

represents the convergence behavior with increasing partition sizes p.

Following the analysis of Figure 1.4 at increasing partition sizes, a clear pattern of instability emerges across all the evaluated methods, particularly in areas where $S < 100$. All methods exhibit a level of instability, with different degrees. Specifically, the Richardson method shows the most significant instability across all partition sizes p with an accuracy of 12% compared to the benchmark. The CD method closely approximates the benchmark at $p = 2000$ and $p = 3000$ with an accuracy within 0.3% the benchmark.

Table A.2 in Appendix A report the Gamma estimates for increased partition sizes p alongside the benchmark values $\Gamma_{\text{Benchmark}}$.

1.4.5 Theta Estimation

Denoted by Θ, it measures the sensitivity of the option's price to a change in the time to expiration $\Theta = \frac{\partial V}{\partial t}$. It represents the rate at which an option's price will change with time, holding everything else constant. Theta is often referred to as the *time decay* of an option because options lose value as they approach expiration, assuming all other factors remain unchanged.

We use the forward difference approximation for a small change in time, where the time increment is $\Delta t = \frac{T}{N_t}$, with N_t representing the total number of time steps over the option's life. We apply Equation (1.4.4) as

$$\Theta_{\text{FD}} = \frac{\tilde{v}(S_0, t_0 + \Delta t) - \tilde{v}(S_0, t_0)}{\Delta t}. \tag{1.4.16}$$

We can also apply the central difference approximation of Equation (1.4.1) for Theta Θ

$$\Theta_{\text{CD}} = \frac{\tilde{v}(S_0, t_0 + \Delta t) - \tilde{v}(S_0, t_0 - \Delta t)}{2\Delta t}, \tag{1.4.17}$$

and the Richardson extrapolation where we compute Θ_1 and Θ_2 corresponding to step sizes of Δt and $\Delta t/2$, respectively

$$\Theta_1 = \frac{\tilde{v}(S_0, t_0 + \Delta t) - \tilde{v}(S_0, t_0)}{\Delta t},$$

$$\Theta_2 = \frac{\tilde{v}(S_0, t_0 - \Delta t/2) - \tilde{v}(S_0, t_0)}{\Delta t/2},$$

$$\Theta_{\text{Rich}} = 2\Theta_2 - \Theta_1. \tag{1.4.18}$$

We compare these Theta estimation methods against the benchmark value with the same FD solution. Figure 1.5 illustrates the Theta Θ estimates of the studied

numerical methods compared to the FD benchmark in the region where $90 \leq S_0 \leq 110$.

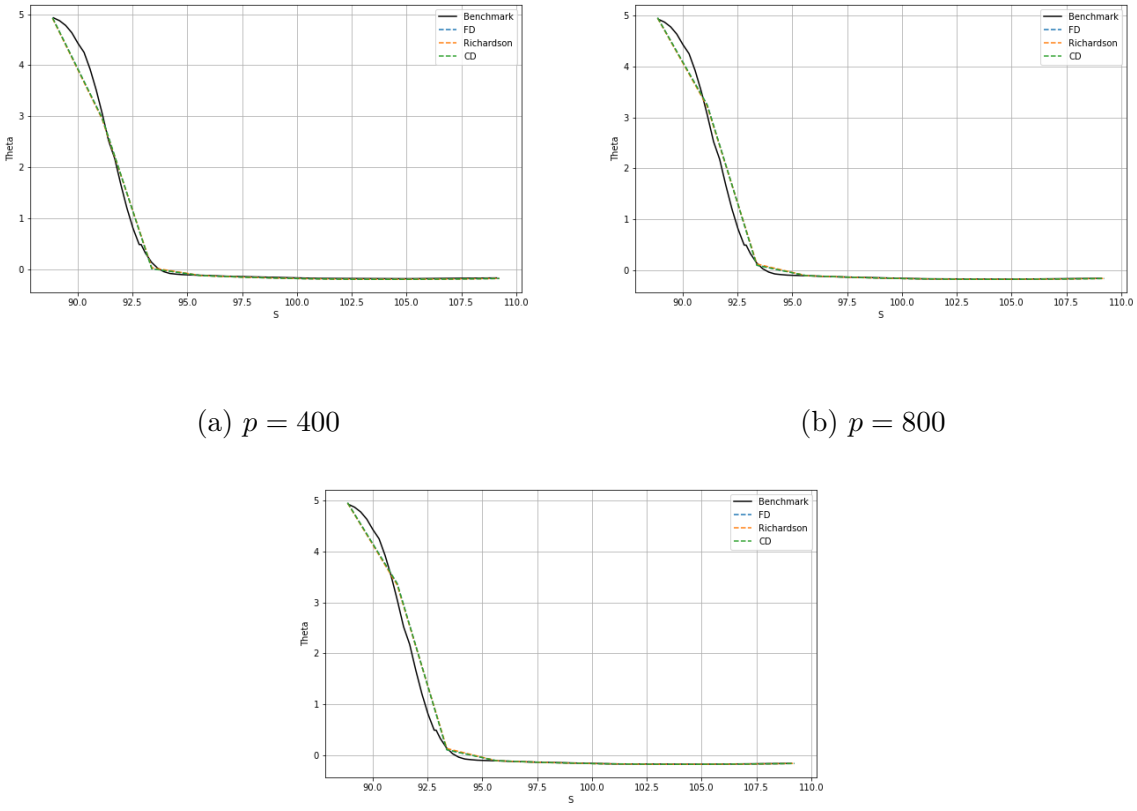

(a) $p = 400$

(b) $p = 800$

(c) $p = 1200$

Figure 1.5: Theta for $S_0 \in [90, 110]$ with different partition sizes p.

The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$. The option is an American put option. Each panel represents the convergence behavior with increasing partition sizes p.

Based on Figure 1.5, each panel presents uniformity in convergence across all evaluated methods, with the central difference (CD) and forward difference (FD) methods demonstrating exact convergence. Table A.3 in Appendix A provide the Theta estimates and benchmark values $\Theta_{\text{Benchmark}}$.

We now conduct an analysis to determine the effect of different time step levels N_t on the accuracy of our Theta estimation. Specifically, we varied N_t to adjust $\Delta t = \frac{T}{N_t}$, to understand its impact on convergence for different moneyness levels. We test the accuracy of the CD method where Theta is Θ_{CD}, since all three approximation methods almost converge at the same rate in this context, against the benchmark from the FD solution $\Theta_{\text{Benchmark}}$ and the same parameters stated previously but with varying $N_t = \{1000, 2000, 3000, 4000\}$. We illustrate the results for each level of moneyness ATM, ITM, and OTM. Figure 1.6 shows the accuracy of Theta estimates for time step levels N_t and partition size p. Specifically, the horizontal axis represents $\Delta t = T/N_t$, the vertical axis represents the partition size p, and the third axis represents the relative error computed using the following formula

$$\text{Relative Error} = \frac{\Theta_{\text{CD}} - \Theta_{\text{Benchmark}}}{\Theta_{\text{Benchmark}}}. \tag{1.4.19}$$

Figure 1.6 provides a three-dimensional visualization of the relative error in estimating Theta using the CD method compared to the FD benchmark. It demonstrates how the accuracy of Theta estimates changes as a function of the time step size $\Delta t = T/N_t$ and the partition size p. We capture the behavior of the CD method for different levels of moneyness.

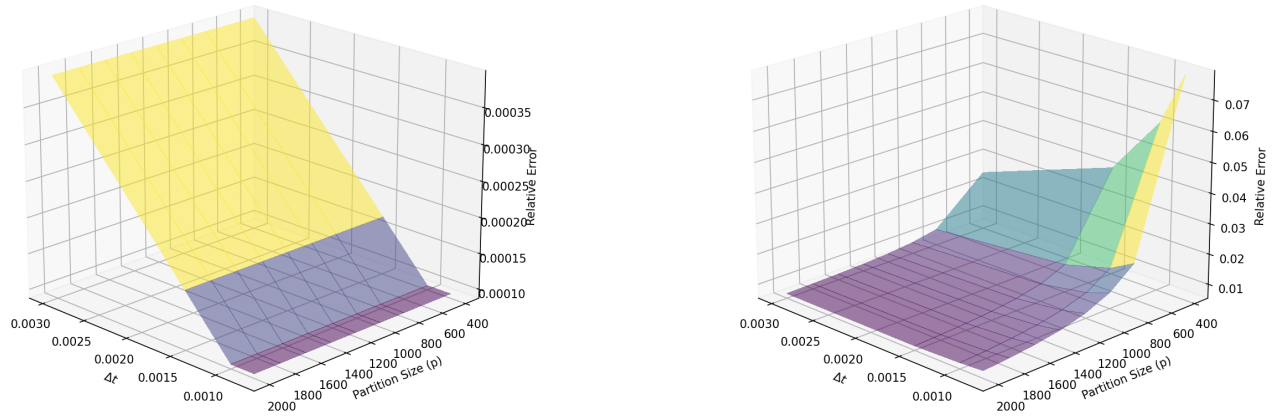

(a) ITM option $S_0/K = 0.85$

(b) ATM option $S_0/K = 1$

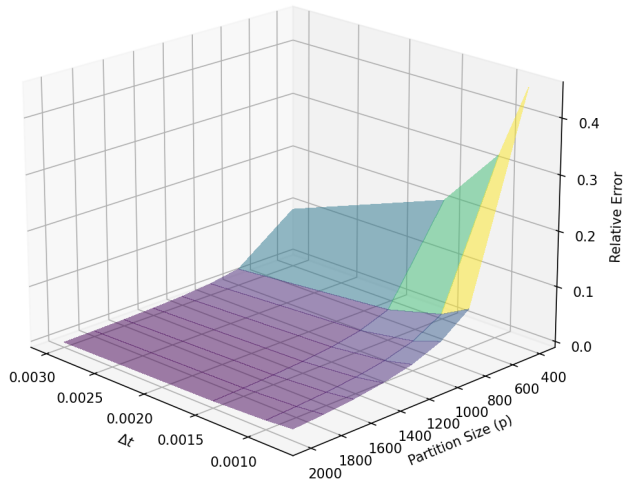

(c) OTM option $S_0/K = 1.15$

Figure 1.6: Relative error of Theta for varying time step levels N_t and partition sizes p.

The parameters of the priced options are: $S_0 \in [85, 115]$ for ITM, ATM, OTM respectively, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, American put option.

For ATM and OTM options, the relative error in the estimation of Theta decreases when both the time increment Δt is reduced and the partition size p is increased. From Table A.4, we can see that the lowest relative error for the ATM and OTM option is at $p = 2000$ and $N_t = 4000$. For the ATM option, we report the lowest relative error of 0.006909 at $p = 2000$ and $N_t = 4000$. Similarly, for the OTM option, we report a relative error of -0.001385. This indicates that smaller time steps combined with a finer partition yield more precise theta estimates for these options. The reduced Δt allows for a better approximation of the time sensitivity, while the larger p improves the resolution of the underlying asset price.

In contrast, for ITM options, the partition size does not appear to affect the accuracy of Theta estimation. The Theta values vary only very slightly as the partition size increases, with changes occurring only at the seventh digit after the

decimal point. Consequently, the primary factor influencing the error is Δt, with smaller time increments resulting in a reduced error, as shown in Table A.4. The lowest relative error is 0.000094. This suggests that, for ITM options, focusing on decreasing the time increment is sufficient to obtain reliable theta estimates, without needing to increase the partition size and face additional computational costs.

1.4.6 Vega Estimation

Denoted by $\mathcal{V}$, Vega represents an option's price sensitivity to changes in the underlying asset's volatility $\mathcal{V} = \frac{\partial V}{\partial \sigma}$. It is crucial for assessing how an option's value changes with fluctuations in market volatility. Under the Black and Scholes (1973) model, volatility is assumed to be constant. However, analyzing Vega remains important because it provides valuable insights into the theoretical sensitivity of the option's price to changes in volatility. This theoretical value of Vega is essential, as it helps quantify how much an option's price would move for a unit change in volatility. By understanding this, one can anticipate how the value of their positions may change with market conditions that deviate from the Black and Scholes (1973) assumptions.

We estimate Vega numerically by calculating the option price at slightly different levels of volatility and then determining the rate of change in price to approximate this sensitivity, allowing for a better understanding of the effect of volatility on the option price. We approximate Vega $\mathcal{V}$ using the FD approximation

$$\mathcal{V}_{\text{CD}} = \frac{\tilde{v}(\sigma + \Delta\sigma) - \tilde{v}(\sigma)}{\Delta\sigma}, \tag{1.4.20}$$

where $\tilde{v}(\sigma)$ and $\tilde{v}(\sigma + \Delta\sigma)$ are the approximated option prices with respective volatilities σ and $\sigma + \Delta\sigma$, and $\Delta\sigma$ is a small change in the volatility.

We use the FD approach of Section (1.3) as a benchmark with $\Delta\sigma = 10^{-6}$ and apply Equation (1.4.20) for the estimation. The choice of $\Delta\sigma = 10^{-6}$ was motivated by the accuracy of the estimate when compared to the closed-form solution of the Black and Scholes (1973) model for a European option. We compare the estimated Vega of SDP where we consider $\Delta\sigma = 10^{-6}$ against its benchmark value within $S_0 \in [90, 110]$, to evaluate the accuracy of the SDP estimate compared to the FD benchmark.

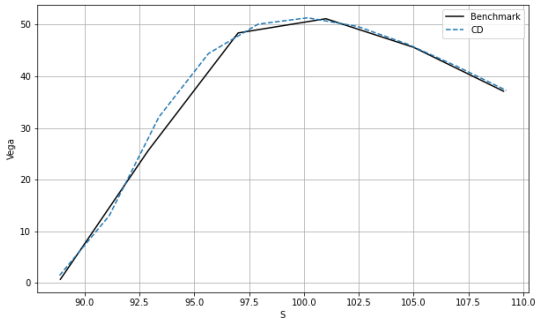

(a) $p = 400$

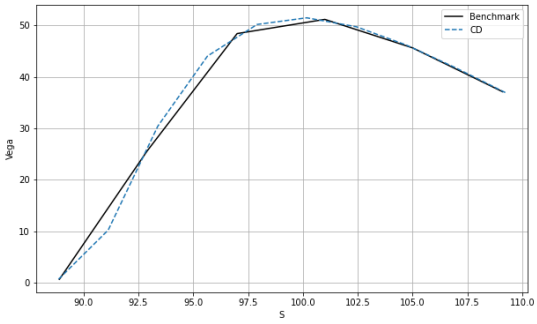

(b) $p = 800$

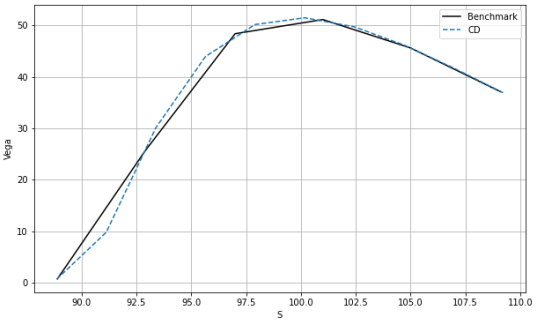

(c) $p = 1200$

Figure 1.7: Vega for $S_0 \in [90, 110]$ with different partition sizes p.

The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$. The option is an American put option. Each panel represents the convergence behavior with increasing partition sizes p.

Figure 1.7 shows that the CD method using SDP provides a good approximation for Vega mainly in the ATM and OTM regions. Table A.5 reports the Vega estimates for SDP and the benchmark.

We now investigate to determine the optimal $\Delta\sigma$ that yields the best convergence of SDP to our benchmark, we perform an analysis across different grid sizes p and volatility increments $\Delta\sigma$. Specifically, we examine the effect of different values of $\Delta\sigma = \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$. The goal is to understand how varying these parameters can influence convergence across options with different moneyness levels. The results of this analysis are illustrated in Figure 1.8, where the horizontal axis

represents $\log(\Delta\sigma)$, the vertical axis represents the partition size p, and the third axis represents the relative error computed using Equation (1.4.19).

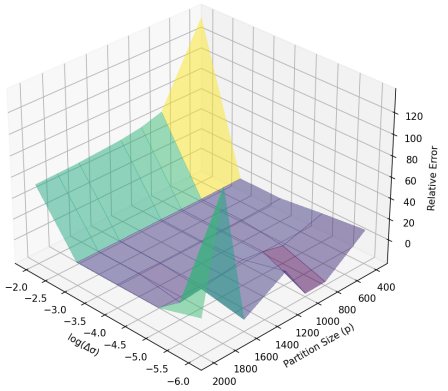

(a) ITM option $S_0/K = 0.85$

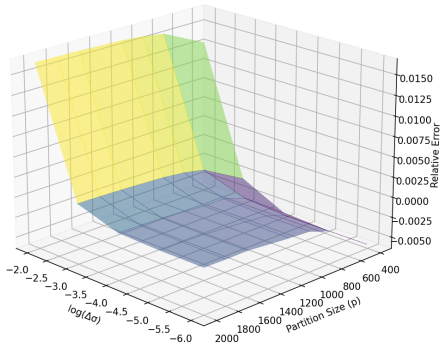

(b) ATM option $S_0/K = 1$

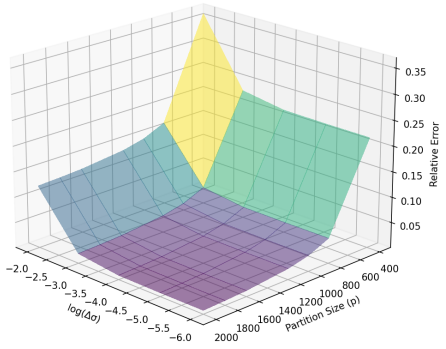

(c) OTM option $S_0/K = 1.15$

Figure 1.8: Relative error of Vega for varying delta sigma $\Delta\sigma$ values and partition sizes p.

The parameters of the priced options are: $S_0 \in [85, 115]$ for ITM, ATM, OTM respectively, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, American put option.

For ATM and OTM options, the Vega estimate shows that the most precise results are obtained with the largest partition size and the smallest change in volatility $\Delta\sigma$. From Table A.6, we can see that at $p = 2000$ and $\Delta\sigma = 10^{-6}$, the relative error is -0.0007330 and 0.007732 for the ATM and OTM option, respectively. This suggests that using finer partitions in conjunction with minimal volatility increments yields a more stable and reliable numerical differentiation.

In contrast, estimating Vega for ITM options presents significant instability, likely involving substantial noise in the results. For $p = 2000$, the relative error begins at 52.965 when $\Delta\sigma = 0.01$, decreases to a range between -0.807 and -0.469, and then rises again to 12.204 as $\Delta\sigma$ approaches 10^{-6}. The most accurate estimations seem to occur when $\Delta\sigma$ is between 10^{-2} and 10^{-3}, with relative errors increasing significantly outside this range, reaching values as high as 138.617. This instability in Vega estimation for ITM options can be attributed to the inherently weaker sensitivity of these options to volatility changes, compared to their ATM and OTM counterparts. Since ITM options are less likely to be exercised, changes in volatility exert less influence on their value. Consequently, the lower sensitivity can make the Vega estimate for these options vulnerable to numerical errors, especially when the changes in volatility are small. The reduced impact of volatility on ITM option prices leads to smaller observed price differences, which in turn can be disproportionately affected by numerical errors during the estimation of Vega.

Finally, it is clear that considering changes in volatility is crucial for accurately pricing options, as volatility shifts can significantly impact their value. This highlights the importance of selecting the right model that accounts for such volatility dynamics. For instance, the Generalized Autoregressive Conditional Heteroskedasticity model (GARCH) of Bollerslev (1986) captures time-varying volatility by modeling it as a function of past returns and volatility, allowing for better adaptation to market fluctuations and observed patterns of volatility over time. Similarly, we can refer to the Heston (1993) model introduces stochastic volatility, where the volatility itself follows a random process, making it more suitable for capturing real market conditions where volatility is not constant.

1.5 Computational Time and Efficiency in Greeks Estimation

In this Section, we conduct a comprehensive analysis of the computational

time required by each numerical method for calculating the Greeks of a put option with parameters: $S_0 = 100$, $K = 100$, $T = 3$, total time steps $N_t = 4000$, exercise opportunities of 50 per year $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, and $\sigma = 0.1$. By comparing the performance of various methods, we aim to assess the trade-offs between accuracy and efficiency.

This analysis focuses on determining which method offers the most computationally efficient solution while maintaining precision in estimating Delta, Gamma, and Theta. The results illustrated in Figure 1.9 will guide the selection of the optimal method. The relative errors are expressed in percentage %, the computational time in seconds and each point on the plot represents computational time versus relative error for different partition sizes p.

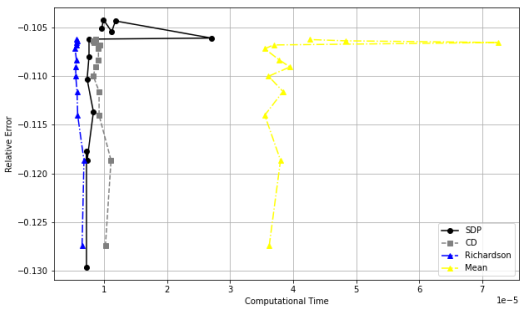

(a) Delta Δ

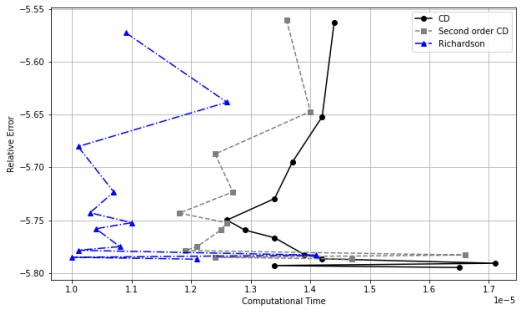

(b) Gamma Γ

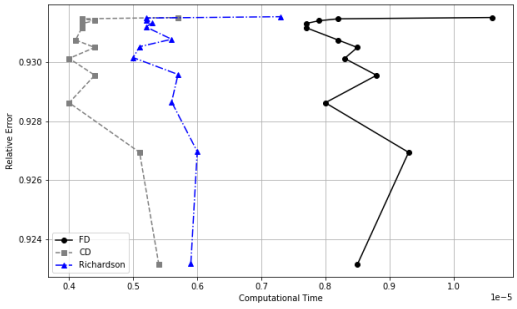

(c) Theta Θ

Figure 1.9: Computational time in seconds for various Greek methods.

The parameters of the American put option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_t = 4000$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, and $\sigma = 0.1$.

Panel (a) reveals that the SDP method stands out as the optimal choice for estimating Delta, offering a balance between minimizing the relative error and optimizing computational efficiency. This method is computed within a time of 9.6×10^{-6}. Moving to Panel (b), the Richardson method proves to be the most effective in calculating Gamma as it is computing within 1.2×10^{-5} seconds. Lastly, Panel (c) highlights the CD method as the most advantageous for Theta calculations in terms of accuracy and computational efficiency within 1.1×10^{-5} seconds.

However, since these computations can be efficiently vectorized and are not as computationally intensive as the transition matrices in Equation (1.2.7), we can say that no single method dominates across each Greek in terms of computational time. Table A.7 in Appendix A reports all times required to compute each method.

To summarize the findings of this Section (1.2), our results demonstrate strong performance in certain areas while showing limitations in others, depending on the approximation method used for the Greeks.

When estimating Delta, the SDP method produces results that closely align with the FD benchmark, especially when the partition size is reasonably large. However, alternative approximations like Richardson tend to become unstable at lower granularity levels, which makes them less reliable in such cases. Estimating Gamma, on the other hand, presents significant challenges. Stability issues arise, particularly when $S < K$, and the relative error remains high even when using finer partitions. The Richardson method exhibits considerable instability, while the central difference (CD) approach shows some improvement when the partition size exceeds $p = 2000$. The Theta estimation is much more stable compared to Gamma. Both FD and CD methods converge efficiently, and reducing the time step Δt consistently increases accuracy. However, for ITM options, increasing the partition size p does not lead to significant improvements, suggesting that unnecessary computational costs can be avoided by optimizing the granularity of the partition. This highlights the importance of balancing precision and efficiency in the numerical approach. Vega estimation is mainly relevant in a theoretical context, as the Black and Scholes (1973) model assumes constant volatility. The analysis indicates that accuracy is highly sensitive to the choice of $\Delta\sigma$ variations, which can limit its practical usefulness. This suggests that while Vega can be computed, its interpretation should be approached with caution, especially in models where volatility is assumed to be fixed.

In conclusion, the SDP approach proves effective for Delta and Theta estimation, with relatively fast convergence. However, Gamma remains a weak point, showing excessive sensitivity to discretization choices. While Vega can be calculated, its practical value is limited due to the constraints of the Black and Scholes (1973) framework. Optimizing the selection of partition size and time step is crucial for improving accuracy while managing computational costs efficiently.

Chapter 2

Pricing a Panel of American Options

In this Section, we show how a single iteration of SDP can be used to price a panel of options with varying moneyness levels and maturities, assuming nothing but the homogeneity property shared by most option pricing models (Létourneau and Stentoft, 2022). What makes this approach versatile is that any time-homogeneous model that is grounded in Markovian state variables can be accommodated. For instance, the Black and Scholes (1973) model, the Heston (1993), and the Cox et al. (2005) (CIR) which is an interest rate model, all exhibit Markovian dynamics and can be utilized within this framework.

To implement our methodology, we first draw inspiration from the concept of price homogeneity introduced by Létourneau and Stentoft (2022). They utilize this concept in the context of simulation-based option pricing, while we apply it within a SDP framework. Price homogeneity allows reformulating the challenge of pricing a panel of options that differ in their strike prices K and underlying asset values S_0, into a scenario where the underlying asset values are still different but the strike prices are identical. This is beneficial because the early-exercise strategies are not dependent on the initial state variable values. This translates into a substantial reduction of the complexity and resource requirements when pricing a set of options.

It is straightforward to determine a singular estimated optimal strategy for early exercise that applies universally across various options, accounting for the differing initial scaling of stock prices. However, this would still require an iterative process of tracing back through each option's paths for payoff calculations, which is characteristic of regression-based methods like LSM of Longstaff and Schwartz (2001). Instead, we leverage the structure of the SDP algorithm, which approximates option values by optimizing the stopping time based on a range of initial state variables. This approach offers significant advantages: only a single valuation run

is necessary for pricing an entire panel of options. In simpler terms, rather than recalculating the price of each option individually, the SDP method provides an efficient, unified framework where options with different moneyness levels are priced through a single iteration, evaluating the continuation value as a function of the current state variables.

To make this method more generalizable, Létourneau and Stentoft (2022) considered pricing options with differing maturities. Leveraging the fact that option pricing models operate on Markovian processes, options with various maturities can be efficiently evaluated using the optimal policy derived from the option with the longest maturity. In the SDP algorithm, the most resource-intensive step is the computation of the transition matrices A_{ki} and B_{ki} in Equation (1.2.7). Therefore, the initial effort invested in calculating the transition matrices for the longest maturity option significantly reduces the additional computational costs for pricing options with shorter maturities.

Price homogeneity plays a central role in simplifying the numerical pricing of a panel of American options. By normalizing asset and strike prices to a single reference strike, the pricing problem is transformed into a standardized form, allowing a single value function to represent multiple contracts. This significantly reduces computational cost and ensures that the same optimal early-exercise strategy applies across all options in the panel. The pricing methodology builds on this property by employing the SDP defined in Chapter 1. The SDP computes the option value function over a standardized state space, which can then be rescaled to recover the prices of individual options with different strikes. For varying maturities, the method efficiently reuses the simulated asset paths generated for the longest-maturity option to price all shorter-maturity contracts. This unified approach enables the entire panel to be priced in a single computation, improving numerical stability and reducing computational time without sacrificing accuracy.

2.1 Pricing Options with Varying Strike Prices

The principle of homogeneity allows us to reformulate the pricing optimization problem in terms of strike prices rather than underlying asset values. We construct a standardized grid where the strike price $K = 1$ and includes all option prices and Greeks, enabling a unified valuation framework throughout the panel of options. This grid provides a structured representation where each point corresponds to a specific combination of state variables and time to maturity. Through this approach, we efficiently approximate the option values and their sensitivities by locating them in the standardized grid. We first clearly establish the concept of price homogeneity

with the following definition.

Definition 2.1.1. *A function is homogeneous of order k if*

$$f(\ell \times x_1, \ell \times x_2, \times, \ell \times x_n) = \ell^k \times f(x_1, x_2, \ldots, x_n) \quad \text{for any } \ell > 0,$$

where f is a function of multiple variables $x_1, x_2, \ldots, x_n$ and ℓ is a positive scalar. The function is said to be homogeneous of order k if scaling all the arguments by ℓ results in the function's value being scaled by ℓ^k.

The homogeneity property applies under specific conditions on both the asset dynamics and the option payoff. In particular, the asset price process must follow a time-homogeneous and Markovian process, such as the geometric Brownian motion used in the Black-Scholes model. Additionally, the payoff function must be homogeneous of degree 1 in the asset price S and the strike price K, as is the case for standard vanilla call and put options. These conditions ensure that scaling both S and K by a common factor $\ell > 0$ results in the option price being scaled by the same factor, thereby satisfying homogeneity of order 1.

Theorem 2.1.1. *Under the Black and Scholes (1973) model, the price of an American or European option is homogeneous of degree 1 in the underlying asset price S and the strike price K, that is, for any $\ell > 0$,*

$$V(\ell S, \ell K, T, \sigma, r, q) = \ell V(S, K, T, \sigma, r, q),$$

where $V(S, K, T, \sigma, r, q)$ denotes the price of the option.

Proof. Consider the price of a European put option $V_{\text{eur}}(S, K, t)$ under the Black and Scholes (1973) model, which satisfies the following PDE

$$\frac{\partial V_{\text{eur}}}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V_{\text{eur}}}{\partial S^2} + (r - q) S \frac{\partial V_{\text{eur}}}{\partial S} - r V_{\text{eur}} = 0.$$

To prove the homogeneity property of order 1 in the strike price K and underlying asset price S, we scale the variables K and S by a factor $\ell > 0$ $V_{\text{eur}}(\ell S, \ell K, t)$, and use the chain rule for partial derivatives

$$\frac{\partial V_{\text{eur}}(\ell S, \ell K, t)}{\partial(\ell S)} = \frac{\partial V_{\text{eur}}}{\partial S} \cdot \frac{\partial S}{\partial(\ell S)} = \frac{1}{\ell} \frac{\partial V_{\text{eur}}}{\partial S},$$

$$\frac{\partial^2 V_{\text{eur}}(\ell S, \ell K, t)}{\partial(\ell S)^2} = \frac{\partial}{\partial(\ell S)} \left(\frac{\partial V_{\text{eur}}(\ell S, \ell K, t)}{\partial(\ell S)} \right) = \frac{\partial}{\partial(\ell S)} \left(\frac{1}{\ell} \frac{\partial V_{\text{eur}}}{\partial S} \right) = \frac{1}{\ell} . \frac{1}{\ell} \frac{\partial^2 V_{\text{eur}}}{\partial S^2} = \frac{1}{\ell^2} \frac{\partial^2 V_{\text{eur}}}{\partial S^2}.$$

We substitute these into the scaled PDE

$$\frac{\partial V_{\text{eur}}}{\partial t} + \frac{1}{2}\sigma^2\ell^2 S^2 \frac{1}{\ell^2}\frac{\partial^2 V_{\text{eur}}}{\partial S^2} + (r-q)\ell S \frac{1}{\ell}\frac{\partial V_{\text{eur}}}{\partial S} - rV_{\text{eur}}(\ell S, \ell K, t) = 0.$$

If we simplify, we get the original PDE for $V_{\text{eur}}(S, K, t)$, indicating that

$$V_{\text{eur}}(\ell S, \ell K, t) = \ell \times V_{\text{eur}}(S, K, t), \text{where } \ell = 1.$$

Thus, the European option pricing function $V_{\text{eur}}(S, K, t)$ is homogeneous of order 1 in S and K. Similar arguments can be used for American options (Létourneau and Stentoft, 2022). □

Consider that the option pricing function P is homogeneous of order one in the underlying asset S and the strike price K for both European and American options, such that

$$P(S \times l, K \times l, T, \sigma, r, q) = l \times P(S, K, T, \sigma, r, q). \tag{2.1.1}$$

Applying the homogeneity principle on Equation (2.1.1), we get

$$P(S, K, T, \sigma, r, q) = K \times P(S/K, 1, T, \sigma, r, q). \tag{2.1.2}$$

Equation (2.1.2) implies that to obtain option prices for a vector of strike prices $\mathbf{K}$, one can alternatively calculate them by assuming varying values of the underlying asset $\mathbf{S}$ while keeping the strike price constant.

Because early exercise strategies do not depend on the starting value of the underlying variables, and we are dealing with models that have the Markov property, it's possible to apply the same optimal strategy for early exercise to price all the options. This means that the complex task of figuring out the best time to exercise the option needs to be performed only once. Price estimates can be acquired by assessing the approximation given in Equation (1.2.2) of Section (1.2.3) for the suitable level of moneyness and then multiplying the result by the strike price K of the specific option.

2.2 Pricing Options with Varying Maturities

It is also possible to price, in one go, options with several different maturities by utilizing the fact that all the required information to price shorter maturity options is in the paths used to price the longest maturity option. If we consider time in terms of *time to maturity* $T - t$ instead of *time from today* t, we can use the same exercise strategy for all options, regardless of their maturity duration.

Because the decision to exercise early is independent of the option's initial maturity and rather dependent on the remaining time to maturity, given that we are utilizing Markovian models, we can employ the same optimal early exercise strategy for pricing all options that have the same remaining time. Consequently, calculating the optimal time to exercise is only necessary once for the option with the longest duration to maturity, which makes it feasible to price multiple options with varying maturities in a single step.

Figure 2.1 effectively illustrates the concept. The horizontal axis represents the discretized time grid used to compute the value function for the longest maturity option $T_{\max}$. The point t_2 corresponds to the maturity date T of a shorter option. Since the SDP value function is computed backwards from $T_{\max}$, the value of the shorter-maturity option is retrieved from the node $T_{\max} - T$ in the grid. The red arrow highlights the time-to-maturity of the shorter option. Although the shorter option reaches its maturity before $T_{\max}$, its boundary condition is treated in the unified SDP grid by indexing appropriately.

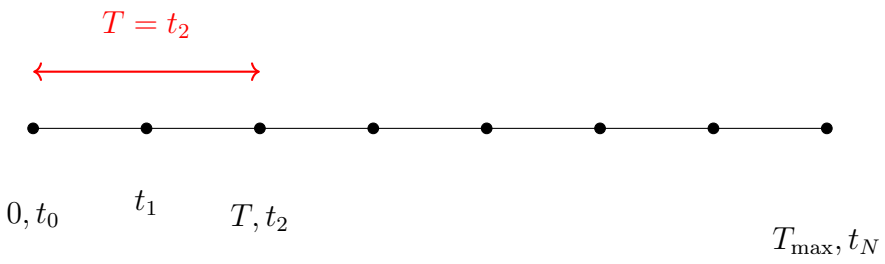

Figure 2.1: Pricing a panel with different maturities.

Therefore, pricing options with shorter maturities involve insignificant additional costs. By leveraging the longest maturity path, we efficiently reuse information to simplify their valuation.

2.3 Selection of the Options Panel and Results

To conduct a comprehensive analysis of option pricing and the behavior of the Greeks (Δ, Γ, Θ, $\mathcal{V}$) under varying conditions, a diverse panel of options was selected. The chosen options cover different levels of moneyness and maturity, which provides a broad spectrum for evaluation. We also allow 50 exercise opportunities $N_{\text{ex}} = 50$ within each year. The selected parameters are as follows

Moneyness level	S_0	K	T
ITM - Short Maturity	85	100	1
ITM - Medium Maturity	85	100	2
ITM - Short Maturity	90	100	1
ATM - Long Maturity	100	100	3
OTM - Long Maturity	115	100	3

Table 2.1: Options selection for the panel pricing.

The selection of these options is intended to cover a range of option characteristics including OTM, ATM, and ITM scenarios. By examining these diverse options, it is possible to observe the impact of different moneyness levels and maturities on the studied Greeks under both low ($\sigma = 0.1$) and high ($\sigma = 0.4$) volatility settings. We include more ITM options in the panel because we aim to verify whether these options are less sensitive to changes in T and S_0, as they are almost certain to be exercised. Additionally, determining their Greeks is more challenging compared to ATM and OTM options. By incorporating multiple ITM options, we aim to investigate whether their Greeks exhibit a reduced variation across different maturities and initial prices.

In our analysis, we employ the FD algorithm as outlined in Section (1.3) to perform a comparative evaluation of our results. The results are visualized using the relative error measure to indicate how far an estimated value is from the benchmark value from Equation (1.4.19).

Each of the Figures below represents a distinct metric, option price V, Delta Δ, Gamma Γ, Theta Θ, and Vega $\mathcal{V}$, evaluated in both low-volatility and high-volatility environments. The Figures are organized into five panels: each panel shows the respective metric as a function of the partition size p, illustrating the evolution of the relative error as p increases for various option scenarios.

We eventually focus on presenting the values only for partition sizes greater than $p = 500$, to effectively visualize which method converges the fastest to zero. In certain panels, the scales are explicitly marked at the top of the vertical axis to denote very fine measurements, underscoring the high precision of estimates.

2.3.1 Low Volatility Setting

Figures 2.2, 2.3, 2.4, 2.5, 2.6 illustrate the option price, Delta, Gamma, Theta and Vega respectively, in a low volatility environment. Each Figure has five panels; (a), (b), and (c) display results for ITM options with short and medium maturities, while panel (d) shows the ATM option with long maturity, and panel (e) illustrates

the OTM option with long maturity. Tables C.1, C.2, C.3, C.4, C.5 of Appendix C report all relative errors for these Figures.

Following Figure 2.2, we can see that the SDP algorithm performs well in estimating the option price V for all options. For ITM options with short and medium maturities, the relative error remains stable and unchanged across all partition sizes p, where we report relative errors of -0.067% and -0.042% respectively, which indicates a reliable performance regardless of computational granularity. The consistent behavior across different maturities suggests that maturity does not significantly impact the option's price under these conditions.

On the other hand, for the ITM option with varying initial asset value S_0, as well as for ATM and OTM options, the relative error steadily decreases as the partition size p increases. This suggests that using a finer partition effectively reduces the estimation error in these scenarios, leading to improved results. Specifically, the relative error reaches -0.5% at $p = 2000$ for the OTM option in panel (e), -0.02% for the ATM option in panel (d), and -0.01% for the ITM option in panel (c).

Based on Figure 2.3 we can see that all Delta methods report a relative error of a 10^{-8} order for panels (a) and (b), which reflects their alignment with the benchmark. Panels (c), (d) and (e) have relative errors ranging between 1% and 0% which is very small. We can say that the SDP method shows remarkable stability for ITM options, as it reaches a relative error of 0%. In the case of ATM and OTM options, Richardson performs best in terms of stability and convergence with a relative error of 0%, followed by the SDP method.

Figure 2.4 reveals that the relative error for Gamma is consistently higher for ITM options compared to ATM and OTM options, irrespective of S_0, T, or the partition size p. We observe that the option in panel (c), which has an increased initial asset price of $S_0 = 90$, exhibits a lower relative error compared to the ITM options in panels (a) and (b), likely because the option in panel (c) approaches the ATM scenario. In low volatility, ITM options are almost certain to be exercised, and their payoff becomes nearly linear as the option price closely tracks its exercise value. This linearity reduces the sensitivity of the option's Delta to changes in the underlying asset price, which is Gamma.

As a result, numerical methods that rely on finite differences are not able to accurately capture the small variations required to compute Gamma. Furthermore, the theoretical Gamma value for ITM options in this setting tends to be very small, and when combined with numerical errors, the relative error becomes disproportionately large. For ATM and OTM cases, Gamma is generally better approximated compared to ITM options in a low volatility setting. This is because these options exhibit higher sensitivity to changes in S_0, making their Gamma values

more stable and easier to capture numerically. The relative errors for ATM and OTM options remain consistently low, ranging between -0.5% and 1% and are all reported in Tables C.4 and C.5.

Figure 2.5 reveals that panels (a) and (b) have a stable relative error of order 10^{-5} for all methods. In contrast, panels (c), (d) and (e) report a relatively higher relative error. The ITM option with a higher initial asset value S_0 has relative errors between 8% to 12% and the CD method performs best, identical to the FD method. For ATM and OTM options, all three methods yield similar results, with relative errors between 0.5% to 3% and the Richardson method showing a slight improvement over the others.

The higher relative errors observed for ITM options with higher S_0, as well as for ATM and OTM options, can be attributed to their greater reliance on the time value component of the option price. Unlike deep ITM options, where exercise value dominates and time value is negligible, these options are more sensitive to changes in the underlying asset price and time to maturity.

Panel (a) of Figure 2.6 shows a very high and stable relative error. This arises because the sensitivity to volatility becomes negligible, as the option's value is mainly dominated by the exercise value in this case. For panel (b), the errors fluctuate from -40% to 60% which indicates that time value starts contributing more to the option price, similar to panel (c) where errors range from -5% to 25%. The ATM option in panel (d) reveals low but fluctuating relative errors as they vary between -0.1% and 0.075%. The Vega is best approximated by SDP for the OTM option as we can see that the relative error goes from -2% to 0% at $p = 2000$. This is because the OTM option is primarily dominated by its time value since the exercise value is effectively zero. The behavior of the time value in the OTM scenarios is smoother and more predictable, especially in a low-volatility setting, where changes in volatility have a limited impact on the option price. This stability makes Vega easier to approximate numerically for OTM options.

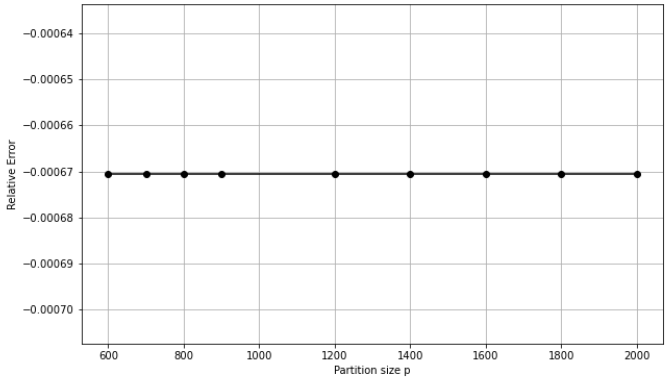

(a) ITM option $S_0 = 85$ - Short Maturity

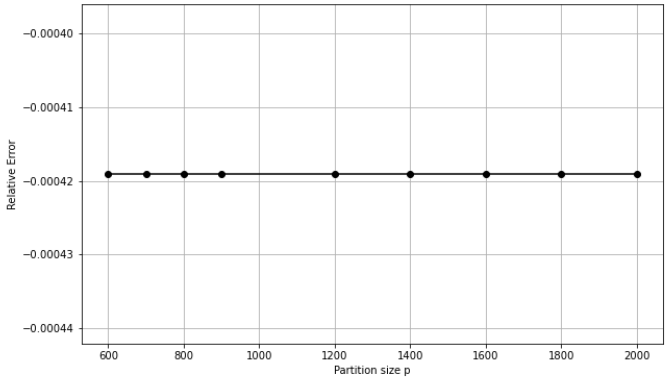

(b) ITM option $S_0 = 85$ - Medium Maturity

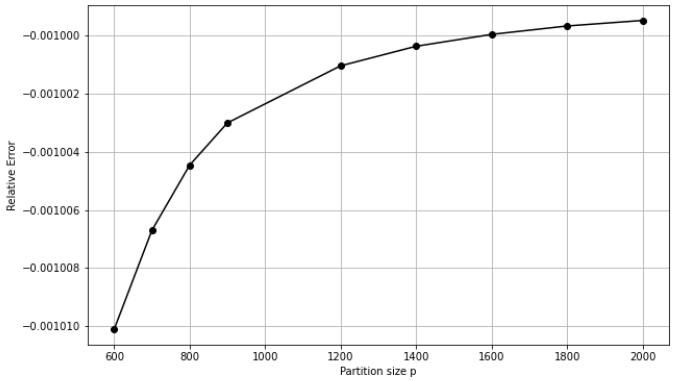

(c) ITM option $S_0 = 90$ - Short Maturity

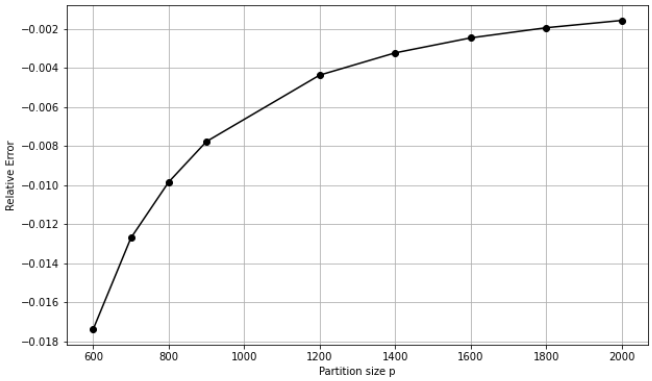

(d) ATM option - Long Maturity

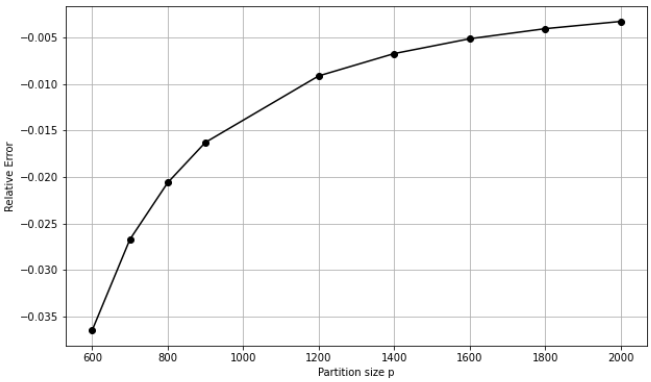

(e) OTM option - Long Maturity

Figure 2.2: Option prices in low volatility.

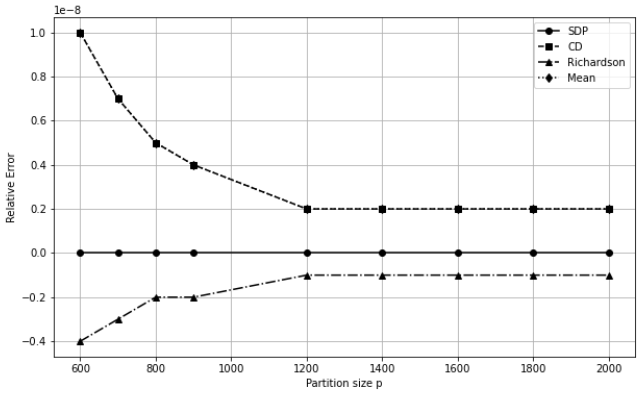

(a) ITM option $S_0 = 85$ - Short Maturity

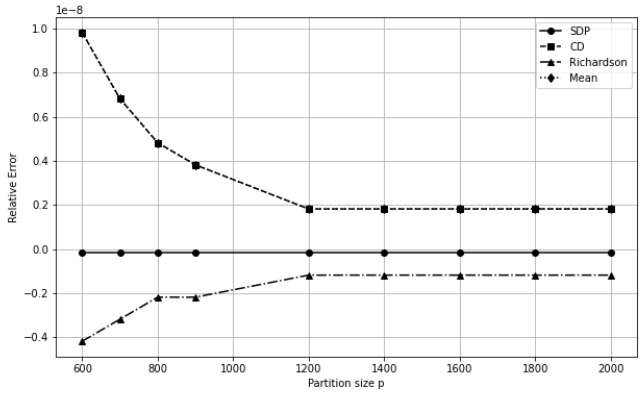

(b) ITM option $S_0 = 85$ - Medium Maturity

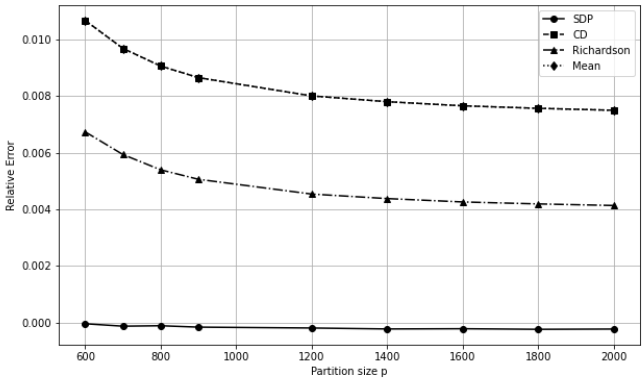

(c) ITM option $S_0 = 90$ - Short Maturity

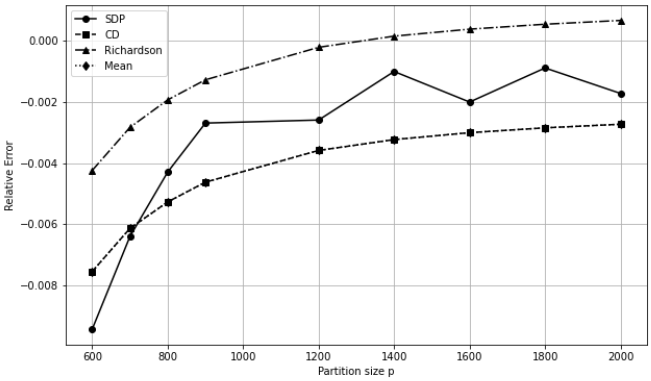

(d) ATM option - Long Maturity

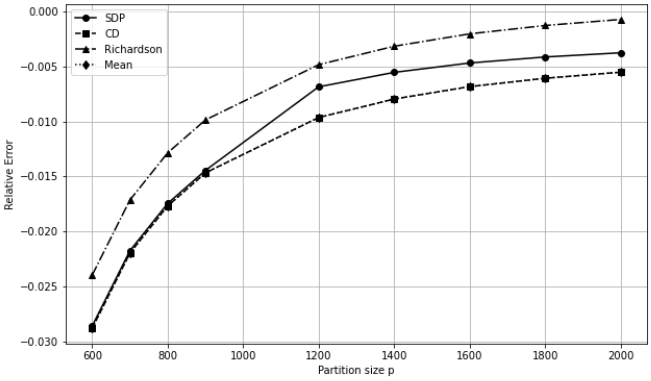

(e) OTM option - Long Maturity

Figure 2.3: Delta in low volatility.

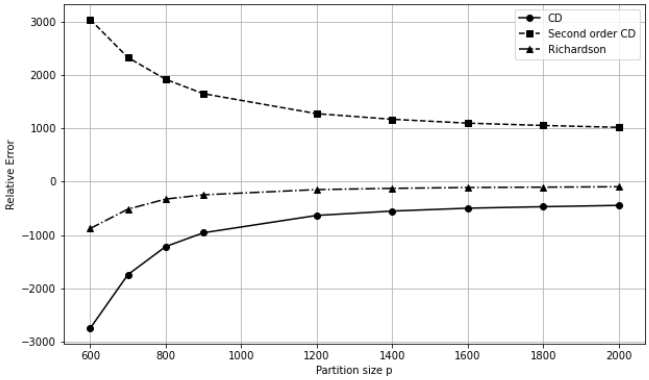

(a) ITM option $S_0 = 85$ - Short Maturity.

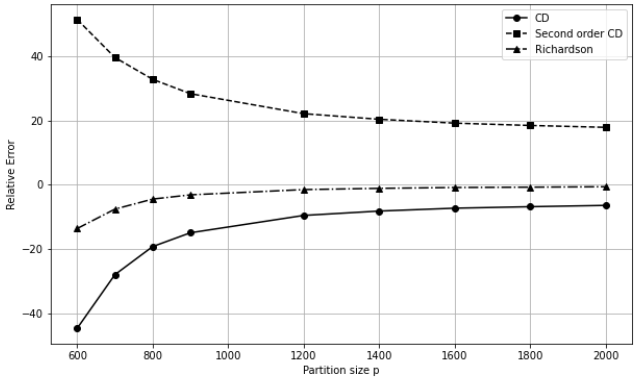

(b) ITM option $S_0 = 85$ - Medium Maturity

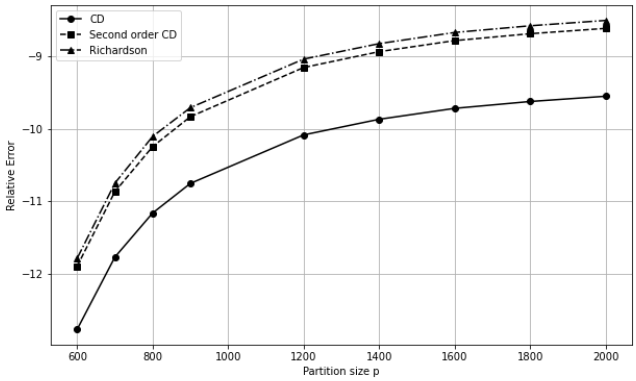

(c) ITM option $S_0 = 90$ - Short Maturity

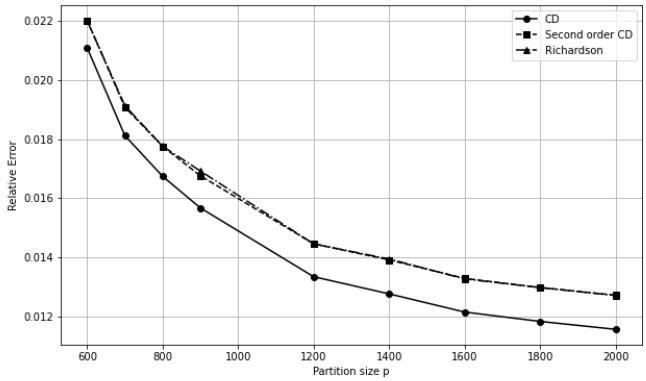

(d) ATM option - Long Maturity

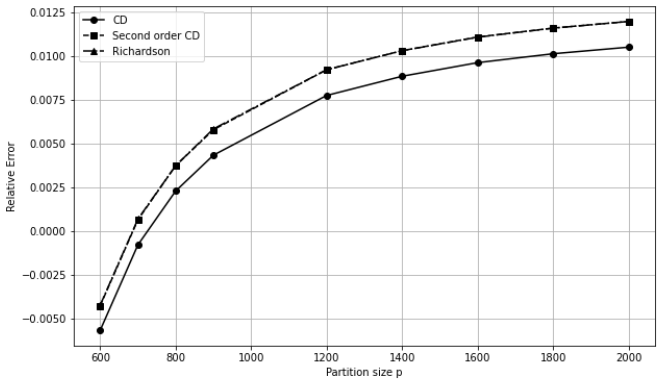

(e) OTM option - Long Maturity

Figure 2.4: Gamma in low volatility.

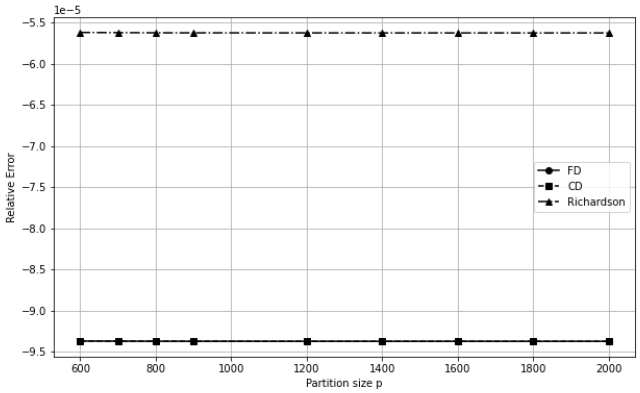

(a) ITM option $S_0 = 85$ - Short Maturity

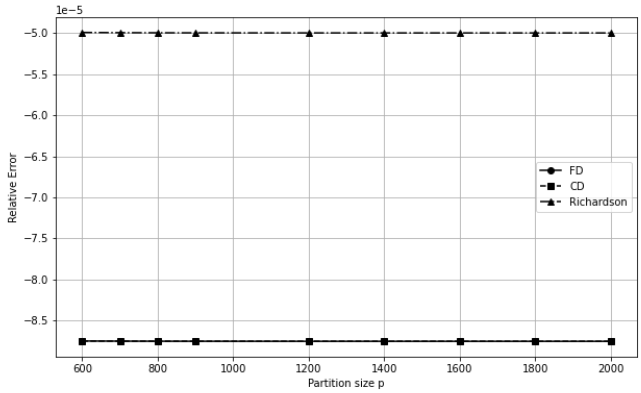

(b) ITM option $S_0 = 85$ - Medium Maturity

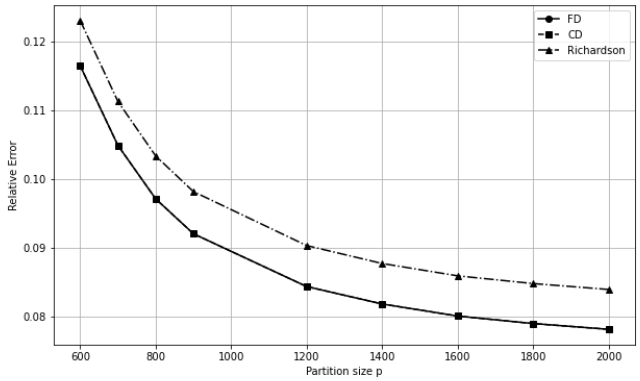

(c) ITM option $S_0 = 90$ - Short Maturity

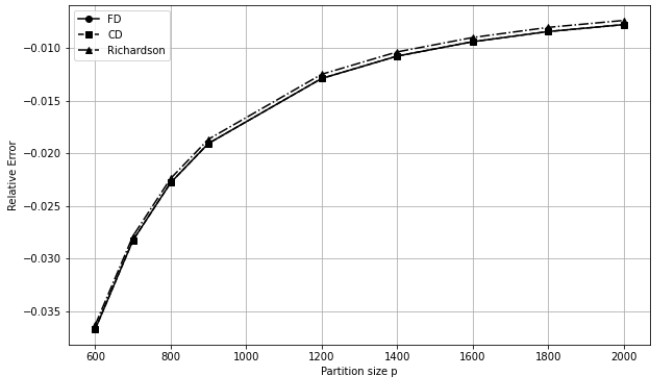

(d) ATM option - Long Maturity

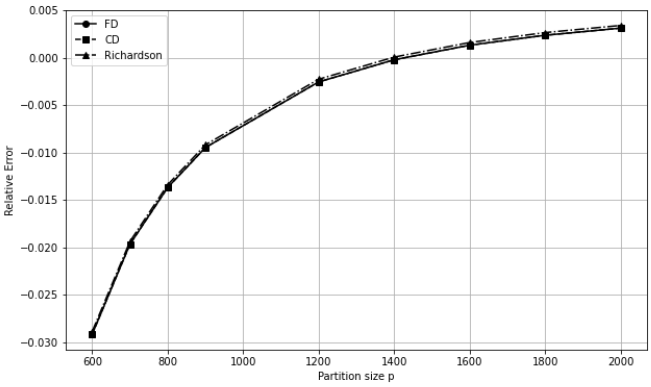

(e) OTM option - Long Maturity

Figure 2.5: Theta in low volatility.

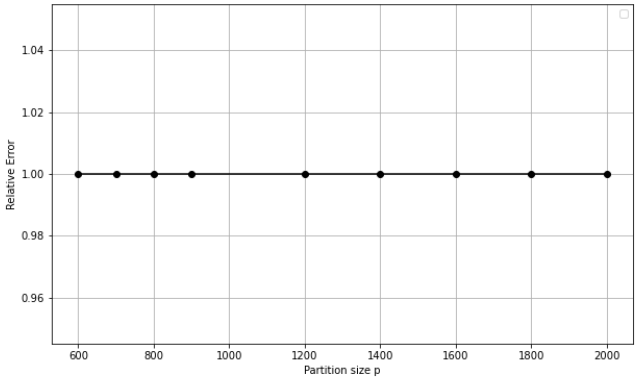

(a) ITM option $S_0 = 85$ - Short Maturity

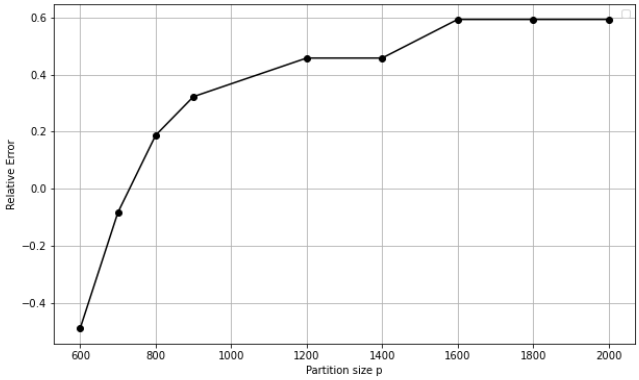

(b) ITM option $S_0 = 85$ - Medium Maturity

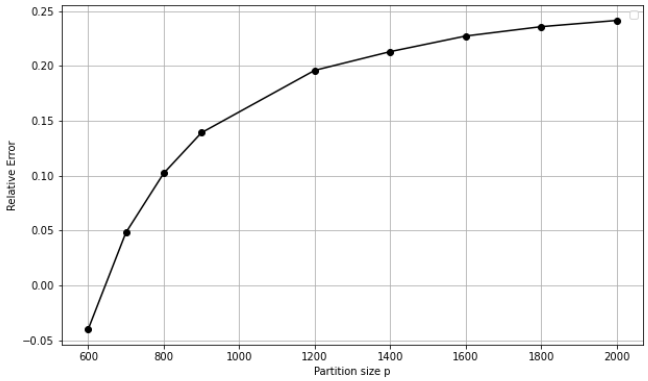

(c) ITM option $S_0 = 90$ - Short Maturity

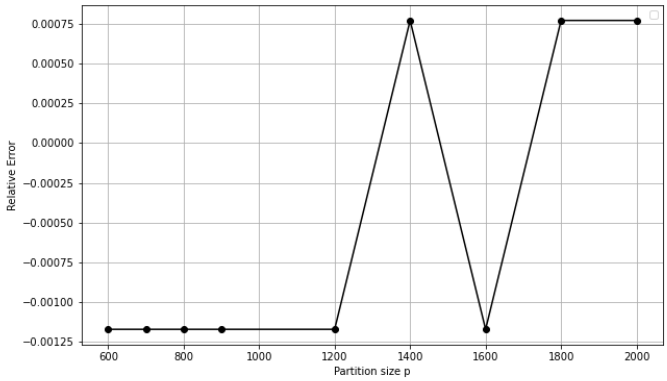

(d) ATM option - Long Maturity

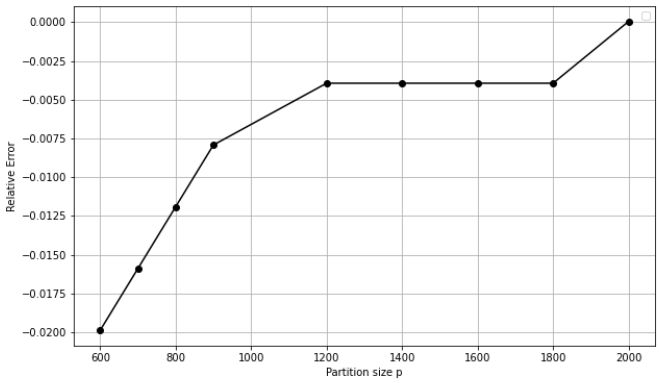

(e) OTM option - Long Maturity

Figure 2.6: Vega in low volatility.

2.3.2 High Volatility Setting

Figures 2.7, 2.8, 2.9, 2.10 and 2.11 illustrate the option price, Delta, Gamma, Theta and Vega respectively, in a high volatility environment. Each Figure has five panels; (a), (b), and (c) display results for ITM options with short and medium maturities, while panel (d) shows the ATM option with long maturity, and panel (e) illustrates the OTM option with long maturity for a high volatility framework. Tables C.6, C.7, C.8, C.9, C.10 in Appendix C report all relative errors for all Figures.

All panels of Figure 2.7 display the same behavior with varying relative error values. Notably, the ITM option with medium duration converges to zero more rapidly than the short-maturity options. Additionally, there is no observed sensitivity to changes in the underlying asset price within the high volatility setting, as the patterns for ITM options with different initial asset prices remain unchanged.

Panels (a), (b) and (c) of Figure 2.8 have the same observed patterns, the SDP method converges the first with a stable relative error of around 0% followed by other methods which are all similar with errors between 8% to 20%, depending on the option's parameters. Panel (d) shows that the SDP method is less stable compared to Panel (e) but remains the best choice to estimate Delta. We can say that Delta is sensitive to the volatility framework, as we have observed very low relative errors in the low volatility setting.

Based on Figure 2.9, in a high volatility setting, the relative errors for Gamma remain elevated for ITM options, ranging between 14% and 36%, but are notably lower compared to the low volatility setting. This improvement occurs because high volatility increases the time value component of ITM options, which makes their payoff less linear and increases the sensitivity of Delta to changes in S_0. This greater sensitivity allows numerical methods to better approximate Gamma. For ATM and OTM options, the relative errors are much smaller, ranging between 0.1% and 0.3%, as these options exhibit a higher sensitivity to changes in S_0, making their Gamma values more stable and easier to estimate numerically.

Across all methods, Gamma estimation is very similar, but the CD method slightly outperforms the others in terms of accuracy. These trends highlight the effect of volatility on the difficulty of Gamma approximation.

Figure 2.10 exhibits significantly higher relative errors for ITM options compared to the low volatility framework. ITM options with short maturities have relative errors ranging between 54% and 57%, whereas ITM options with medium maturity report a bit lower relative errors around 27%. This may be because a longer maturity can mitigate the sensitivity to time decay. In contrast, ATM and OTM

options exhibit much smaller relative errors, ranging between -1.7% and 0.25%, as their time value is more stable and predictable even in high volatility. The greater sensitivity of ITM options to volatility and the dominant role of time value in their pricing under high volatility can explain the higher reported errors in these cases.

In a high volatility setting, Figure 2.11 illustrates that Vega exhibits consistently low relative errors for all options, ranging between -1.2% and -0.01%, regardless of their moneyness levels or maturity. This significant improvement compared to the low volatility case is due to the increased sensitivity of the option price to volatility in high volatility environments, which increases the numerical stability and accuracy of Vega approximation. For ITM options, where the exercise value dominates in low volatility, the higher volatility introduces more time value, making Vega less negligible and more predictable. Similarly, for ATM and OTM options, the smoother and more pronounced impact of volatility on the option price leads to stable and accurate numerical approximations.

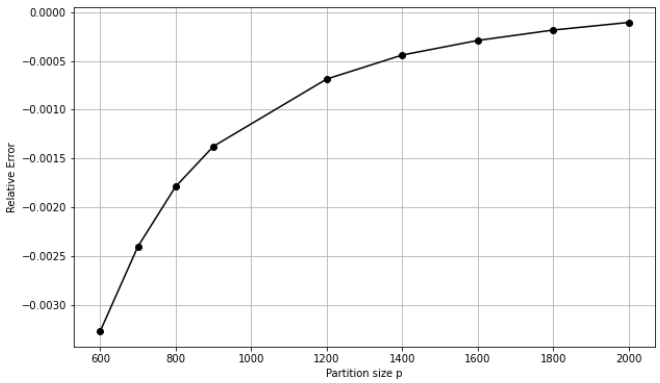

(a) ITM option $S_0 = 85$ - Short Maturity

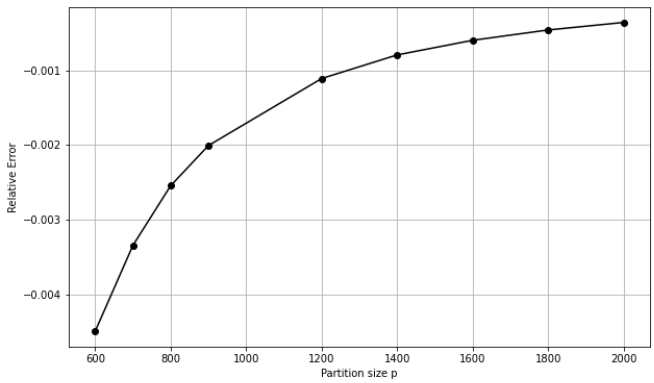

(b) ITM option - Medium Maturity

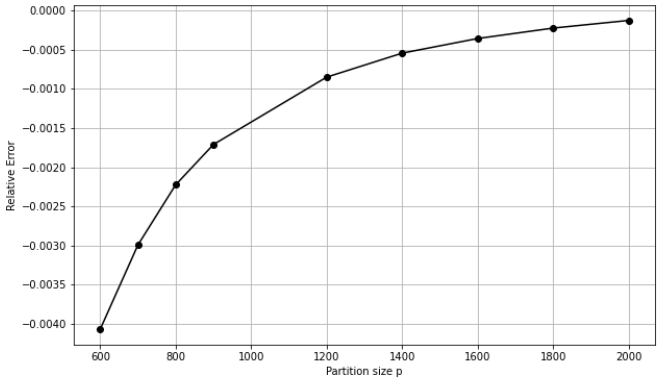

(c) ITM option $S_0 = 90$ - Short Maturity

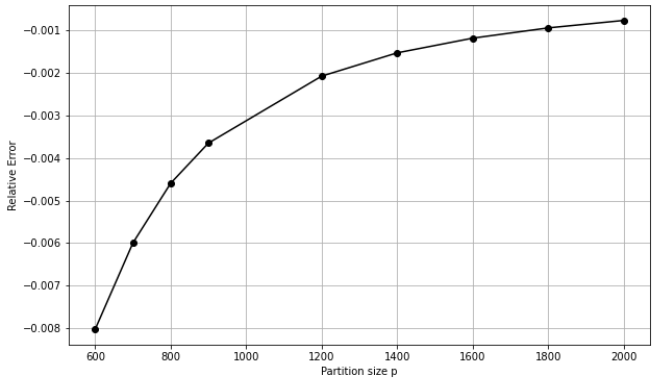

(d) ATM option - Long Maturity

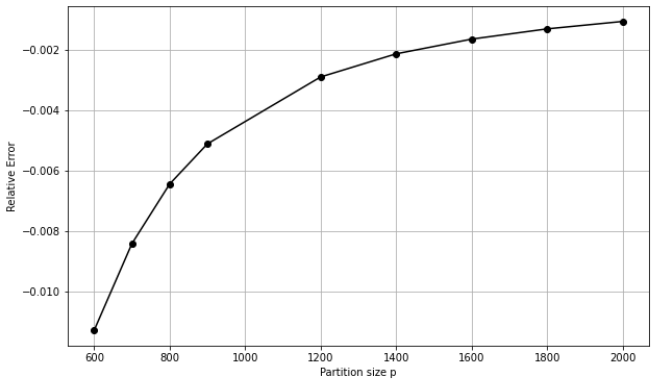

(e) OTM option - Long Maturity

Figure 2.7: Option prices in high volatility.

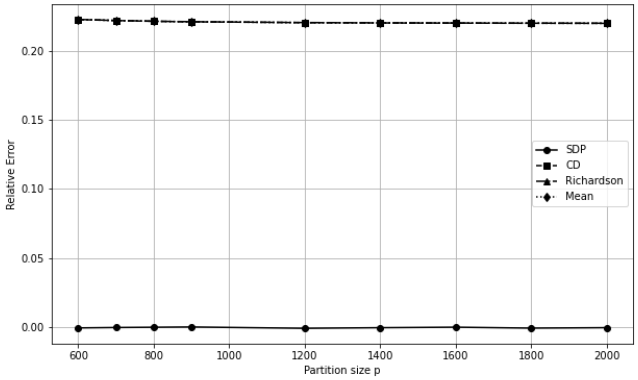

(a) ITM option $S_0 = 85$ - Short Maturity

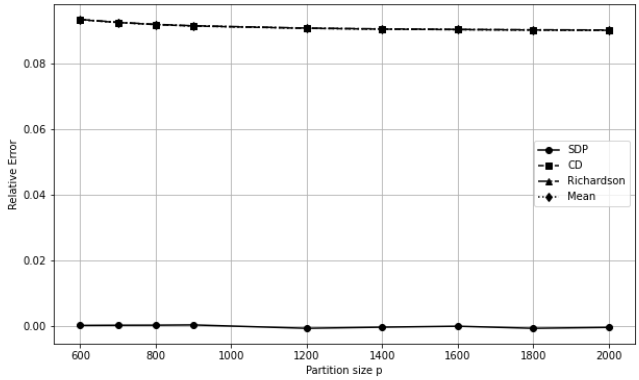

(b) ITM option - Medium Maturity

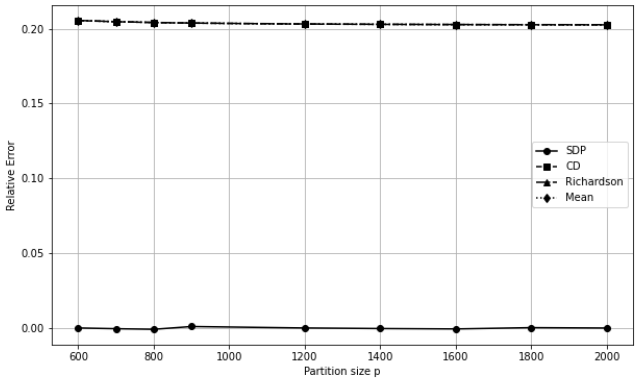

(c) ITM option $S_0 = 90$ - Short Maturity

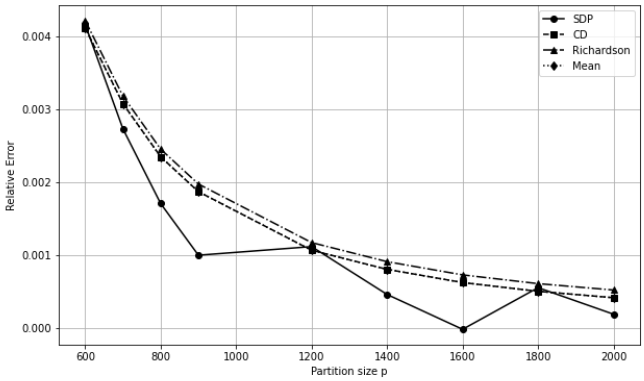

(d) ATM option - Long Maturity

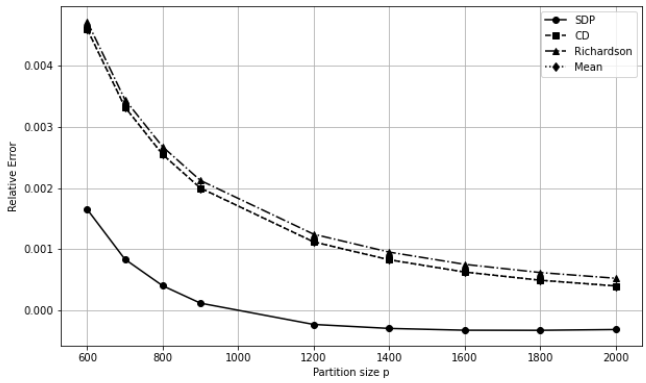

(e) OTM option - Long Maturity

Figure 2.8: Delta in high volatility.

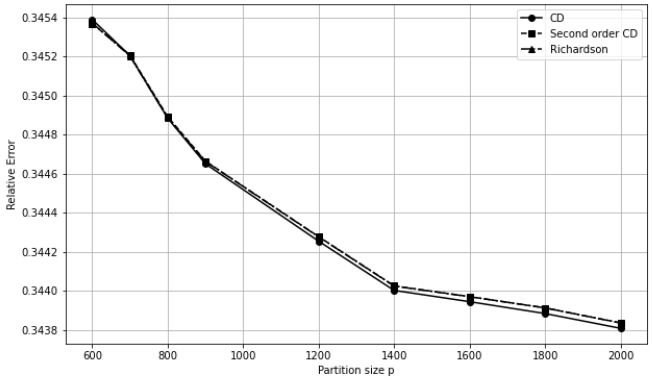

(a) ITM option $S_0 = 85$ - Short Maturity

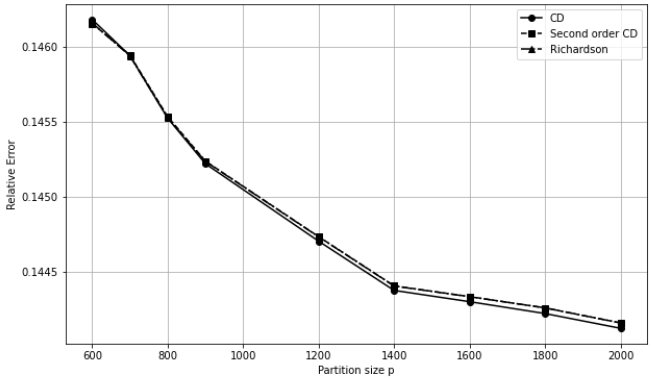

(b) ITM option - Medium Maturity

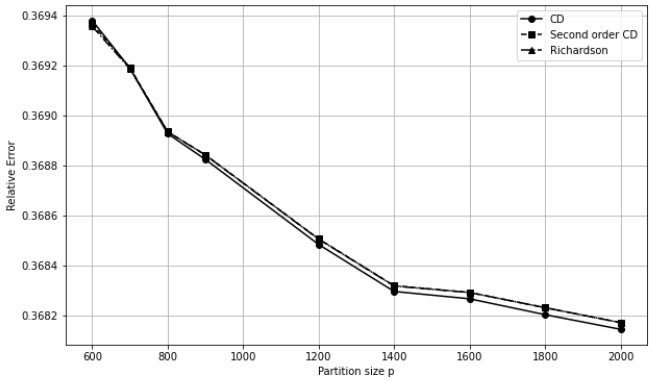

(c) ITM option $S_0 = 90$ - Short Maturity

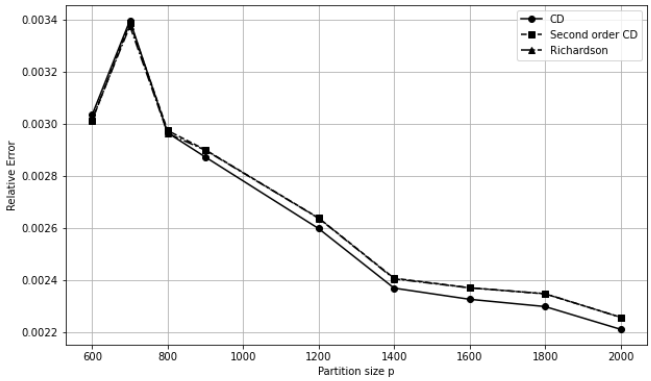

(d) ATM option - Long Maturity

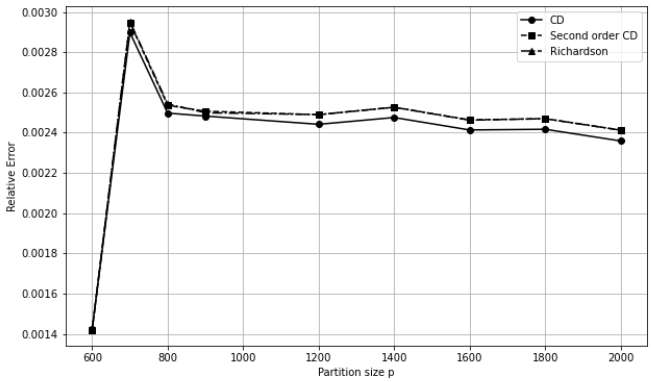

(e) OTM option - Long Maturity

Figure 2.9: Gamma in high volatility.

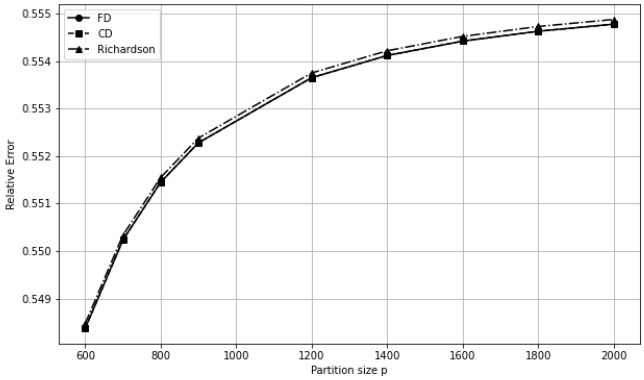

(a) ITM option $S_0 = 85$ - Short Maturity

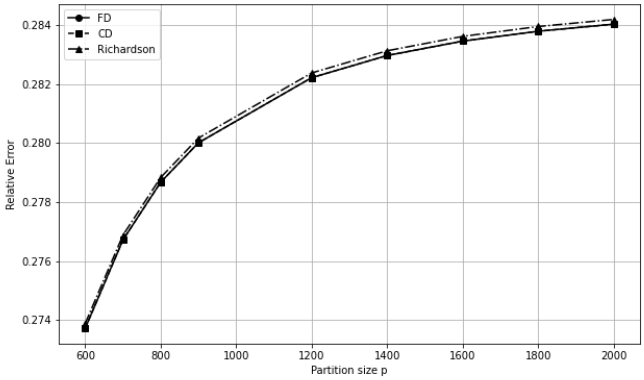

(b) ITM option - Medium Maturity

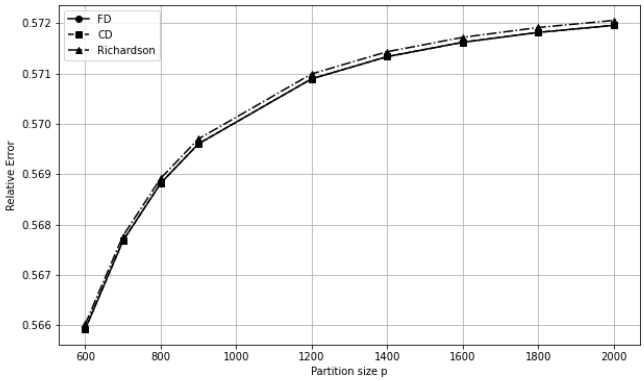

(c) ITM option $S_0 = 90$ - Short Maturity

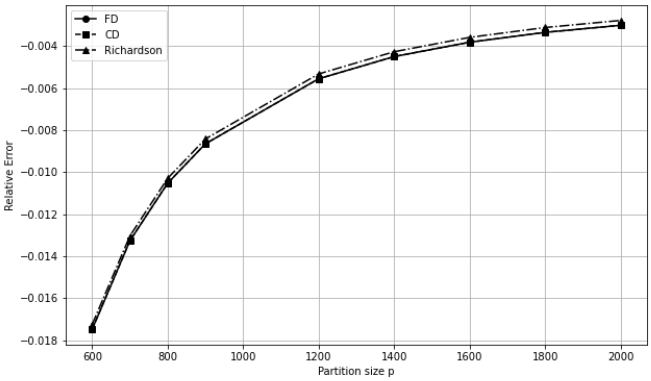

(d) ATM option - Long Maturity

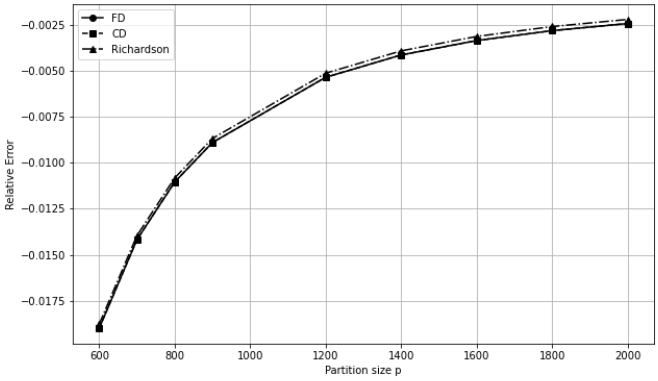

(e) OTM option - Long Maturity

Figure 2.10: Theta in high volatility.

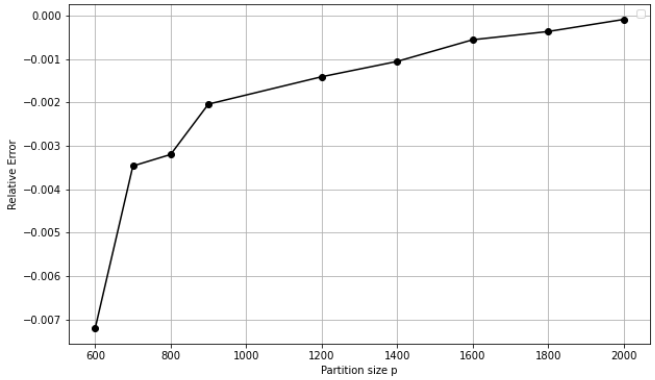

(a) ITM option $S_0 = 85$ - Short Maturity

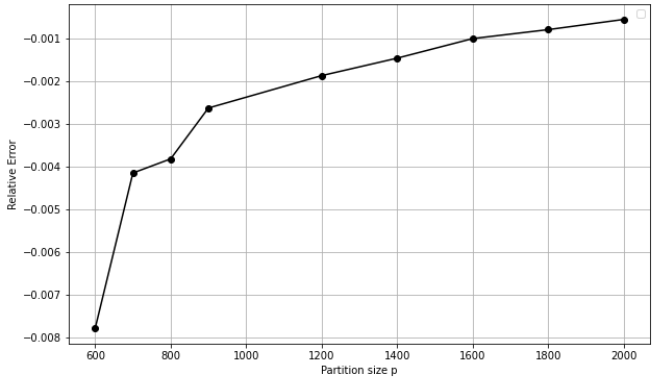

(b) ITM option - Medium Maturity

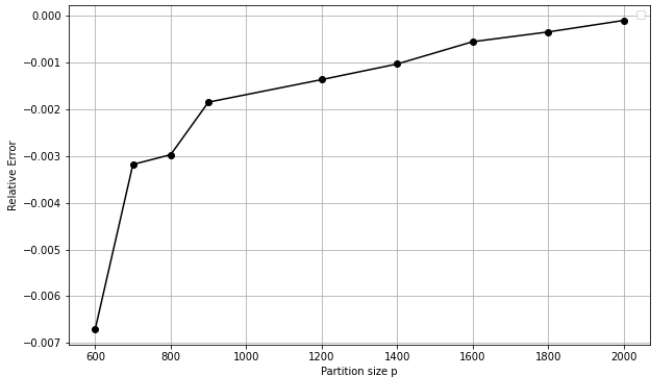

(c) ITM option $S_0 = 90$ - Short Maturity

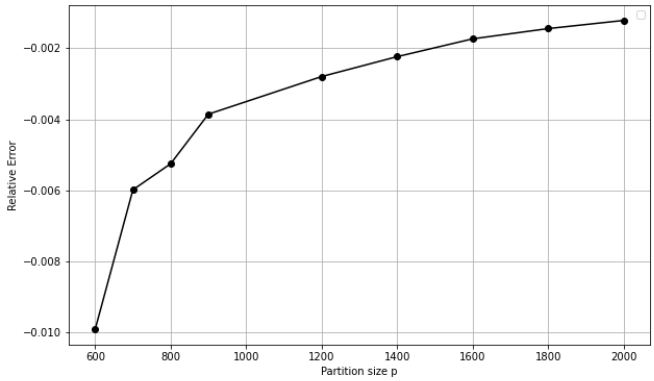

(d) ATM option - Long Maturity

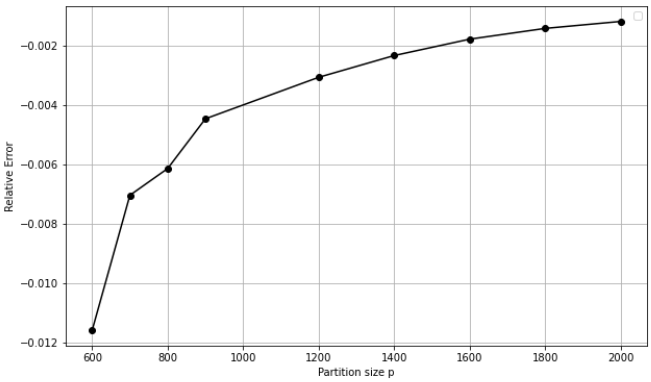

(e) OTM option - Long Maturity

Figure 2.11: Vega in high volatility.

2.3.3 Greeks Recommendations

Based on the analysis, we recommend numerical approaches for estimating Delta, Gamma and Theta under varying conditions of volatility, maturity, and initial asset price.

For Delta, in both low and high-volatility environments, the SDP method is generally the best to employ due to its computational efficiency, low relative error, and zero additional cost as it is already integrated within the algorithm. The Richardson method may surpass it in some low-volatility cases, particularly for ATM and OTM options, but the SDP method remains preferable.

For Gamma, in both low and high-volatility environments, the CD method is generally recommended for ATM and OTM options due to its efficiency and relatively low errors. However, for ITM options, Gamma is highly unstable across all methods because the option's payoff becomes nearly linear, especially for low-volatility cases, reducing sensitivity to changes in the underlying asset price. The Richardson method remains a competitive alternative in high volatility due to its lower computational time and robust performance, particularly for OTM and ITM options. When estimating Theta, Richardson's method is the most suitable in all environments, regardless of the option's maturity.

More sophisticated numerical approaches may offer marginal improvements in accuracy, but they often come with increased computational costs that may not be justified, especially in the case of Delta estimation for instance. In practice, the balance between efficiency and accuracy makes methods like SDP and CD preferable for most scenarios. It is essential to consider the specific conditions of volatility and option type when selecting the most appropriate method. The estimation of Vega reveals that in low-volatility environments, its sensitivity to changes becomes more pronounced, leading to greater variability in the estimation depending on the option type (ATM, OTM). This suggests that the specific characteristics of the option have a more significant impact on Vega under low volatility. In contrast, the impact of option characteristics diminishes in high-volatility environments, leading to a more uniform and stable estimation of Vega across different option types. This indicates that under such conditions, the variability in Vega's sensitivity is reduced, making its estimation less dependent on the nuances of the option's position relative to the underlying asset's price.

In summary, low volatility settings require more careful consideration of the option characteristics. Lastly, when addressing the impact of duration, it has a minimal effect on the convergence of option prices and Greeks, unlike the initial asset price, which significantly influences both prices and Greeks in low volatility settings.

Chapter 3

Calibration to Market Data

We now proceed to the calibration procedure where the primary objective is to estimate the volatility parameter σ in a way that is consistent with either market prices or the theoretical dynamics of the model. This is accomplished by adjusting the volatility parameter σ, which is assumed constant as per the Black and Scholes (1973) model. First, we use a direct optimization technique to minimize the Sum of Squared Errors (SSE) between market prices and the model's prices. Subsequently, we employ the Feynman-Kac Theorem in conjunction with an Ordinary Least Squares (OLS) regression to iteratively update volatility until convergence.

3.1 Data Selection

To support an effective calibration process, we began by acquiring options and security data for the Royal Bank of Canada (RBC). We choose RBC over other Canadian banks because of its strong presence in the United States markets, which makes it the largest investment bank in Canada. Options are sourced from the OptionMetrics database through WRDS. We selected the year 2017, which is the last available year in the WRDS dataset to conduct our analysis. We have specifically chosen to analyze prices on Wednesdays, as it is widely regarded as the most stable trading day of the week. Markets are not open over the weekend, and Mondays and Fridays are known for extreme events or reactions that can significantly impact trading. By choosing Wednesday, we aim to avoid the market fluctuations and potential distortions associated with the beginning and end of the trading week to ensure more stable data for analysis.

The data extraction procedure focuses on two key tables from the *optionm* library of the OptionMetrics Ivy-DB US dataset. The first table (*opprcd*) allows the extraction of option-related information including trade date, expiration date, strike price, bid and ask prices, and option type. This is filtered for the security identifier

of RBC and limited to trades recorded in 2017. The strike prices are standardized (*divided by* 1000), and the mid-price of each option is calculated by averaging the bid and ask prices. Simultaneously, the second table (*secprd*) is used to retrieve the closing prices of the underlying asset for the same security, filtered by the same year. The extracted data from these two tables is then merged based on the trade date to combine the option prices with the corresponding underlying asset prices. Finally, the time to maturity for each option is computed as the difference between the expiration date and the trade date. This merged dataset forms the foundation for calibrating option prices.

We eventually filtered the options to exclude those with less than thirty days to maturity, as are options that violate various no-arbitrage conditions. Additionally, options with extreme moneyness of less than 0.9 or more than 1.1 were omitted from the sample. These filtering rules follow Bakshi et al. (1997). Options violating no-arbitrage conditions are those where the option price V does not satisfy the lower bound condition $V \geq \max(0, S_0 - K)$ for call options, or $V \geq \max(0, K - S_0)$ for put options. This condition ensures that there are no arbitrage opportunities due to mispricing. We finally arrange the remaining 1354 option contracts, as in Christoffersen et al. (2009).

	$30 <$ DTM < 90	**$90 <$ DTM < 180**	**DTM > 180**	**Total**
$S_0/K < 0.975$	176	152	216	544
$0.975 < S_0/K < 1$	76	62	78	216
$1 < S_0/K < 1.025$	44	38	60	142
$1.025 < S_0/K < 1.05$	54	42	60	156
$1.05 < S_0/K < 1.075$	56	54	72	182
$S_0/K > 1.075$	32	32	50	114
Total	**438**	**380**	**536**	**1354**

Table 3.1: Number of options for calibration.

Table 3.1 summarizes our data set of 1354 contracts. Options are split up into three days-to-maturity (DTM) categories and six moneyness S_0/K categories. The choice of filtering options with a moneyness range between 0.975 and 1.05 is made to ensure a balanced selection of near-the-money options. By choosing a range close to 1, we focus on options that are neither too deeply ITM nor too deeply OTM. This range helps in maintaining liquidity and stability in option pricing, as deeply ITM or OTM options are generally more sensitive to changes in moneyness and can be less liquid. These options typically have better pricing stability and less erratic behavior compared to extreme moneyness cases (Christoffersen et al., 2009). Deep ITM or OTM options tend to have higher bid-ask spreads and lower trading volume, which

makes them less reliable for calibration and modeling purposes.

To initiate the calibration, we compute an annualized volatility estimate denoted as σ_0 . This initial guess is derived from daily log returns calculated on adjusted security prices, which are closing prices corrected for corporate actions such as dividends or stock splits. The daily log returns are then used to compute the return variance, from which the daily volatility is extracted and annualized with the approximate number of trading days in a year to obtain σ_0, the starting point for calibration.

We derive the dividend yield q based on the characteristics of dividend-paying securities as described in the OptionMetrics IvyDB US Reference Manual.[3] According to the manual, the dividend yield is estimated as the most recently announced dividend payment divided by the most recent closing price for the security

$$q = \frac{\text{Most Recent Dividend Payment} \times \text{Payment Frequency}}{\text{Current Security Price}},$$

under the assumption that the dividend yield remains constant over the remaining term of the option.

In the calibration process, the risk-free rate is not assumed to be a fixed constant across subsamples. It is dynamically determined through an interpolation of zero-coupon bond yields to match the maturity profile of the options being analyzed. The zero-coupon bond yields denoted as $r(t, T)$, are obtained from market data, where t represents the current date or the trade date of the option and T denotes the maturity of the bond. These yields are sourced from Optionmetrics financial databases, provided through WRDS.

We summarize the procedure as follows:

- Identify all unique trade dates $t = [t_1, t_2, \ldots, t_m]$ in the dataset. These dates represent distinct days on which options are traded for the first time.
- For each unique date t_j, the days to maturity $T - t_j$ are calculated for all options sharing this date.
- The zero-coupon bond data is retrieved for each unique trade date t_j, providing a set of risk-free rates $[r_1, r_2, \ldots, r_n]$ corresponding to maturities $[d_1, d_2, \ldots, d_n]$ where d_i represents the days to maturity of the bonds available in the market on the same release date t_j.

[3]OptionMetrics IvyDB US Reference Manual, 2023. Accessed via Wharton Research Data Services (WRDS): https://wrds-www.wharton.upenn.edu/.

On a given trade date, the zero-coupon bond data provides risk-free rates for specific maturities d_i, which might not align exactly with the maturities $T - t_j$ of the options. For this purpose, interpolation is needed to estimate the appropriate risk-free rate of the options.

The relationship can be visualized as follows, the maturities of the zero-coupon bond $[d_1, d_2, \ldots, d_n]$ are viewed as a grid of maturities for which risk-free rates $[r_1, r_2, \ldots, r_n]$ are known. The maturity of each option lies between two points d_i and d_{i+1} on this grid. Using a simple linear interpolation, the risk-free rate $r(T - t_j)$ for the option with $T - t_j$ days to maturity is estimated as

$$r(T - t_j) = r_i + \frac{(r_{i+1} - r_i)}{(d_{i+1} - d_i)} \cdot ((T - t_j) - d_i),$$

where r_i and r_{i+1} are the corresponding risk-free rates to the lower bond maturity d_i, and the upper bond maturity d_{i+1}, respectively. The interpolated rates are then assigned to all options with the same release date t_j. Given the substantial amount of available data, we found that higher-order interpolation methods, such as cubic interpolation, do not provide significant improvements in accuracy.

Through this process, each option in the dataset is assigned a risk-free rate based on its trade date t_j and its maturity $T - t_j$. Moreover, our analysis of each subsample confirmed that all options within a given subsample exhibit very similar risk-free rates. Consequently, we assigned the mean of these rates as the representative rate for each subsample.

3.2 Calibration Procedure

We now move to the calibration procedure where the primary goal is to align the theoretical model of SDP with the observed market data by adjusting the volatility parameter σ. The calibration starts with an initial volatility estimate σ_0, derived from historical price data. This volatility is then iteratively adjusted to minimize the SSE between the model's predicted option prices and the observed market prices, for each subsample in Table 3.1.

3.2.1 Numerical Optimization

We employ an optimization algorithm to iteratively update the volatility parameter σ. The approach minimizes the SSE between observed market option prices $\tilde{V}$ and theoretical option prices $\tilde{v}$ generated by the SDP. We define the objective function as

$$f(\sigma) = \sum_{i=1}^{n} \left(\tilde{V}_i - \tilde{v}_i(\sigma) \right)^2, \tag{3.2.1}$$

where $\tilde{V}_i$ represents the observed market price of the i-th option, $\tilde{v}_i(\sigma)$ is the theoretical price of the same option computed under the SDP model with volatility σ, and n is the number of options in the subsample. This objective function quantifies the total error between market and theoretical prices. The smaller the value of $f(\sigma)$, the closer the theoretical model is to replicate the observed market behavior. The optimization process seeks to find the value of σ that minimizes $f(\sigma)$, ensuring that the theoretical model is calibrated as closely as possible to market data.

The gradient of the objective function $\nabla f(\sigma)$ provides information about how $f(\sigma)$ changes with respect to the volatility

$$\nabla f(\sigma) = -2 \sum_{i=1}^{n} \left(\tilde{V}_i - \tilde{v}_i(\sigma) \right) \times \mathcal{V}_i(\sigma),$$

where $\mathcal{V}_i(\sigma)$ represents the sensitivity of the theoretical price to changes in volatility, or Vega, computed using Equation (1.4.20). The gradient helps the optimizer determine the direction and magnitude of the adjustment needed to reduce the objective function $f(\sigma)$. A zero gradient indicates that the optimization has reached a critical point, which, in this context, is the minimum of $f(\sigma)$.

The calibration utilizes the Limited-memory BFGS-B (L-BFGS-B) (Byrd et al., 1995) which is an optimization algorithm in the family of quasi-Newton methods that approximates the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method using a limited amount of computer memory. L-BFGS-B [4] is implemented in Python's `scipy.optimize.minimize` function and is suitable for the calibration due to its efficiency and ability to handle constraints. [5] This algorithm refines the initial estimate σ_0, to find the minimum of the differentiable function $f(\sigma)$. At each iteration k, it computes the gradient $\nabla f(\sigma_k)$, which indicates the direction of the steepest descent. Unlike standard quasi-Newton methods that store a full $n \times n$ approximation of the inverse Hessian matrix H_k, L-BFGS-B maintains a limited history of the most recent m updates of position differences and gradient differences $\sigma_{k+1} - \sigma_k$.

To evaluate the goodness of the optimization, we use the Root Mean Square

[4] Convergence is reached when either the gradient tolerance criterion is met, indicating a small gradient norm or the function tolerance criterion is satisfied, ensuring minimal relative change in the objective function between iterations.

[5] The optimization includes the constraint $\sigma > 10^{-3}$, because volatility must always be non-negative. This small positive constant is chosen to ensure numerical stability while remaining close to zero.

Error (RMSE) between the evaluated prices as a metric

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\tilde{V}_i - \tilde{v}_i)^2}, \tag{3.2.2}$$

where $\tilde{V}_i$ are the market values, $\tilde{v}_i$ are the predicted SDP values, and is the number of data points.

The following section introduces a new calibration method designed for option pricing. This approach aims to establish an estimable relationship between observed market option prices and theoretical Greeks computed using the SDP methodology. The method leverages the fundamental PDE governing option prices. We then compare its performance against the traditional numerical optimization approach.

3.2.2 Feynman-Kac Theorem

In this Section, we employ the Feynman-Kac (FK) Theorem to establish a functional relationship between the model's parameters and the observed data. Specifically, as in Biagini et al. (2008), any derivative product $\overline{V}(S,t)$ obeys the fundamental PDE under the Black and Scholes (1973) model

$$\frac{\partial \overline{V}(S,t)}{\partial t} + (r-q)S\frac{\partial \overline{V}(S,t)}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 \overline{V}(S,t)}{\partial S^2} = r\overline{V}(S,t). \tag{3.2.3}$$

By employing the FK Theorem, we can interpret the solution $\overline{V}(S,t)$ in Equation (3.2.3) as the expected value under the risk-neutral measure $\mathbb{Q}$ of the discounted payoff of the derivative product. Mathematically, the solution to Equation (3.2.3) is Equation (3.2.4).

Theorem 3.2.1. *(Oksendal, 2013)*

Let S_t follow the risk-neutral stochastic differential equation defined in Equation (1.1.1) and let $h(S)$ be a measurable payoff function satisfying regularity and growth conditions to ensure the expectation is finite. Then, the unique solution to Equation (3.2.3) is given by the Feynman-Kac representation

$$\overline{V}(S_t,t) = \sup_{\tau\in[t,T]} \mathbb{E}^{\mathbb{Q}}\left[e^{-r(\tau-t)}h(S_\tau) \mid S_t\right], \tag{3.2.4}$$

where $h(S_\tau)$ is the payoff function evaluated at the optimal stopping time τ and $\mathbb{E}^{\mathbb{Q}}$ denotes the expectation under the risk-neutral measure. This result provides a probabilistic interpretation of the PDE, which allows to express the derivative price as a conditional expectation under the probability measure $\mathbb{Q}$.

Sketch of Proof (Oksendal, 2013). We outline the intuition behind the Feynman-Kac representation. Starting from the PDE in Equation (3.2.3), we assume that the solution $\overline{V}(S_t, t)$ is sufficiently smooth.

By applying Itô's Lemma to the function $\overline{V}(S_t, t)$, where S_t follows the risk-neutral stochastic differential Equation

$$dS_t = (r - q)S_t\, dt + \sigma S_t\, dW_t^{\mathbb{Q}},$$

we obtain the dynamics of $\overline{V}(S_t, t)$. Substituting the PDE into Itô's formula cancels the drift term, leading to a representation where the discounted process $e^{-rt}\overline{V}(S_t, t)$ becomes a martingale.

Using the Optional Sampling Theorem for stopping times $\tau \in \mathcal{T}_{t,T}$ where $\mathcal{T}_{t,T}$ is the set of all stopping times that take values in the interval $[t, T]$, and assuming appropriate integrability and boundary conditions, we obtain

$$\overline{V}(S_t, t) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(\tau - t)}h(S_\tau) \mid S_t\right].$$

Maximizing over all stopping times $\tau \in \mathcal{T}_{t,T}$ yields the optimal stopping formulation in Equation (3.2.4). □

The PDE (3.2.3) can be rewritten as

$$\Theta_{\text{SDP}} + (r - q)S\Delta_{\text{SDP}} + \frac{1}{2}\sigma^2 S^2 \Gamma_{\text{SDP}} \approx r\tilde{V}, \tag{3.2.5}$$

where Θ_{SDP}, Δ_{SDP}, Γ_{SDP} are the calculated Greeks from SDP, $\tilde{V}$ are the market option prices and S denotes the underlying asset prices at which the options are evaluated. Equation (3.2.5) creates a link between the SDP theoretical option prices $\tilde{v}$ and the real market prices $\tilde{V}$ through the SDP estimated Greeks.

Ordinary Least Squares (OLS) is a classical statistical technique employed within linear regression modeling to determine the most suitable parameter estimates (Wooldridge, 2016). By adhering to the least squares principle, OLS seeks to identify parameter values that minimize the total sum of squared differences between the observed values of the dependent variable and the corresponding estimated values generated by a linear function of the independent variables.

In our context, one can define the variables for an OLS regression at each grid point (S_i, t_n) as

$$Y_i = r\tilde{V}_i - \Theta_{\text{SDP},i} - (r - q)S_i\Delta_{\text{SDP},i}, \tag{3.2.6}$$

$$X_i = \frac{1}{2}S_i^2\Gamma_{\text{SDP},i}, \tag{3.2.7}$$

where Y_i and X_i are the dependent and independent variables at asset price S_i and time t_n. The index i refers to the discretized asset price grid, and t_n is the fixed time level at which the regression is performed.

The relationship between Y and X in the OLS regression is considered linear and passes through the origin because it assumes a proportional relationship between the dependent variable Y and the independent variable X. In other words, the model assumes that changes in X produce consistent, proportional changes in Y, following a straight line. The linear model is written as $Y = \beta_1 \times X + \epsilon$, where ϵ represents the error term and β_1 corresponds to σ^2. In this regression model, Y is the dependent variable constructed from market prices and SDP-estimated Theta and Delta, while X is the independent variable built from Gamma, also estimated via SDP. The linear form follows directly from the Black-Scholes PDE structure.

Figure 3.1 visually illustrates the relationship between the dependent variable Y and the independent variable X, defined in Equations (3.2.7)–(3.2.6), for a selected subsample from the panel of options. Similar Figures for other subsamples are available in Appendix D, Figure D.1. The Figure consists of two panels: panel (a) illustrates the relationship using the initial volatility estimate explained in Section (3.1), while panel (b) shows the relationship after applying the final estimated volatility of the FK method. In many cases, we observe an improved fit after applying FK.

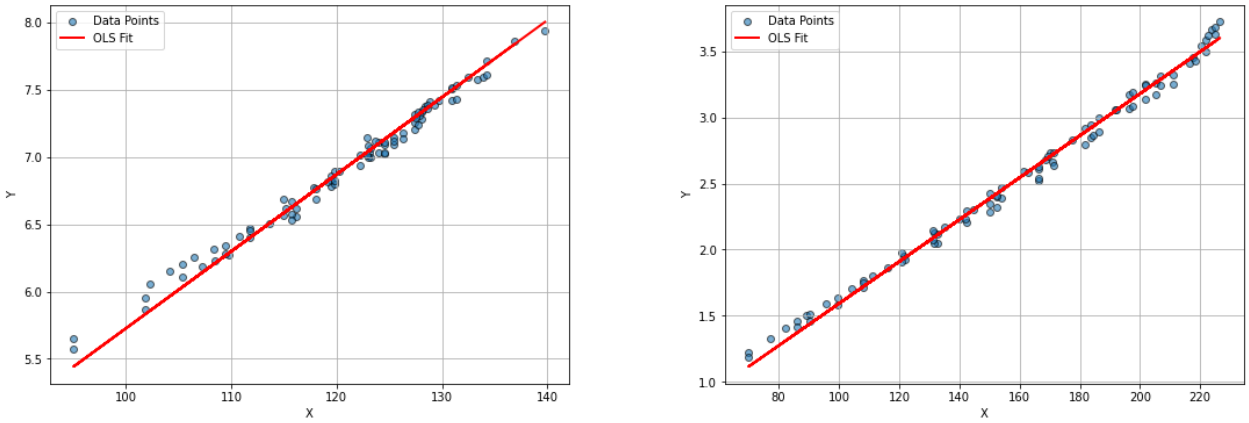

(a) Initial volatility estimate

(b) Final estimated volatility

Figure 3.1: OLS regression fits for $S_0/K < 0.975$ and $30 < \text{DTM} < 90$.

The OLS regression aims to estimate the coefficient β_1 that best fits the linear relationship between X and Y. Mathematically, we solve

$$\hat{\beta}_1 = \arg\min_{\beta_1} \sum_{i=1}^{n} (Y_i - \beta_1 X_i)^2,$$

and the volatility parameter σ appears in Equation (3.2.5) as $\frac{1}{2}\sigma^2$, so it is estimated as

$$\hat{\sigma}^{(k+1)} = \sqrt{\hat{\beta}_1^{(k)}},$$

where k is the current iteration.

To refine the estimate of σ, we employ an iterative optimization method using the L-BFGS-B algorithm, which updates $\hat{\sigma}$ until convergence. The procedure is as follows:

1. *Initialization:* Choose an initial guess $\sigma_0^{(k)}$ at $k = 0$.
2. *Compute Greeks:* Using SDP, calculate $\Theta_{\text{SDP}}, \Delta_{\text{SDP}}, \Gamma_{\text{SDP}}$ based on $\hat{\sigma}^{(k)}$.
3. *Perform OLS Regression:* Compute the dependent variable Y and independent variable X.
4. Estimate $\hat{\beta}_1^{(k)}$ using OLS.
5. *Update:*
$$\hat{\sigma}^{(k+1)} = \sqrt{\hat{\beta}_1^{(k)}}.$$
6. *Check Convergence:* If
$$|\hat{\sigma}^{(k+1)} - \hat{\sigma}^{(k)}| < 10^{-5}, \tag{3.2.8}$$
stop; otherwise, repeat from step 2.

This optimization minimizes the SSE

$$\mathcal{L}(\sigma) = \sum_{i=1}^{n} (Y_i - \hat{Y}_i)^2, \tag{3.2.9}$$

where

$$\hat{Y}_i = \frac{1}{2}(\hat{\sigma}^{(k)})^2 S_i^2 \Gamma_{i,\text{SDP}},$$

and

$$Y_i = r\tilde{V}_i - \Theta_{i,\text{SDP}} - (r - q) S_i \Delta_{i,\text{SDP}}.$$

The optimization process terminates when either Equation (3.2.8), or when the gradient tolerance or function tolerance criteria are met.

The L-BFGS-B method presented in Section (3.2.1) focuses on minimizing the SSE between market prices and model prices. In contrast, the combination of the FK Theorem with OLS regression aims to minimize the squared residuals within the regression itself. This distinction is crucial: rather than searching for the volatility that best matches observed option prices, the FK approach seeks the volatility that produces a set of Greeks satisfying the PDE in Equation (3.2.5).

This perspective shifts the focus from *implied volatility* to the concept of *implied Greeks*. The FK approach would likely offer better results when implementing a Delta-Gamma-Theta hedging strategy since it ensures internal consistency in the dynamics of option pricing and risk management. On the other hand, the L-BFGS-B approach is designed to minimize pricing discrepancies, leading to a more precise calibration of market prices. As such, each method has its strengths: L-BFGS-B is more suitable for pricing options, while the FK approach is better suited for risk management.

An important question arises: does minimizing pricing errors necessarily yield accurate estimates of the Greeks? The volatility that aligns the model prices with market prices may not correspond to the true underlying volatility that governs risk sensitivities, which can potentially lead to discrepancies in hedging strategies. To assess this, we will compare the FK and L-BFGS-B methods using 3D plots of implied volatility and implied Greeks. This will allow us to determine whether the FK method provides more reliable Greek estimates, a more meaningful implied volatility surface, or ultimately no significant advantage over the L-BFGS-B approach.

To evaluate the goodness of the FK optimization, we use the RMSE between the errors in the regression as a metric with the formula

$$\overline{\text{RMSE}} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(Y_i - \tilde{Y}_i)^2}, \tag{3.2.10}$$

where Y_i are the observed values, $\tilde{Y}_i$ are the predicted values, and n is the number of data points.

Table 3.3 reports the $\overline{\text{RMSE}}$ for samples defined in Table 3.1 reported both with the initial estimate as $\overline{\text{RMSE}}_0$ and with the final estimate as $\overline{\text{RMSE}}_{\text{FK}}$ computed with Equation (3.2.10).

Moneyness	DTM	$\overline{\text{RMSE}}_0$	$\overline{\text{RMSE}}_{\text{FK}}$
$S_0/K < 0.975$	$30 < \text{DTM} < 90$	0.611424	0.004721
	$90 < \text{DTM} < 180$	2.654845	0.014288
	$\text{DTM} > 180$	1.139653	0.030941
$0.975 < S_0/K < 1$	$30 < \text{DTM} < 90$	8.508569	0.332804
	$90 < \text{DTM} < 180$	10.45708	0.354525
	$\text{DTM} > 180$	3.598528	0.260309
$1 < S_0/K < 1.025$	$30 < \text{DTM} < 90$	12.80109	1.035917
	$90 < \text{DTM} < 180$	13.71561	0.698900
	$\text{DTM} > 180$	4.202216	0.509184
$1.025 < S_0/K < 1.05$	$30 < \text{DTM} < 90$	3.388763	0.354273
	$90 < \text{DTM} < 180$	5.727184	0.500429
	$\text{DTM} > 180$	2.894602	0.566090
$1.05 < S_0/K < 1.075$	$30 < \text{DTM} < 90$	0.415632	0.221672
	$90 < \text{DTM} < 180$	1.580094	0.283842
	$\text{DTM} > 180$	1.981228	0.378788
$S_0/K > 1.75$	$30 < \text{DTM} < 90$	0.413965	0.231664
	$90 < \text{DTM} < 180$	0.688452	0.259223
	$\text{DTM} > 180$	2.299326	0.335852

Table 3.3: RMSEs using FK.

As expected, we observe a significant reduction in $\overline{\text{RMSE}}$ from the initial estimate with $\overline{\text{RMSE}}_0$ to the final estimate with $\overline{\text{RMSE}}_{\text{FK}}$. This result is natural, as the calibration procedure is specifically designed to minimize the errors in the linear regression, which serve as the loss function. Specifically, we achieve a substantial improvement in fit as we see $\overline{\text{RMSE}}$ decreasing from 13.71561 to 0.698900.

In the following Section, we compare both calibration methods in terms of their ability to find the best price that matches the market price. This is done by comparing the RMSE values calculated using the optimal volatility obtained from each method. Since the L-BFGS-B approach explicitly minimizes pricing errors as a loss function, it is expected to achieve better accuracy in this regard. However, we aim to quantify how much accuracy is sacrificed when using the FK approach, which instead focuses on identifying Greeks that satisfy the fundamental PDE and minimize a different loss function.

Beyond this comparison, we focus on analyzing the performance of both calibration methods by evaluating their ability to identify the optimal volatility. This is done by examining how the RMSE evolves as a function of volatility. Additionally, we visualize the resulting volatility and Greek surfaces to better understand the

implications of each calibration approach on risk management and option pricing.

3.2.3 Comparison and Discussion of Results

This section analyzes the call options results of the local implied volatility (IV) for each subsample in Table 3.1 using the L-BFGS-B and FK methods. We use a partition size of $p = 1000$ to ensure fast results. Table 3.4 provides the local IV, and RMSE reported with the initial estimate σ_0 as RMSE_0, with the optimized volatility for the L-BFGS-B method as $\text{RMSE}_{\text{L-BFGS-B}}$ and with the optimized volatility for the FK method as RMSE_{FK} computed with Equation (3.2.2). $\text{IV}_{\text{L-BFGS-B}}$ represents the implied volatility obtained using the numerical optimization method of Section (3.2.1), while IV_{FK} corresponds to the IV computed using the FK Theorem of Section (3.2.2). Each subsample is calibrated within 4 to 6 iterations for each method and only takes a few minutes.

Table 3.4 suggests that the local IVs vary slightly across different subsamples for the L-BFGS-B method. These variations are expected due to differences in moneyness and maturity of the options. For instance, for $1.05 < S_0/K < 1.075$, the implied volatility for options with $30 < \text{DTM} < 90$ was estimated at 0.158129, while for those with $\text{DTM} > 180$ it was 0.115641. This suggests that volatility estimates tend to be higher for options with shorter maturities and lower for those with longer maturities, which aligns with market behavior, as shorter-term options are typically more sensitive to volatility.

For the L-BFGS-B method, the RMSE values show a significant reduction from using the initial approximation (RMSE_0) to the optimized parameter ($\text{RMSE}_{\text{L-BFGS-B}}$), which indicates the success of the optimization in minimizing the difference between theoretical and market prices. For example, for $S_0/K < 0.975$, the RMSE decreased from 0.861927 to 0.0878799 which reflects the effectiveness of the optimization algorithm. Similar improvements were observed for all other samples.

As shown in Table 3.4, the FK method prioritizes aligning Greeks with the PDE rather than directly minimizing pricing errors, which inherently creates a trade-off. Consequently, FK does not necessarily produce the best price fit. One key observation from Table 3.4 is that the local IV obtained through FK calibration is often lower than $\text{IV}_{\text{L-BFGS-B}}$. Since FK estimates volatility by regressing the dependent variable Y in Equation (3.2.6) on the independent variable X in Equation (3.2.7) based on Equation (3.2.5) instead of minimizing price errors, the resulting IV may not be optimal for price matching. This suggests that while the volatility structure from FK is more theoretically consistent, it may be less aligned with actual market prices.

When RMSE_{FK} is lower than RMSE_0 (as seen in some cases in Table 3.4), FK still provides a better fit than the initial estimate, despite not explicitly minimizing RMSE. This indicates that the volatility structure estimated by FK is at least partially aligned with observed prices. However, in cases where RMSE_{FK} remains close to or higher than RMSE_0, the FK method's volatility estimate does not necessarily bring model prices closer to market prices. As expected, this occurs because FK minimizes a different loss function, ensuring PDE consistency rather than directly reducing pricing errors. The discrepancy varies across moneyness levels, with no strict pattern of FK performing significantly worse in ITM or OTM cases.

Overall, Table 3.4 confirms that the L-BFGS-B method is the most effective at minimizing pricing errors. While FK ensures theoretical consistency through the PDE, it does not always lead to a lower RMSE compared to the initial estimate. The effectiveness of FK depends on whether the estimated volatility structure sufficiently aligns with market prices. If RMSE_{FK} is lower than RMSE_0, FK offers some improvement while maintaining PDE consistency. However, if RMSE_{FK} remains high, it suggests that enforcing the PDE comes at the cost of pricing accuracy.

Ultimately, this presents a trade-off: should we prioritize pricing accuracy with L-BFGS-B or theoretical consistency with FK? The choice depends on whether the primary objective is market fit or maintaining adherence to the fundamental PDE and prioritizing hedging and risk management.

Moneyness	DTM	$IV_{\text{L-BFGS-B}}$	IV_{FK}	$RMSE_0$	$RMSE_{\text{L-BFGS-B}}$	$RMSE_{\text{FK}}$
$S_0/K < 0.975$	$30 < \text{DTM} < 90$	0.125263	0.094692	0.861927	0.087899	0.332709
	$90 < \text{DTM} < 180$	0.123029	0.047697	1.806286	0.131569	0.887013
	$\text{DTM} > 180$	0.120716	0.060106	2.974491	0.201032	1.872191
$0.975 < S_0/K < 1$	$30 < \text{DTM} < 90$	0.130652	0.009985	1.269261	0.169377	1.237294
	$90 < \text{DTM} < 180$	0.125904	0.011313	2.114466	0.190820	2.158299
	$\text{DTM} > 180$	0.118018	0.007836	3.455149	0.323697	3.628298
$1 < S_0/K < 1.025$	$30 < \text{DTM} < 90$	0.122805	0.008186	1.225202	0.196294	1.220356
	$90 < \text{DTM} < 180$	0.124727	0.010043	1.936785	0.164998	2.133884
	$\text{DTM} > 180$	0.115590	0.024850	3.159726	0.233167	3.262074
$1.025 < S_0/K < 1.05$	$30 < \text{DTM} < 90$	0.146112	0.025085	0.940733	0.212798	0.841833
	$90 < \text{DTM} < 180$	0.132113	0.097436	1.798773	0.227970	1.635702
	$\text{DTM} > 180$	0.120720	0.010180	2.828025	0.278923	2.494748
$1.05 < S_0/K < 1.075$	$30 < \text{DTM} < 90$	0.158129	0.125159	0.666872	0.110897	0.538853
	$90 < \text{DTM} < 180$	0.131396	0.037183	1.639580	0.231945	1.147482
	$\text{DTM} > 180$	0.115641	0.018386	3.012915	0.320401	2.208325
$S_0/K > 1.075$	$30 < \text{DTM} < 90$	0.177039	0.107390	0.457839	0.239156	0.478474
	$90 < \text{DTM} < 180$	0.127749	0.077977	1.398258	0.283234	0.764219
	$\text{DTM} > 180$	0.105171	0.027370	2.745839	0.289531	1.435964

Table 3.4: IVs and RMSEs.

To further compare the validity of the results, we select a range of volatility values and compute the RMSE for each method with Equations (3.2.2)-(3.2.10). By examining the shape of the RMSE curves, we can assess the ability of each method to accurately find the optimal volatility.

Figure 3.2 presents the RMSE for both calibration methods applied to a selected subsample. Similar Figures for other subsamples are available in Appendix D, Figures D.2, D.3. The Figure is divided into two panels: panel (a) displays the RMSE obtained using the L-BFGS-B method as $\text{RMSE}_{\text{L-BFGS-B}}$ computed with Equation (3.2.2), while panel (b) shows the RMSE corresponding to the FK approach as RMSE_{FK} computed with Equation (3.2.10).

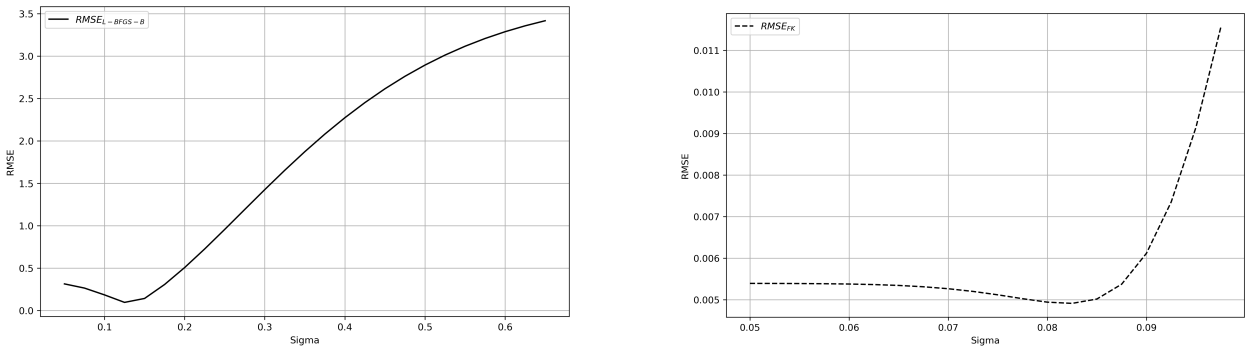

(a) L-BFGS-B (b) FK

Figure 3.2: RMSE across volatility values for $S_0/K < 0.975$ and $30 < \text{DTM} < 90$.

Following Figure 3.2, the RMSE curve for the L-BFGS-B method is convex, with a clear global minimum. This behavior aligns with the theoretical expectations of numerical optimization, which confirms the method's ability to minimize the SSE and reliably converge to an optimal volatility value. Similarly, the FK method also exhibits a well-defined global minimum, which closely matches the optimized local implied volatility reported in Table 3.3.

This consistency across both methods reinforces their reliability in volatility estimation. However, it is important to emphasize that while both approaches provide stable results, they do not serve the same purpose. This distinction motivates the next step in our analysis: comparing the IV surfaces obtained from each method to assess whether one provides a more meaningful representation of market conditions. Additionally, we extend this comparison to the Greeks surfaces to evaluate how each calibration approach impacts risk management and hedging strategies.

Figure 3.3 illustrates the volatility surfaces obtained with each method. Panel (a) corresponds to the L-BFGS-B approach of Section (3.2.1), while panel (b) represents the IV surface using the FK method of Section (3.2.2).

This comparison reveals that panel (b) exhibits a smoother and more continuous peak, while panel (a) features a sharper, more pronounced peak. The surface in panel (b) extends naturally towards the edges, with a stable transition, whereas panel (a) shows a steeper decay at the left edge.

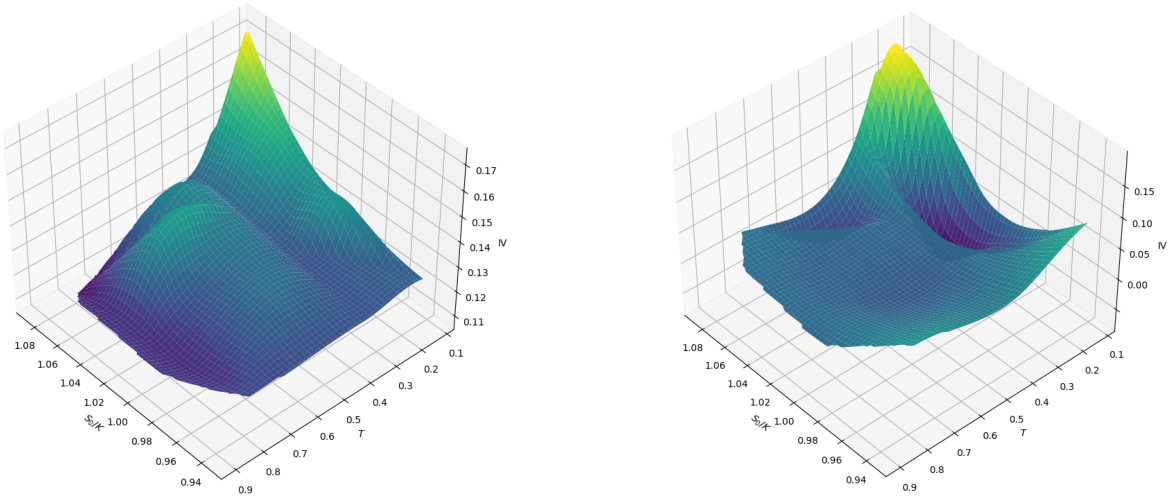

(a) L-BFGS-B (b) FK

Figure 3.3: Implied volatility surfaces.

Examining the IV surfaces in Figure 3.3, we notice two key observations. Panel (b) exhibits a more pronounced volatility smile than panel (a). Specifically, for FK, IV is lowest near $S_0/K = 1$ and increases for both ITM $S_0/K < 1$ and OTM $S_0/K > 1$ regions. This is because FK estimates volatility based on PDE consistency, which can naturally introduce a stronger smile-looking structure. This effect is reinforced by Gamma, which is highest for ATM options, increasing price sensitivity to volatility and requiring frequent hedging, thereby driving demand. However, high liquidity and continuous hedging in ATM options smooth out price fluctuations and provide lower implied volatility. Vega further amplifies the impact by increasing option price reactions to volatility changes.

Another important observation in both IV surfaces is the plateau (or flat region) around $T = 0.2$, where IV appears relatively stable before fluctuating again. This could be due to several factors. First, it may reflect market behavior, where certain

maturities are more actively traded, which leads to more stable IV estimates. Another possibility is that the plateau results from data limitations, where fewer options exist beyond certain maturities, which provides a reduced variation in the estimated IVs.

Overall, the FK approach does not appear to produce a noisy or poorly structured volatility surface when compared to the L-BFGS-B method, where we minimize the SSE between prices. This is precisely what we aimed to verify, as existing literature already supports the idea that minimizing SSE between prices yields well-behaved implied volatility surfaces. Our findings further confirm that a calibration method based on an iterated regression of the FK PDE can also provide a meaningful volatility surface.

In addition, we generate 3D plots to visualize the Greeks surfaces computed using the initial estimate σ_0, the volatilities obtained from the L-BFGS-B method (reported in Table 3.4) and the volatilities of the FK method (reported in Table 3.3). These visualizations enable a direct comparison of each method's behavior and help assess both their accuracy and stability in estimating Greeks under varying market conditions. We present the implied Delta surface, while the implied Gamma and Theta surfaces are provided in Figures D.4 and D.5, respectively, in Appendix D.

Figure 3.4 provides the implied Delta surfaces using the initial volatility approximation defined in Section (3.1), as well as the optimal IV obtained with the L-BFGS-B method and the FK method. All three surfaces exhibit a similar structure, with a smooth gradient from lower values in purple to higher values in yellow. The implied Delta increases as S_0/K rises and as T decreases, which is an expected behavior in option pricing.

Panel (a) can serve as a reference for evaluating the two calibration approaches since it is computed with an initial volatility guess σ_0. Panel (b) presents a slightly sharper curvature along the S_0/K axis, with more pronounced ridges, compared to panel (a). Panel (c) exhibits an even stronger curvature along both axes, with steeper variations in Delta values across different regions, which can potentially make it more sensitive to shifts in market parameters. These differences in curvature have important implications for hedging strategies. Since panel (c), produces more localized variations in Delta, it may offer better precision for dynamic hedging, which allows for a more responsive adjustment to small changes in the underlying asset. In contrast, panel (b) provides a smoother Delta surface, which could lead to more stable hedge ratios across different maturities and moneyness levels. A smoother Delta estimate can reduce sensitivity to market noise, but it might not capture local fluctuations as effectively as the FK method.

Ultimately, if the goal is to obtain a stable Delta estimate that performs consistently across a wide range of conditions, the L-BFGS-B approach could be preferable. However, if the objective is to capture small variations in Delta for a

more precise hedging strategy, the FK method appears to be a better choice.

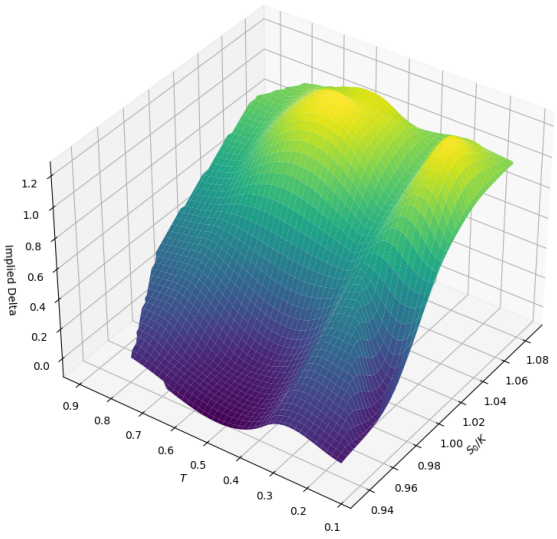

(a) Initial estimate

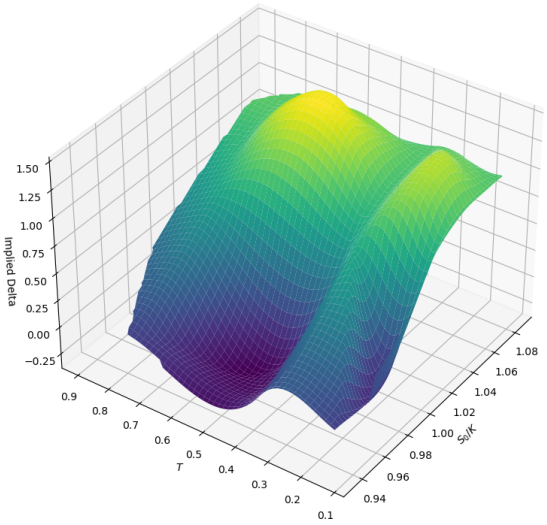

(b) L-BFGS-B

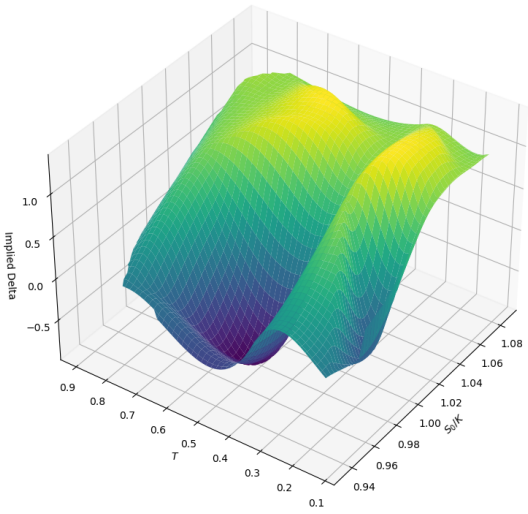

(c) FK

Figure 3.4: Implied Delta surfaces.

To build on our findings regarding the behavior of the FK and L-BFGS-B calibration methods, it is crucial to further assess the impact of market conditions on the stability of volatility estimation. While our results confirm that the FK approach does not produce a noisy or poorly structured volatility surface compared to the L-BFGS-B method, it is important to consider potential sources of variability in financial data.

One key factor is the presence of heteroskedasticity, where the variance of errors changes with the level of the independent variables. To formally assess this, we applied the Breusch-Pagan (BP) test, which examines whether the variance of residuals systematically depends on the independent variables. The test results outlined in Table D.3 of Appendix D provide strong evidence of heteroskedasticity, indicating that the error variance is not constant across the dataset.

Given the presence of heteroskedasticity, it is important to further investigate its nature to understand how the error variance evolves with respect to volatility. Specifically, we seek to determine whether the variance follows a linear, exponential, or more complex pattern and whether there exist effects where the variance exhibits sudden changes. To explore this, we examine the relationship between the squared residuals $(Y - \hat{Y})^2$ and Vega $\mathcal{V}$ for a selected sample of the dataset with $S_0/K < 0.975 - 30 < \text{DTM} < 90$.

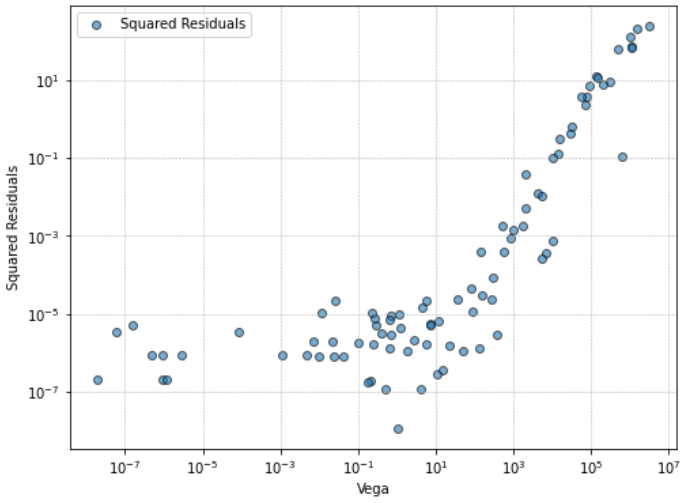

Figure 3.5: Vega vs squared residuals.

Figure 3.5 suggests that the relationship between Vega and error variance is non-linear. For low Vega values, the variance remains relatively stable. However, for larger Vega values ($> 10^1$), the error variance increases sharply, following what appears to be an exponential growth pattern. This suggests the presence of a critical

threshold around $\mathcal{V} > 10^1$, beyond which the model becomes significantly less precise.

To formally confirm this trend, we performed a log-log linear regression between the log of squared errors and the log of Vega, which helps linearize potential exponential relationships. The results indicate a significant regression coefficient, which confirms that error variance increases systematically with Vega. Specifically, the model achieves an R^2 value of 0.647, meaning that 64.7% of the variance in errors is explained by Vega, which confirms the strong relationship between the two variables. Furthermore, the p-value of the regression coefficient is 3.90×10^{-21}, which is extremely low, indicating that the effect of Vega on the error variance is highly significant.

We alternatively apply a spline regression to assess whether it provides a better fit. A spline is a piecewise polynomial function joined at specific points called knots, allowing for flexibility while avoiding the instability of high-degree polynomial regressions. Since the relationship between Vega and squared errors appears nonlinear, we compare a standard linear regression to a spline model to determine which better captures the underlying pattern. A spline regression adapts to local variations without introducing excessive instability, and its results are also all significative, with $R^2 = 0.92$. Figure 3.6 illustrates these two fits.

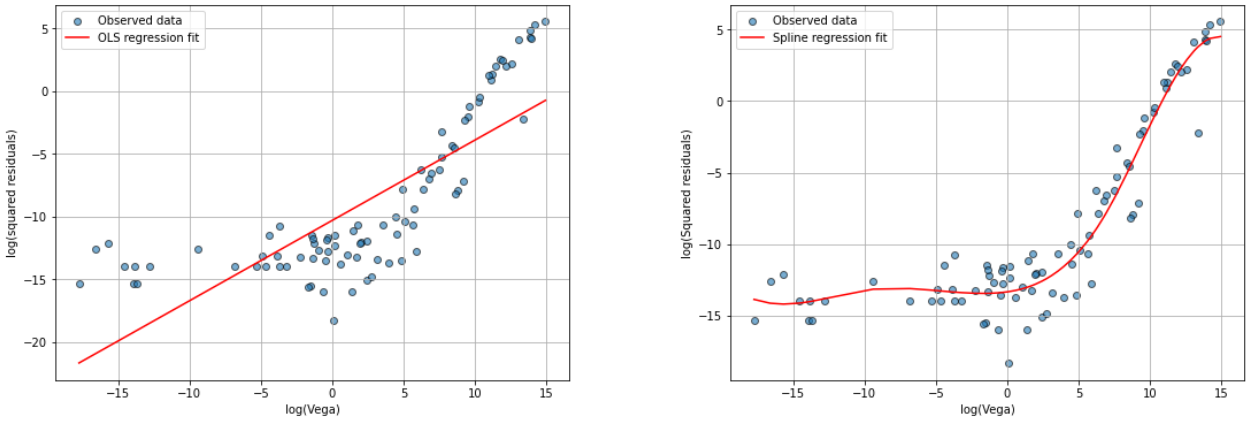

(a) Linear (b) Spline

Figure 3.6: Comparison of linear and spline regressions for residual variance.

Panel (a) shows that for large values of log(Vega), some points deviate significantly from the fitted line. This suggests that the relationship may not be strictly linear and that a more flexible model or transformation could be necessary.

Panel (b) demonstrates that spline regression better captures the overall trend compared to the linear model. Unlike a single global fit, splines adapt to local variations, providing a more accurate representation of the nonlinear relationship between Vega and error variance. This approach can be particularly useful for adjusting weights in a Weighted Least Squares (WLS) regression based on the estimated variance structure.

Our current results suggest that the current model underestimates errors for options with high Vega, which reinforces the need for an appropriate correction strategy. These findings are crucial for addressing heteroskedasticity through WLS, where the weights could be defined as $w = \frac{1}{\hat{\mathbb{V}}(\epsilon|\mathcal{V})}$, where the estimated variance $\hat{\mathbb{V}}(\epsilon|\mathcal{V})$ is obtained by a spline regression of log squared residuals ϵ of the initial model on log Vega $\mathcal{V}$.

We compute the RMSE for both the FK and L-BFGS-B methods across different values of volatility to assess whether the optimal volatility estimate changes when applying weighted adjustments to account for heteroskedasticity. Consequently, RMSEs are no longer computed with Equations (3.2.2) and (3.2.10). Instead, they are evaluated using their weighted formulations which incorporate the variance structure estimated through spline regression.

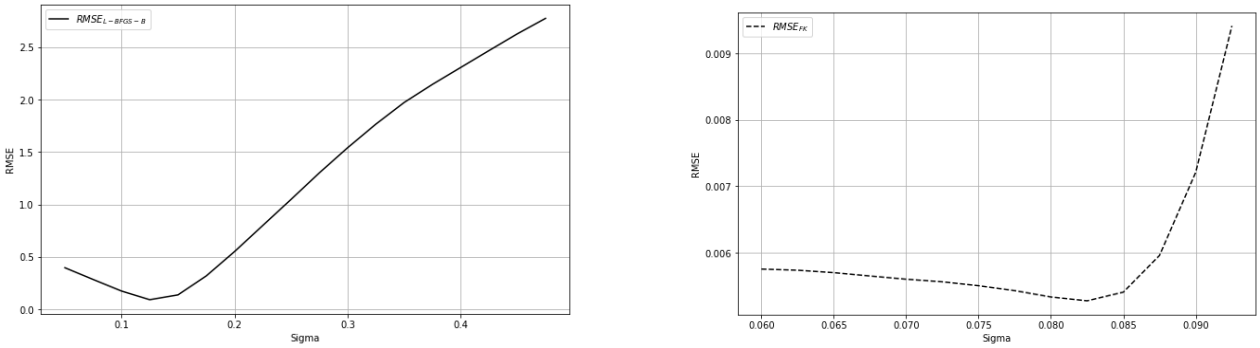

(a) L-BFGS-B (b) FK

Figure 3.7: RMSE across volatility values with WLS for $S_0 < 0.975$ and $30 < \text{DTM} < 90$.

We observe that both RMSE curves remain unchanged after applying the weights. This occurs when the WLS weights do not significantly impact the influence of observations or when the heteroskedasticity pattern does not strongly affect the relationship in the model. Although OLS estimates remain unbiased, WLS increases efficiency by improving the reliability of standard errors. Even if the estimated

coefficient stays the same, WLS still helps by refining inference through better standard error estimation.

When implementing the WLS approach for calibration, small Vegas yields a numerically ill-conditioned regression under various conditional heteroskedasticity specifications. Specifically, they cause a systematic floating point overflow of regression weights, which compromises the accurate computation of the inverse covariance matrix and results in an ill-conditioned regression problem. Overall, although promising, this WLS method requires careful selection of weights and the use of numerically stable techniques such as sparse matrix operations. Furthermore, improvements in calibration accuracy may not be substantial enough to justify the added computational cost. Altogether, these questions go beyond the scope of this thesis and are left for future research.

Another possible approach to mitigate the impact of market noise and heteroskedasticity is to use estimation techniques that are designed for non-constant error variance, such as Generalized Least Squares (GLS) or robust regression methods. GLS (Aitken, 1936) transforms the regression equation using an estimated covariance matrix to stabilize the variance of errors and ensure that observations with higher variance do not disproportionately affect the results. Therefore, it produces more efficient and reliable parameter estimates when heteroskedasticity is present. Methods such as Huber regression (Huber, 1992) and Tukey's bisquare regression (Tukey et al., 1977) reduce the influence of extreme residuals by assigning lower weights to observations with unusually large residuals. This makes the estimation more robust to market fluctuations and variability.

On the other side, the Black and Scholes (1973) model, which despite its foundational role in option pricing, has several notable limitations. The model assumes that volatility remains constant throughout the life of the option, an assumption that does not hold in real financial markets where volatility is stochastic and varies over time. These fluctuations significantly impact the pricing and optimal exercise strategy of American options since these options can be exercised at any time before expiration and their value is highly sensitive to volatility dynamics.

Furthermore, the Black and Scholes (1973) model relies on a single volatility parameter to estimate option prices, which oversimplifies market dynamics. The model assumes that asset prices follow a continuous GBM, ignoring the presence of sudden jumps in asset prices caused by market events. Incorporating discrete jumps into the model can provide greater flexibility and better capture real-world price movements and market behavior. More advanced models, such as stochastic volatility models (Ball and Roma, 1994) that incorporate stochastic volatility or jump-diffusion

models (Kou, 2002) that consider asset price jumps offer greater flexibility as they incorporate additional parameters and can better adjust to market conditions and provide a better fit to observed market data to construct accurate risk management strategies.

Summary and Future Directions

This thesis focuses on pricing a panel of American-style options and the estimation of Greeks in a way that addresses the computational challenges inherent to these tasks and provides solutions that improve both efficiency and accuracy. This was achieved through the combination of a Stochastic Dynamic Program (Ben-Ameur et al., 2002) with the homogeneity property (Létourneau and Stentoft, 2022).

On one hand, SDP allows for the determination of the option price by focusing its evaluation on specific points in time where exercise is allowed. This involves discretizing time only at these points, significantly reducing computational effort. Additionally, SDP calculates the transition matrices, which model the dynamics of the underlying asset's price movements, only once which is the most computationally intensive task in the entire procedure. This program also simplifies the calculation of Greeks as they are either integrated within the algorithm or computed using numerical differentiation which requires minimal additional computational effort.

On the other hand, leveraging the homogeneity property significantly increases the efficiency of pricing a panel of options. This property ensures that the early exercise strategies depend only on current values and strike prices, not on the initial state variables. As a result, the same optimal exercise strategy can be applied across all options in the panel. This eliminates the need for individual recalculations, as the precomputed SDP grid can be reused for multiple options with varying maturities and moneyness levels. This means that pricing a panel of options with SDP incurs the same cost as pricing a single option, regardless of the number of options in the panel. Remarkable!

The proposed combination demonstrates accuracy while significantly reducing computational time, resulting in greater efficiency compared to traditional methods in pricing a panel of options and estimating their Greeks under varying market conditions.

A key contribution of this work is the study of the convergence properties of the Greeks computed within the SDP framework, demonstrating the robustness of sensitivity estimates. Additionally, we developed a new calibration method based on an iterated regression of the Feynman-Kac PDE, offering an alternative to traditional approaches by leveraging the relationship between option prices and their sensitivities rather than minimizing price discrepancies.

The SDP, while efficient for low-dimensional problems, introduces approximation errors due to the discretization of the state space. Although increasing the grid resolution can improve accuracy with minimal additional cost in one or two dimensions, the method becomes computationally intensive in higher-dimensional settings due to the curse of dimensionality. A notable computational bottleneck is the construction of transition matrices, which remains the most time-consuming part of the algorithm. Identifying faster methods (Cosma et al., 2020) for generating these matrices represents a key area for potential improvement.

Similarly, the FK calibration technique presented in this thesis is subject to limitations stemming from its reliance on the underlying model. Its effectiveness is highly dependent on the degree to which the model reflects real market dynamics. The single-regressor structure, as in the Black-Scholes PDE, may not sufficiently capture the complexity of observed option prices. This factor can limit the accuracy and robustness of the calibration, especially in the presence of market noise. Addressing these limitations would likely require richer models, such as those with multiple sources of randomness or time-varying coefficients, rather than adjustments to the calibration technique itself.

Our analysis confirmed the presence of heteroskedasticity, as indicated by the BP test, meaning that the OLS estimator is no longer BLUE (Best Linear Unbiased Estimator). This suggests that the FK calibration method, which relies on OLS, may be affected by non-constant error variance. Future research could explore alternative regression techniques, such as WLS using the proposed weighting scheme in this thesis to increase accuracy and stability. Additionally, more robust regression methods could be investigated to further improve model reliability in the presence of heteroskedasticity and market noise.

Another promising direction would be to explore the implications of the FK calibration approach in hedging strategies. Unlike traditional optimization methods that minimize pricing errors, FK calibration ensures internal consistency in sensitivity estimates. This could lead to greater stability in Greeks for dynamic hedging, particularly for portfolios exposed to Delta and Gamma variations. Investigating its effectiveness in reducing hedging errors and stabilizing portfolio adjustments would provide valuable insights for risk and option market practitioners.

On a broader level, while this research was conducted within the Black and Scholes (1973) model, which provides a solid theoretical foundation, its assumptions introduce well-known limitations. The Black and Scholes (1973) assumes constant volatility and relies on a single parameter to estimate option prices. This simplicity limits its flexibility in capturing real-world dynamics. More advanced models that incorporate stochastic volatility or asset price jumps can offer greater flexibility as they incorporate additional parameters and can provide a better fit to observed market data to construct accurate risk management strategies.

In conclusion, while the Black and Scholes (1973) model provided a strong foundation for this research and a solid starting point, extending the framework to incorporate more sophisticated models presents a promising avenue for further research. Future work could explore these extensions to build upon the computational efficiencies achieved in this thesis.

Appendix A

Appendix for Chapter 1

Table A.1 reports the Delta estimates for the methods presented in Section (1.4.3) for different partition sizes p as well as the benchmark from Section (1.3). The underlying asset values S are interpolated linearly from the finite grid G. Delta estimates are obtained as follows: Δ_{CD} using Equation (1.4.9), Δ_{Rich} with Equation (1.4.10), Δ_{mean} based on Figure (1.2), and Δ_{SDP} directly from the SDP. The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

S	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Δ_{SDP}	$\Delta_{\text{Benchmark}}$
			$p = 400$		
91.12	-1.034300	-0.832039	-0.975204	-0.948226	-0.962473
93.39	-0.818788	-0.839937	-0.765252	-0.770038	-0.800586
95.65	-0.624289	-0.505041	-0.588930	-0.601985	-0.611717
97.91	-0.472340	-0.468573	-0.449424	-0.473003	-0.467114
100.17	-0.374794	-0.373642	-0.371835	-0.365738	-0.362310
102.44	-0.279469	-0.357545	-0.268921	-0.283141	-0.277662
104.70	-0.233948	-0.232470	-0.225941	-0.218779	-0.213754
106.96	-0.172999	-0.172106	-0.167536	-0.167970	-0.162873
109.22	-0.130677	-0.130106	-0.126797	-0.129513	-0.126179
111.49	-0.105735	-0.094287	-0.102736	-0.099891	-0.094681
			$p = 1200$		
91.12	-0.990591	-1.001834	-0.972792	-0.967316	-0.962473
93.39	-0.822746	-0.825754	-0.803922	-0.797155	-0.800586
95.65	-0.629328	-0.628902	-0.616788	-0.612465	-0.611717
97.91	-0.470931	-0.445257	-0.462975	-0.471939	-0.467114
100.17	-0.368232	-0.383246	-0.362823	-0.363690	-0.362310
102.44	-0.282095	-0.280137	-0.278426	-0.277915	-0.277662
104.70	-0.214769	-0.214753	-0.212242	-0.213262	-0.213754
106.96	-0.165891	-0.165007	-0.164091	-0.163547	-0.162873
109.22	-0.127005	-0.120595	-0.125711	-0.124812	-0.126179
111.49	-0.096018	-0.092169	-0.095086	-0.095325	-0.094681

Table A.1: Delta for $S \in [90, 110]$ for different partition sizes p.

Table A.2 reports the Gamma estimates for the methods presented in Section (1.4.4) for different partition sizes p as well as the benchmark from Section (1.3). The underlying asset values S are interpolated linearly from the finite grid G. Gamma estimates are obtained as follows: Γ using Equation (1.4.13), $\tilde{\Gamma}$ with Equation (1.4.14) and Γ_{Rich} with Equation (1.4.15). The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

S	Γ	$\tilde{\Gamma}$	Γ_{Rich}	$\Gamma_{\text{Benchmark}}$
		$p = 400$		
91.12	-0.099090	-0.152152	0.172110	0.039858
93.39	0.139612	0.140079	1.019616	0.089746

S	Γ	$\tilde{\Gamma}$	Γ_{Rich}	$\Gamma_{\text{Benchmark}}$
95.65	0.212781	0.250186	0.636731	0.072772
97.91	0.070843	0.075322	0.070480	0.055350
100.17	0.042377	0.042346	0.042290	0.040552
102.44	0.037761	0.040159	0.053044	0.032988
104.70	0.032644	0.035743	0.032526	0.025251
106.96	0.022233	0.022170	0.022171	0.019584
109.22	0.015771	0.015735	0.015735	0.014644
111.49	0.000480	-0.002407	0.015211	0.011588
		$p = 1200$		
91.12	-0.038371	-0.056629	-0.244460	0.039858
93.39	0.143697	0.160356	0.437571	0.089746
95.65	0.107393	0.117618	0.304188	0.072772
97.91	0.048711	0.045758	0.077424	0.055350
100.17	0.037845	0.037823	-0.030370	0.040552
102.44	0.040550	0.042621	0.040382	0.032988
104.70	0.029181	0.029675	0.078907	0.025251
106.96	0.021682	0.022369	0.021619	0.019584
109.22	0.021061	0.022759	0.013224	0.014644
111.49	0.007550	0.006212	0.012169	0.011588
		$p = 2000$		
91.12	0.067022	0.071292	0.364735	0.039858
93.39	0.104856	0.10798	0.142337	0.089746
95.65	0.104718	0.113216	0.103969	0.072772
97.91	0.062004	0.063568	0.085416	0.055350
100.17	0.033348	0.030296	0.061923	0.040552
102.44	0.032449	0.032382	0.003202	0.032988
104.70	0.034911	0.037651	0.058045	0.025251
106.96	0.01691	0.016445	0.023102	0.019584
109.22	0.018617	0.019614	0.031398	0.014644
111.49	0.011515	0.011338	0.023972	0.011588
		$p = 3000$		
91.12	0.050109	0.052771	0.130657	0.039858
93.39	0.095626	0.098213	0.038908	0.089746
95.65	0.074731	0.074564	0.113752	0.072772
97.91	0.055812	0.055701	0.136731	0.055350

S	Γ	$\tilde{\Gamma}$	Γ_{Rich}	$\Gamma_{\text{Benchmark}}$
100.17	0.040392	0.039523	0.021365	0.040552
102.44	0.030133	0.029775	0.01651	0.032988
104.70	0.025293	0.025471	0.030854	0.025251
106.96	0.018415	0.018085	0.022512	0.019584
109.22	0.015632	0.015948	0.004657	0.014644
111.49	0.013046	0.013443	0.018982	0.011588

Table A.2: Gamma for $S \in [90, 110]$ for different partition sizes p.

Table A.3 reports the Theta estimates for the methods presented in Section (1.4.5) for different partition sizes p as well as the benchmark from Section (1.3). The underlying asset values S are interpolated linearly from the finite grid G. Theta estimates are obtained as follows: Θ_{FD} using Equation (1.4.16), Θ_{CD} with Equation (1.4.17) and Θ_{Rich} with Equation (1.4.18). The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

S	Θ_{FD}	Θ_{Rich}	Θ_{CD}	$\Theta_{\text{Benchmark}}$
		$p = 400$		
91.12	2.963171	2.942274	2.963171	3.201416
93.39	0.019477	0.038868	0.019477	0.201059
95.65	-0.114475	-0.114837	-0.114475	-0.106142
97.91	-0.149926	-0.149860	-0.149926	-0.140846
100.17	-0.173367	-0.173301	-0.173367	-0.162061
102.44	-0.184519	-0.184448	-0.184519	-0.173069
104.70	-0.186453	-0.186383	-0.186453	-0.174780
106.96	-0.181106	-0.181041	-0.181106	-0.169532
109.22	-0.170866	-0.170808	-0.170866	-0.159374
111.49	-0.157572	-0.157522	-0.157572	-0.146007
		$p = 1200$		
91.12	3.369446	3.341021	3.369446	3.201416
93.39	0.109517	0.133795	0.109517	0.201059
95.65	-0.105693	-0.106155	-0.105693	-0.106142
97.91	-0.141238	-0.141171	-0.141238	-0.140846
100.17	-0.163477	-0.163413	-0.163477	-0.162061
102.44	-0.174525	-0.174457	-0.174525	-0.173069
104.70	-0.176050	-0.175984	-0.176050	-0.174780
106.96	-0.170673	-0.170612	-0.170673	-0.169532
109.22	-0.160316	-0.160261	-0.160316	-0.159374
111.49	-0.146967	-0.146920	-0.146967	-0.146007

Table A.3: Theta for $S \in [90, 110]$ for different partition sizes p.

Table A.4 reports the relative errors computed using Equation (1.4.19) for the Theta estimates of different moneyness levels options at different partition sizes p and total time steps N_t. The parameters of the priced option are: $S_0 = 85, 100, 115$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

p	N_t	ITM	ATM	OTM
400	1000	0.000394	0.077924	0.453351
400	2000	0.000169	0.060577	0.320944
400	3000	0.000094	0.041103	0.208797
400	4000	0.000094	0.024936	0.098557
800	1000	0.000394	0.022931	0.098043
800	2000	0.000169	0.018710	0.071441
800	3000	0.000094	0.014540	0.045333
800	4000	0.000094	0.010711	0.020042
1200	1000	0.000394	0.012958	0.039183
1200	2000	0.000169	0.011201	0.027811
1200	3000	0.000094	0.009418	0.016580
1200	4000	0.000094	0.008191	0.005865
1600	1000	0.000394	0.009442	0.019143
1600	2000	0.000169	0.008550	0.012836
1600	3000	0.000094	0.007645	0.006673
1600	4000	0.000094	0.007319	0.000903
2000	1000	0.000394	0.007809	0.009929
2000	2000	0.000169	0.007324	0.005949
2000	3000	0.000094	0.006827	0.002085
2000	4000	0.000094	0.006909	-0.001385

Table A.4: Relative errors of Theta estimates at different partition sizes p and time steps N_t for different market scenarios.

Table A.5 reports the Vega estimate for the method presented in Section (1.4.6) for different partition sizes p as well as the benchmark from Section (1.3). The underlying asset values S are interpolated linearly from the finite grid G. The vega estimate is obtained using Equation (1.4.20). The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

S_{SDP}	$\mathcal{V}_{\text{CD}}$	$\mathcal{V}_{\text{Benchmark}}$
	$p = 400$	
91.12	13.015892	4.843706
93.39	32.102595	24.801207
95.65	44.405668	41.666816
97.91	50.095418	49.134366
100.17	51.322093	50.633498
102.44	49.653807	48.644655
104.70	46.261225	44.679772
106.96	41.948419	39.786170
109.22	37.260289	34.632518
111.49	32.558042	29.576249
	$p = 1200$	
91.12	9.764564	4.843706
93.39	30.10082	24.801207
95.65	43.95258	41.666816
97.91	50.17965	49.134366
100.17	51.484398	50.633498
102.44	49.725947	48.644655
104.70	46.180932	44.679772
106.96	41.712053	39.786170
109.22	36.892175	34.632518
111.49	36.892175	29.576249

Table A.5: Vega for $S \in [90, 110]$ for different partition sizes p.

Table A.6 reports the relative errors computed using Equation (1.4.19) for the Vega estimate of different moneyness levels options at different partition sizes p and total time steps N_t. The parameters of the priced option are: $S_0 = 85, 100, 115$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

p	$\Delta\sigma$	ITM	ATM	OTM
400	10^{-2}	138.617	0.012915	0.362647
400	10^{-3}	-0.182457	-0.003087	0.231403
400	10^{-4}	-0.235180	-0.005650	0.218038
400	10^{-5}	5.067469	-0.005936	0.216863
400	10^{-6}	9.554216	-0.005921	0.216754
800	10^{-2}	65.659	-0.015473	0.165552
800	10^{-3}	-0.741590	-0.000155	0.061433
800	10^{-4}	-1.158554	0.001910	0.050666
800	10^{-5}	-5.746987	-0.002302	0.049703
800	10^{-6}	-19.457831	-0.002299	0.049609
1200	10^{-2}	57.043233	0.016126	0.136923
1200	10^{-3}	-0.775855	0.000663	0.033570
1200	10^{-4}	-1.053012	-0.001062	0.022950
1200	10^{-5}	1.906024	-0.001139	0.021977
1200	10^{-6}	12.204819	-0.001195	0.021886
1600	10^{-2}	54.248009	0.016349	0.127087
1600	10^{-3}	-0.815421	0.000939	0.023929
1600	10^{-4}	-1.185060	-0.000677	0.013298
1600	10^{-5}	-4.426506	-0.000765	0.012218
1600	10^{-6}	-11.506024	-0.000743	0.012061
2000	10^{-2}	52.965826	0.016462	0.122569
2000	10^{-3}	-0.807469	0.001085	0.019494
2000	10^{-4}	-0.788915	-0.000560	0.008832
2000	10^{-5}	-0.469879	-0.0007235	0.007823
2000	10^{-6}	12.204819	-0.0007330	0.007732

Table A.6: Relative errors of Vega estimate at different partition sizes p and delta sigma $\Delta\sigma$ for different market scenarios.

Table A.7 reports the computational time, as well as the relative errors expressed in percentage % computed with Equation (1.4.19) for all Greeks methods of an ATM option. The parameters of the priced option are: $S_0 = 100$, $K = 100$, $T = 3$, $N_{\text{ex}} = 50$, $r = 0.05$, $q = 0$, $\sigma = 0.1$, put option - American (US).

p	Δ_{SDP}time	Δ_{SDP}	Δ_{CD} time	Δ_{CD}	Δ_{Rich} time	Δ_{Rich}	Δ_{mean} time	Δ_{mean}
300	7.2×10^{-6}	-0.129639	1.02×10^{-5}	-0.127432	6.5×10^{-6}	-0.124040	3.6×10^{-5}	-0.127432
400	7.2×10^{-6}	-0.117739	1.11×10^{-5}	-0.118644	6.8×10^{-6}	-0.115209	3.8×10^{-5}	-0.118644
500	7.3×10^{-6}	-0.118636	9.20×10^{-6}	-0.114012	5.8×10^{-6}	-0.110437	3.6×10^{-5}	-0.114012
600	8.3×10^{-6}	-0.113666	9.20×10^{-6}	-0.111578	5.7×10^{-6}	-0.107923	3.8×10^{-5}	-0.111578
700	7.3×10^{-6}	-0.110314	8.30×10^{-6}	-0.110018	5.5×10^{-6}	-0.106366	3.6×10^{-5}	-0.110018
800	7.6×10^{-6}	-0.107982	8.60×10^{-6}	-0.109062	5.5×10^{-6}	-0.105371	3.9×10^{-5}	-0.109062
900	7.6×10^{-6}	-0.106209	9.10×10^{-6}	-0.108345	5.6×10^{-6}	-0.104645	3.8×10^{-5}	-0.108345
1200	2.7×10^{-5}	-0.106099	9.10×10^{-6}	-0.107198	5.4×10^{-6}	-0.103469	3.6×10^{-5}	-0.107198
1400	1.18×10^{-5}	-0.104343	9.30×10^{-6}	-0.106806	5.6×10^{-6}	-0.103065	3.7×10^{-5}	-0.106806
1600	1.12×10^{-5}	-0.105448	8.40×10^{-6}	-0.106555	5.6×10^{-6}	-0.102813	7.3×10^{-5}	-0.106555
1800	9.90×10^{-6}	-0.104216	8.30×10^{-6}	-0.106381	5.7×10^{-6}	-0.102634	4.8×10^{-5}	-0.106381
2000	9.6×10^{-6}	-0.105137	8.70×10^{-6}	-0.106251	5.6×10^{-6}	-0.102498	4.3×10^{-5}	-0.106251

p	Γ_{CD} time	Γ_{CD}	$\tilde{\Gamma}$ time	$\tilde{\Gamma}$	Γ_{Rich} time	Γ_{Rich}	Θ_{CD} time	Θ_{CD}
300	1.4×10^{-5}	-5.562849	1.4×10^{-5}	-5.559877	1.1×10^{-5}	-5.572200	8×10^{-6}	0.923131
400	1.4×10^{-5}	-5.651841	1.4×10^{-5}	-5.646989	1.3×10^{-5}	-5.638179	9×10^{-6}	0.926943
500	1.4×10^{-5}	-5.694789	1.2×10^{-5}	-5.687277	1.0×10^{-5}	-5.680195	8×10^{-6}	0.928629
600	1.3×10^{-5}	-5.729536	1.3×10^{-5}	-5.723187	1.1×10^{-5}	-5.723260	9×10^{-6}	0.929564
700	1.3×10^{-5}	-5.749937	1.2×10^{-5}	-5.743323	1.0×10^{-5}	-5.742981	8×10^{-6}	0.930137
800	1.3×10^{-5}	-5.759370	1.3×10^{-5}	-5.752418	1.1×10^{-5}	-5.752391	8×10^{-6}	0.930509
900	1.3×10^{-5}	-5.766706	1.3×10^{-5}	-5.759248	1.0×10^{-5}	-5.758140	8×10^{-6}	0.930762
1200	1.4×10^{-5}	-5.782669	1.2×10^{-5}	-5.775025	1.1×10^{-5}	-5.775059	8×10^{-6}	0.931180
1400	1.4×10^{-5}	-5.786651	1.2×10^{-5}	-5.778769	1.0×10^{-5}	-5.778556	8×10^{-6}	0.931325
1600	1.7×10^{-5}	-5.790870	1.7×10^{-5}	-5.783064	1.4×10^{-5}	-5.783183	8×10^{-6}	0.931418
1800	1.3×10^{-5}	-5.793050	1.2×10^{-5}	-5.785174	1.0×10^{-5}	-5.785198	8×10^{-6}	0.931483
2000	1.7×10^{-5}	-5.794879	1.5×10^{-5}	-5.786976	1.2×10^{-5}	-5.787028	1.1×10^{-5}	0.931528

p	Θ_{FD} time	Θ_{FD}	Θ_{Rich} time	Θ_{Rich}
300	5×10^{-6}	0.923131	6×10^{-6}	0.923160
400	5×10^{-6}	0.926943	6×10^{-6}	0.926971
500	4×10^{-6}	0.928629	6×10^{-6}	0.928657
600	4×10^{-6}	0.929564	6×10^{-6}	0.929592
700	4×10^{-6}	0.930137	5×10^{-6}	0.930164
800	4×10^{-6}	0.930509	5×10^{-6}	0.930536
900	4×10^{-6}	0.930762	6×10^{-6}	0.930789
1200	4×10^{-6}	0.931180	5×10^{-6}	0.931208
1400	4×10^{-6}	0.931325	5×10^{-6}	0.931352
1600	4×10^{-6}	0.931418	5×10^{-6}	0.931444
1800	4×10^{-6}	0.931483	5×10^{-6}	0.931510
2000	6×10^{-6}	0.931528	7×10^{-6}	0.931555

Table A.7: ATM option relative errors and computational time in seconds.

Appendix B

Numerical Implementation of the Finite-Difference Method

In the absence of arbitrage opportunities, Itô's lemma allows us to express that any derivative product $\overline{V}$ following an underlying asset S governed by a GBM obeys the fundamental PDE under the Black and Scholes (1973) model

$$\frac{\partial \overline{V}(s,t)}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 \overline{V}(s,t)}{\partial S^2} + (r-q)S\frac{\partial \overline{V}(s,t)}{\partial S} = r\overline{V}(s,t), \tag{B.0.1}$$

where σ^2 represents the variance of the underlying asset's return, $(r-q)S$ reflects the contribution of the asset's drift to the option's price sensitivity, incorporating both the risk-free rate r and the convenience yield q. The implicit FD method then involves discretizing this PDE in a way that it is resolved it as an algorithm with the Crank-Nicolson method.

We define a grid $\overline{G}$ consisting of discrete grid points, where the passage of time is represented on the horizontal axis, and changes in the underlying asset's price on the vertical axis. The grid has J vertical intervals ΔS representing price increments and I horizontal intervals $\Delta t = t_{n+1} - t_n$ representing constant time increments with $n = 0, 1, \ldots, N$ and $N = T$. The upper barrier of $\overline{G}$ for the underlying asset price is approximately four times greater than the exercise price of the option and is denoted as $S_{\max}$. The lower boundary $S_{\min}$, is typically set to zero, representing the minimum possible asset price. We define $\Delta S = (S_{\max} - S_{\min})/M$ and consider a total of $M+1$ equally spaced stock prices $\{0, \Delta S, 2\Delta S, \ldots, S_{\max}\}$.

The time points and stock prices define a mesh $\overline{G}$ that consists of a total of $(M+1)(N+1)$ points, where the point (i,j) corresponds to time $t = T - i\Delta t$ and stock price $S_t = j\Delta S$ with $0 \leq j \leq J$ and $0 \leq i \leq I$. We will also consider that $V_{i,j}$ is the value of the option price at the (i,j) mesh point. We can clearly see that

i illustrates the remaining time before maturity T, which makes a recursion-based algorithm more practical at the time of valuation because it allows to iteratively determine the option price from maturity back to the present time.

For an interior point (i, j) on the grid, $\frac{\partial V_{i,j}}{\partial S}$ can be approximated as

$$\frac{\partial V_{i,j}}{\partial S} = \Delta_V \approx \frac{V_{i,j+1} - V_{i,j-1}}{2\Delta S}, \tag{B.0.2}$$

where Δ_V, known as Delta, represents the sensitivity of the option value to changes in the underlying asset price.

For $\frac{\partial V_{i,j}}{\partial t}$, we use a forward difference approximation so that the value at time $T - i\Delta t$ is related to the value at time $T - (i+1)\Delta t$

$$\frac{\partial V_{i,j}}{\partial t} = \Theta_V \approx \frac{V_{i+1,j} - V_{i,j}}{\Delta t}, \tag{B.0.3}$$

where Θ, known as Theta, represents the sensitivity of the option value to the passage of time.

A backward difference approximation for $\frac{\partial V_{i,j+1}}{\partial S}$ at the $(i, j+1)$ point is

$$\frac{\partial V_{i,j+1}}{\partial S} = \Delta_V \approx \frac{V_{i,j+1} - V_{i,j}}{\Delta S}, \tag{B.0.4}$$

and for the point (i, j)

$$\frac{\partial V_{i,j}}{\partial S} = \Delta_V \approx \frac{V_{i,j} - V_{i,j-1}}{\Delta S}. \tag{B.0.5}$$

A central difference approximation for $\frac{\partial V^2_{i,j}}{\partial S^2}$ at point (i, j) is found by combining equations (B.0.4) and (B.0.5) :

$$\frac{\partial V^2_{i,j}}{\partial S^2} = \Gamma_V \approx \frac{V_{i,j+1} - 2V_{i,j} + V_{i,j-1}}{(\Delta S)^2}, \tag{B.0.6}$$

where Γ, known as Gamma, is the sensitivity of the Delta Δ_V to changes in the underlying asset price.

For both explicit and implicit methods, the FD requires boundary conditions to initiate the algorithm. In this case, we provide boundary conditions for put options, as the benchmark is focused on put options only. Boundary conditions for call options can be derived in a similar manner if needed.

The value of an option at time T is defined in equation (1.2.5) and is denoted as $V_{0,j}$ with $j = 0, 1, ..., J$

$$V_{0,j} = \max(K - S_T, 0) = \max(K - j\Delta S, 0). \tag{B.0.7}$$

The lower boundary condition, when the stock price is zero and $i = 0, 1, ..., I$ is

$$V_{i,0} = K. \tag{B.0.8}$$

The upper boundary condition as the stock price approaches infinity, or $S_{\max}$, is expressed as

$$V_{i,J} = 0. \tag{B.0.9}$$

Equations (B.0.7), (B.0.8) and (B.0.9) define the options values along the three main edges of the grid when $t = T$, $S = 0$, and $S = S_{\max}$, respectively.
Another relevant boundary condition would be that in areas that are strongly In-The-Money (ITM) , the second derivative $\frac{\partial V_{i,j}^2}{\partial S^2}$, or Gamma Γ, tends to zero since the payoff becomes linear in S, and, in terms of programming, it is written as

$$V_{i,0} = 2V_{i,1} - V_{i,2}.$$

Back to our PDE, we substitute Equations (B.0.2), (B.0.3), and (B.0.6) into Equation (B.0.1) with $S = j\Delta S$, yielding

$$\frac{V_{i+1,j} - V_{i,j}}{\Delta t} + (r - q)j\Delta S\frac{V_{i,j+1} - V_{i,j-1}}{2\Delta S} + \frac{1}{2}\sigma^2 j^2 \Delta S^2 \frac{V_{i,j+1} - 2V_{i,j} + V_{i,j-1}}{(\Delta S)^2} = rV_{i,j}, \tag{B.0.10}$$

for $j = 1, 2, ..., M - 1$ and $i = 0, 1, ..., N - 1$.

We can rearrange terms to obtain

$$V_{i+1,j} = a_{i,j}V_{i,j-1} + b_{i,j}V_{i,j} + c_{i,j}V_{i,j+1}, \tag{B.0.11}$$

where we set

$$\begin{aligned} a_{i,j} &= \tfrac{1}{2}(r - q)j\Delta t - \tfrac{1}{2}\sigma^2 j^2 \Delta t, \\ b_{i,j} &= 1 + \sigma^2 j^2 \Delta t + r\Delta t, \\ \hat{c}_{i,j} &= -\tfrac{1}{2}(r - q)j\Delta t - \tfrac{1}{2}\sigma^2 j^2 \Delta t. \end{aligned}$$

It remains now to use the Equation (B.0.11) to arrive at the value of V at all other points. We first consider points corresponding to time $T - \Delta t$ using Equation (B.0.11) with $i = N - 1$ and get

$$V_{N,j} = a_{i,j}V_{N-1,j-1} + b_{i,j}V_{N-1,j} + \hat{c}_{i,j}V_{N-1,j+1}, \tag{B.0.12}$$

for $j = 1, 2, \ldots, M-1$. The left-hand side of the equation $V_{N,j}$ is just $V_{0,j}$ which is a boundary condition since $N = T$. Furthermore, from Equations (B.0.8) and (B.0.9), we can derive $V_{N-1,0} = K$ and $V_{N-1,J} = 0$ respectively.
Equation (B.0.12) is therefore $M-1$ simultaneous equations that can be solved for the $M-1$ unknowns $V_{N-1,1}, V_{N-1,2}, \ldots, V_{N-1,M-1}$.

Each value of $V_{N-1,j}$ is then compared with the payoff function $K - j\Delta S$. If $V_{N-1,j} < K - j\Delta S$, early exercise at time $T - \Delta t$ is optimal, which means that the value of holding the option is less than the payoff from exercising it immediately. In such a case, early exercise is chosen as the optimal decision, and $V_{N-1,j}$ is set equal to the immediate payoff. The nodes corresponding to $T - 2\Delta t$ are handled in a similar way, and so on until obtaining $V_{0,1}, V_{0,2}, V_{0,3}, \ldots, V_{0,M-1}$ which are approximations of the option price at different points in time until maturity. The analogue of Equation (B.0.11) for the implicit FD is

$$V_{i,j} = \alpha_{i+1,j} V_{i+1,j-1} + \beta_{i+1,j} V_{i+1,j} + \gamma_{i+1,j} V_{i+1,j+1}, \tag{B.0.13}$$

where $\alpha_{i+1,j}$, $\beta_{i+1,j}$, and $\gamma_{i+1,j}$ are $a_{i,j}$, $b_{i,j}$, and $c_{i,j}$ respectively evaluated at $i+1$.

Equation (B.0.13) is evaluated by approximating $\frac{\partial V_{i+1,j}}{\partial S}$ and $\frac{\partial V^2_{i+1,j}}{\partial S^2}$ of Equation (B.0.10) at $i+1$ as

$$\frac{V_{i+1,j} - V_{i,j}}{\Delta t} + (r-q) j \Delta S \frac{V_{i+1,j+1} - V_{i+1,j-1}}{2\Delta S} + \frac{1}{2}\sigma^2 j^2 \Delta S^2 \frac{V_{i+1,j+1} - 2V_{i+1,j} + V_{i+1,j-1}}{(\Delta S)^2} = r V_{i+1,j}.$$

Combining Equations (B.0.11) and (B.0.13), we obtain

$$V_{i+1,j} + V_{i,j} = a_{i,j} V_{i,j-1} + b_{i,j} V_{i,j} + \hat{c}_{i,j} V_{i,j+1} + \alpha_{i+1,j} V_{i+1,j-1} + \beta_{i+1,j} V_{i+1,j} + \gamma_{i+1,j} V_{i+1,j+1}.$$

We set $h^*_{i,j} = -a_{i,j} V_{i,j-1} - b_{i,j} V_{i,j} - \hat{c}_{i,j} V_{i,j+1}$, which can be computed directly since it only involves elements of the same step t_i and can be written in the same way for the implicit FD equation (B.0.13). Considering the boundary conditions of Equations (B.0.8) and (B.0.9) at $j = 0$ and $j = J$, the Crank-Nicolson algorithm is a system of $J-1$ equations and $J-1$ variables. This system is then resolved with iterative methods such as *Jacobi*, *Gauss-Seidel (GS)*, or *Successive Over-Relaxation (SOR)* of Young (1954).

We use the SOR method due to the computational cost of using the Jacobi and GS that resolve the linear system at each time step with the matrix inverse A^{-1} of a given matrix A. The chosen method is indeed an extension of GS, which itself is an improvement on Jacobi. The main idea behind SOR is to accelerate the convergence

by using a relaxation factor ω. Below are listed forms of the GS and SOR methods respectively

$$x_i^{(k+1)} = \frac{1}{a_{ii}} \left(b_i - \sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)} \right), \tag{B.0.14}$$

$$x_i^{(k+1)} = (1-\omega) x_i^{(k)} + \frac{\omega}{a_{ii}} \left(b_i - \sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)} \right), \tag{B.0.15}$$

where x_i^{k+1} is the value of the i-th variable at the $(k+1)$-th iteration. The term a_{ii} is the i-th diagonal element of matrix A, which is also the i-th element of the diagonal matrix D. The coefficient of the j-th variable in the i-th equation is a_{ij}. The term b_i is the i-th element of the constant terms vector b which is on the right-hand side of the system of linear equations represented by the matrix equation $Ax = b$. Each element b_i corresponds to the constant term in the i-th equation, and the range of b is defined by the number of equations, which is the same as the number of rows in matrix A, meaning $i = 1, \ldots, n$. The parameter ω is the relaxation factor used in the SOR method, chosen to optimize the rate of convergence. When $\omega = 1$, the SOR method reduces to GS.

Matrix A is the coefficient matrix where each element a_{ij} represents the coefficient of the j-th variable in the i-th linear equation of the system. Matrix D plays a pivotal role in analyzing the convergence of these iterative methods, as it influences the rate at which these methods converge to a solution. The matrix D for the SOR method, same as GS is

$$D = \begin{bmatrix} a_{11} & 0 & \cdots & 0 \\ 0 & a_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{nn} \end{bmatrix}.$$

More detailed explanations about how each of of the operation and convergence properties of the Jacobi, GS, and SOR methods operate can be found in the renowned iterative methods for linear systems textbook of Saad (2003).

We note $V_j^{(u)}$ as the value of $V_{i+1,j}$ after u iterations that is obtained using equation (B.0.15) which satisfies

$$g_j^{(u)} = V_j^{(u-1)} + \frac{\omega}{\beta_{i+1,j}} \left(h_j^* - \alpha_{i+1,j} V_{j-1}^{(u)} - \beta_{i+1,j} V_j^{(u-1)} - \gamma_{i+1,j} V_{j+1}^{(u-1)} \right), \tag{B.0.16}$$

where the last term on the right hand side of equation (B.0.16) stands for the correction made during the u-th iteration, multiplied by ω, the relaxation factor to

optimize convergence. To initialize the iteration, we take the value of the derivative product $\overline{V}$ at the initial time step $\overline{V}_{0,j} = \hat{V}_{i,j}$.Numerical convergence is achieved when a certain level of precision is reached, that is when

$$\sqrt{\sum_{k=1}^{J-1} \left(g_k^{(u)} - g_k^{(u-1)} \right)^2} < \epsilon,$$

where ϵ is an arbitrary low threshold of tolerance set to 10^{-6}.

Since we are evaluating American options, it is easy to modify our iterative procedure as

$$\max\{V_i^{(u-1)} + \frac{\omega}{\beta_{i+1,j}}(h_i - \alpha_{i+1,j}V_{i-1}^{(u)} - \beta_{i+1,j}V_i^{(u-1)} - \gamma_{i+1,j}V_{i+1}^{(u-1)}), v_i^e\},$$

where v_i^e is the exercise value of the option when the underlying asset is S_i.

We outline the steps of the algorithm as follows:

1. **Initialization:** Set up the computational grid $\overline{G}$ for asset prices and time, with boundary conditions at the extremes of the grid with Equations (B.0.7), (B.0.8) and (B.0.9).

2. **Discretization:** Use the Crank-Nicolson FD method to discretize the PDE over $\overline{G}$ by approximating the partial derivatives of Equations (B.0.3), (B.0.2), (B.0.6) with respect to time and the underlying asset price respectively.

3. **Backward Induction:** Starting from the maturity date $N = T$, move backward through time $i = N - 1, \ldots, 0$ solving Equation (B.0.13) at each time step.

4. **SOR Iteration:** At each time step i and grid point j, solve the discretized Equation (B.0.16) with SOR.

5. **Early Exercise Check:** At certain time steps, corresponding to potential early exercise dates N_{ex}, compare the option value from solving the PDE to the payoff function $h(S_t)$. If the intrinsic value is higher, the option should be exercised, and the propagated option value is updated accordingly.

6. **Convergence:** Repeat the SOR iterations until numerical convergence.

7. **Result:** Once the iteration through all time steps is complete, the option value at the initial asset price and time zero is the approximated price of the option.

Appendix C

Appendix for Chapter 2

Tables from C.1 to C.5 report the relative errors for different market scenarios options in a low volatility environment. The corresponding Figures can be found in Section (2.3.1).

Tables from C.6 to C.10 report the relative errors for different market scenarios options in a high volatility environment. The corresponding Figures can be found in Section (2.3.2).

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.000670502	-2.88E-12	5.86E-07	-1.02E-07	5.86E-07	-147773.3628	-21930.59194	-66145.61953	-8.97E-05	-8.97E-05	-5.03E-05	-11.87499219
400	-0.000670502	-2.88E-12	6.90E-08	-2.00E-08	6.90E-08	-17825.74853	7354.533767	-6707.486945	-9.34E-05	-9.34E-05	-5.57E-05	1
500	-0.000670502	-2.88E-12	2.20E-08	-7.00E-09	2.20E-08	-5839.605767	4492.871503	-2156.974821	-9.37E-05	-9.37E-05	-5.61E-05	1
600	-0.000670502	-2.88E-12	1.00E-08	-4.00E-09	1.00E-08	-2755.109147	3050.312248	-879.7821146	-9.37E-05	-9.37E-05	-5.62E-05	1
700	-0.000670502	-2.88E-12	7.00E-09	-3.00E-09	7.00E-09	-1746.49048	2334.896682	-515.0374573	-9.37E-05	-9.37E-05	-5.62E-05	1
800	-0.000670502	-2.88E-12	5.00E-09	-2.00E-09	5.00E-09	-1218.724899	1924.412341	-326.2146604	-9.37E-05	-9.37E-05	-5.62E-05	1
900	-0.000670502	-2.88E-12	4.00E-09	-2.00E-09	4.00E-09	-959.533358	1654.665488	-246.463417	-9.37E-05	-9.37E-05	-5.62E-05	1
1200	-0.000670502	-2.88E-12	2.00E-09	-1.00E-09	2.00E-09	-633.49151	1279.365519	-147.9471752	-9.37E-05	-9.37E-05	-5.62E-05	1
1400	-0.000670502	-2.88E-12	2.00E-09	-1.00E-09	2.00E-09	-551.3946418	1173.812403	-124.4909271	-9.37E-05	-9.37E-05	-5.62E-05	1
1600	-0.000670502	-2.88E-12	2.00E-09	-1.00E-09	2.00E-09	-497.4452713	1101.098034	-108.4233972	-9.37E-05	-9.37E-05	-5.62E-05	1
1800	-0.000670502	-2.88E-12	2.00E-09	-1.00E-09	2.00E-09	-469.2977736	1057.703975	-102.0902102	-9.38E-05	-9.38E-05	-5.62E-05	1
2000	-0.000670502	-2.88E-12	2.00E-09	-1.00E-09	2.00E-09	-444.6687131	1022.519603	-94.23236712	-9.38E-05	-9.38E-05	-5.62E-05	1

Table C.1: Relative errors for option with $S_0 = 85$, $K = 100$, $T = 1$ in low volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.000419056	1.38E-08	5.86E-07	-1.02E-07	5.86E-07	-2442.39284	-361.6305246	-1092.709176	-8.34E-05	-8.34E-05	-4.40E-05	-25.1347633
400	-0.000419055	2.82E-09	6.88E-08	-2.02E-08	6.88E-08	-293.7585013	122.5878818	-109.9222781	-8.71E-05	-8.71E-05	-4.94E-05	-4.958184379
500	-0.000419055	8.25E-10	2.18E-08	-7.18E-09	2.18E-08	-95.57219318	75.27138552	-34.68129226	-8.74E-05	-8.74E-05	-4.98E-05	-1.437439064
600	-0.000419055	-1.75E-10	9.82E-09	-4.18E-09	9.82E-09	-44.57121566	51.41921732	-13.563397	-8.75E-05	-8.75E-05	-4.99E-05	-0.489546095
700	-0.000419055	-1.75E-10	6.82E-09	-3.18E-09	6.82E-09	-27.89408993	39.59009326	-7.532482932	-8.75E-05	-8.75E-05	-5.00E-05	-0.083306251
800	-0.000419055	-1.75E-10	4.82E-09	-2.18E-09	4.82E-09	-19.16768693	32.80289093	-4.410369859	-8.75E-05	-8.75E-05	-5.00E-05	0.187520312
900	-0.000419055	-1.75E-10	3.82E-09	-2.18E-09	3.82E-09	-14.88205346	28.34272939	-3.091713406	-8.75E-05	-8.75E-05	-5.00E-05	0.322933593
1200	-0.000419055	-1.75E-10	1.82E-09	-1.18E-09	1.82E-09	-9.491075605	22.13728726	-1.462784846	-8.75E-05	-8.75E-05	-5.00E-05	0.458346875
1400	-0.000419055	-1.75E-10	1.82E-09	-1.18E-09	1.82E-09	-8.133635138	20.39200666	-1.074944713	-8.75E-05	-8.75E-05	-5.00E-05	0.458346875
1600	-0.000419055	-1.75E-10	1.82E-09	-1.18E-09	1.82E-09	-7.241602832	19.18970225	-0.809274222	-8.75E-05	-8.75E-05	-5.00E-05	0.593760156
1800	-0.000419055	-1.75E-10	1.82E-09	-1.18E-09	1.82E-09	-6.776194672	18.472198	-0.704557386	-8.75E-05	-8.75E-05	-5.00E-05	0.593760156
2000	-0.000419055	-1.75E-10	1.82E-09	-1.18E-09	1.82E-09	-6.368962532	17.8904378	-0.574630941	-8.75E-05	-8.75E-05	-5.00E-05	0.593760156

Table C.2: Relative errors for option with $S_0 = 85$, $K = 100$, $T = 2$ in low volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.001091317	0.001177	0.024652	0.019843	0.024652	-24.0174	-23.4788	-23.0434	0.273052	0.273052	0.279322	-1.40657
400	-0.001033726	0.000520	0.015788	0.011241	0.015788	-17.2758	-16.5012	-16.2006	0.177967	0.177967	0.184984	-0.491675
500	-0.001016925	0.000130	0.012318	0.008108	0.012318	-14.3701	-13.5471	-13.3957	0.135871	0.135871	0.142696	-0.187658
600	-0.001010113	-3.51E-05	0.010672	0.006739	0.010672	-12.7600	-11.8909	-11.7823	0.116550	0.116550	0.123136	-0.039911
700	-0.001006703	-0.000120	0.009685	0.005939	0.009685	-11.7659	-10.8646	-10.7483	0.104961	0.104961	0.111353	0.048169
800	-0.001004455	-0.000105	0.009062	0.005399	0.009062	-11.1595	-10.2458	-10.1055	0.097185	0.097185	0.103426	0.102154
900	-0.001002999	-0.000154	0.008653	0.005064	0.008653	-10.7522	-9.83378	-9.70947	0.092104	0.092104	0.098237	0.139090
1200	-0.001001033	-0.000185	0.008008	0.004543	0.008008	-10.0849	-9.15765	-9.04173	0.084425	0.084425	0.090386	0.195916
1400	-0.001000358	-0.000219	0.007803	0.004384	0.007803	-9.86976	-8.93934	-8.82884	0.081888	0.081888	0.087779	0.212964
1600	-0.000999949	-0.000211	0.007659	0.004265	0.007659	-9.71961	-8.78728	-8.67548	0.080132	0.080132	0.085981	0.227170
1800	-0.00099966	-0.000231	0.007570	0.004195	0.007570	-9.62514	-8.69235	-8.58331	0.079040	0.079040	0.084858	0.235694
2000	-0.000999469	-0.000222	0.007501	0.004140	0.007501	-9.55296	-8.61890	-8.50973	0.078201	0.078201	0.083997	0.241377

Table C.3: Relative errors for option with $S_0 = 90$, $K = 100$, $T = 1$ in low volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.066436622	-0.023926	-0.021926	-0.018852	-0.021926	0.045330	0.045762097	0.043970	-0.131390	-0.131390	-0.130960	-0.005050
400	-0.038510853	-0.013141	-0.013961	-0.010847	-0.013961	0.032385	0.033090	0.034372	-0.075284	-0.075284	-0.074867	-0.001170
500	-0.024629504	-0.013953	-0.009762	-0.006522	-0.009762	0.026137	0.027230	0.028260	-0.050463	-0.050463	-0.050051	-0.003110
600	-0.017370438	-0.009448	-0.007556	-0.004243	-0.007556	0.021082	0.022006095	0.021995	-0.036705	-0.036705	-0.036299	-0.001170
700	-0.012701606	-0.006410	-0.006142	-0.002832	-0.006142	0.018115	0.019076999	0.019127	-0.028271	-0.028271	-0.027868	-0.001170
800	-0.009844654	-0.004297	-0.005275	-0.001930	-0.005275	0.016743	0.017754	0.017758	-0.022796	-0.022796	-0.022396	-0.001170
900	-0.007773678	-0.002689	-0.004625	-0.001271	-0.004625	0.015676	0.016760535	0.016922	-0.019076	-0.019076	-0.018678	-0.001170
1200	-0.004377703	-0.002589	-0.003586	-0.000206	-0.003586	0.013354	0.01446545	0.014460	-0.012913	-0.012913	-0.012515	-0.001170
1400	-0.003239331	-0.000998	-0.003230	0.000160	-0.003230	0.012774	0.013920788	0.013952	-0.010790	-0.010790	-0.010394	0.000770
1600	-0.00247372	-0.002000	-0.003003	0.000389	-0.003003	0.012161	0.013296118	0.013279	-0.009424	-0.009424	-0.009028	-0.001170
1800	-0.00195607	-0.000883	-0.002845	0.000551	-0.002845	0.011843	0.012989105	0.012986	-0.008462	-0.008462	-0.008066	0.000770
2000	-0.001582412	-0.001718	-0.002727	0.000675	-0.002727	0.011577	0.012727018	0.012719	-0.007803	-0.007803	-0.007407	0.000770

Table C.4: Relative errors for option with $S_0 = 100$, $K = 100$, $T = 3$ in low volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.144697	-0.101854	-0.101935	-0.097005	-0.101935	-0.056791	-0.055686	-0.055828	-0.135579	-0.135579	-0.135268	-0.075643
400	-0.082028	-0.060076	-0.060282	-0.055345	-0.060282	-0.027065	-0.025750	-0.026191	-0.073720	-0.073720	-0.073421	-0.043772
500	-0.052274	-0.039461	-0.039659	-0.034797	-0.039659	-0.013058	-0.011753	-0.011857	-0.044648	-0.044648	-0.044354	-0.027837
600	-0.036479	-0.028611	-0.028815	-0.023977	-0.028815	-0.005668	-0.004265	-0.004231	-0.029151	-0.029151	-0.028861	-0.019869
700	-0.026716	-0.021754	-0.021969	-0.017137	-0.021969	-0.000790	0.000646	0.000683	-0.019664	-0.019664	-0.019376	-0.015885
800	-0.020627	-0.017454	-0.017679	-0.012847	-0.017679	0.002286	0.003734	0.003732	-0.013682	-0.013682	-0.013395	-0.011901
900	-0.016296	-0.014446	-0.014668	-0.009851	-0.014668	0.004330	0.005784	0.005832	-0.009468	-0.009468	-0.009182	-0.007917
1200	-0.009156	-0.006836	-0.009631	-0.004822	-0.009631	0.007742	0.009209	0.009214	-0.002555	-0.002555	-0.002270	-0.003933
1400	-0.006758	-0.005522	-0.007947	-0.003143	-0.007947	0.008838	0.010300	0.010295	-0.000219	-0.000219	0.000065	-0.003933
1600	-0.005166	-0.004667	-0.006818	-0.002015	-0.006818	0.009615	0.011081	0.011067	0.001317	0.001317	0.001600	-0.003933
1800	-0.004084	-0.004119	-0.006056	-0.001255	-0.006056	0.010124	0.011589	0.011584	0.002364	0.002364	0.002648	-0.003933
2000	-0.003305	-0.003737	-0.005504	-0.000703	-0.005504	0.010499	0.011968	0.011968	0.003117	0.003117	0.003400	0.000051

Table C.5: Relative errors for option with $S_0 = 115$, $K = 100$, $T = 3$ in low volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.012445788	0.004899	0.230704	0.230771	0.230704	0.339557	0.338836391	0.338681	0.527276	0.527276	0.527375	-0.028663
400	-0.007318206	0.003396	0.226330	0.226387	0.226330	0.344666	0.344425031	0.344337	0.539547	0.539547	0.539647	-0.010469
500	-0.004711523	0.002637	0.224064	0.224135	0.224064	0.345600	0.345530750	0.345506	0.545280	0.545280	0.545380	-0.009805
600	-0.003266157	-0.000710	0.222776	0.222843	0.222776	0.345388	0.345368376	0.345370	0.548372	0.548372	0.548472	-0.007199
700	-0.002398713	-0.000416	0.222041	0.222108	0.222041	0.345203	0.345206441	0.345202	0.550232	0.550232	0.550332	-0.003462
800	-0.001785728	-0.000229	0.221495	0.221563	0.221495	0.344884	0.344892138	0.344885	0.551449	0.551449	0.551549	-0.003198
900	-0.001374607	-0.000031	0.221131	0.221198	0.221131	0.344651	0.344664683	0.344663	0.552280	0.552280	0.552380	-0.002039
1200	-0.000686579	-0.001005	0.220525	0.220592	0.220525	0.344255	0.344278361	0.344277	0.553651	0.553651	0.553751	-0.001409
1400	-0.000441194	-0.000549	0.220312	0.220380	0.220312	0.344003	0.344026802	0.344025	0.554118	0.554118	0.554218	-0.001056
1600	-0.000291933	-0.000182	0.220179	0.220247	0.220179	0.343945	0.343971801	0.343971	0.554422	0.554422	0.554521	-0.000557
1800	-0.000185470	-0.000870	0.220086	0.220154	0.220086	0.343885	0.343913878	0.343914	0.554630	0.554630	0.554729	-0.000367
2000	-0.000107525	-0.000535	0.220017	0.220085	0.220017	0.343809	0.343837402	0.343836	0.554779	0.554779	0.554878	-0.000089

Table C.6: Relative Errors for option with $S_0 = 85$, $K = 100$, $T = 1$ in high volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.016916631	0.007755	0.102758	0.102836	0.102758	0.138578	0.137637543	0.137435	0.239802	0.239802	0.239962	-0.030119
400	-0.009906039	0.005066	0.097657	0.097723	0.097657	0.145241	0.144926862	0.144812	0.259536	0.259536	0.259696	-0.011482
500	-0.006411402	0.003710	0.095014	0.095096	0.095014	0.146459	0.146369064	0.146338	0.268756	0.268756	0.268916	-0.010994
600	-0.004494326	0.000358	0.093512	0.093590	0.093512	0.146183	0.146157278	0.146160	0.273728	0.273728	0.273888	-0.007794
700	-0.003343356	0.000407	0.092654	0.092733	0.092654	0.145942	0.145946064	0.145940	0.276719	0.276719	0.276879	-0.004150
800	-0.002544133	0.000426	0.092018	0.092097	0.092018	0.145525	0.145536116	0.145527	0.278675	0.278675	0.278836	-0.003819
900	-0.002006240	0.000503	0.091593	0.091671	0.091593	0.145221	0.145239444	0.145237	0.280012	0.280012	0.280172	-0.002625
1200	-0.001107844	-0.000534	0.090886	0.090965	0.090886	0.144705	0.144735560	0.144734	0.282217	0.282217	0.282377	-0.001865
1400	-0.000791607	-0.000183	0.090638	0.090717	0.090638	0.144377	0.144407449	0.144406	0.282968	0.282968	0.283128	-0.001456
1600	-0.000595587	0.000106	0.090483	0.090562	0.090483	0.144301	0.144335710	0.144334	0.283456	0.283456	0.283616	-0.000996
1800	-0.000457290	-0.000533	0.090374	0.090453	0.090374	0.144222	0.144260161	0.144260	0.283790	0.283790	0.283950	-0.000786
2000	-0.000356432	-0.000252	0.090293	0.090372	0.090293	0.144124	0.144160413	0.144159	0.284030	0.284030	0.284190	-0.000547

Table C.7: Relative Errors for option with $S_0 = 85$, $K = 100$, $T = 2$ in high volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.015477984	0.005700	0.214307	0.214393	0.214307	0.364320	0.363809048	0.363696	0.546023	0.546023	0.546121	-0.026244
400	-0.009123756	0.002465	0.209303	0.209374	0.209303	0.367364	0.367029767	0.366913	0.557600	0.557600	0.557697	-0.009216
500	-0.005870693	0.000698	0.206746	0.206820	0.206746	0.369554	0.369519577	0.369513	0.563003	0.563003	0.563100	-0.009774
600	-0.004067520	-0.000230	0.205391	0.205463	0.205391	0.369379	0.369366047	0.369356	0.565919	0.565919	0.566016	-0.006707
700	-0.002991861	-0.000727	0.204616	0.204690	0.204616	0.369189	0.369190822	0.369187	0.567675	0.567675	0.567773	-0.003182
800	-0.002223969	-0.001137	0.204048	0.204122	0.204048	0.368928	0.368935547	0.368935	0.568821	0.568821	0.568918	-0.002974
900	-0.001710113	0.000740	0.203667	0.203741	0.203667	0.368825	0.368843223	0.368842	0.569604	0.569604	0.569701	-0.001850
1200	-0.000850301	-0.000237	0.203041	0.203115	0.203041	0.368486	0.368507975	0.368507	0.570896	0.570896	0.570994	-0.001369
1400	-0.000544199	-0.000618	0.202826	0.202900	0.202826	0.368298	0.368321190	0.368320	0.571336	0.571336	0.571434	-0.001035
1600	-0.000357446	-0.000889	0.202687	0.202762	0.202687	0.368269	0.368294507	0.368294	0.571622	0.571622	0.571719	-0.000558
1800	-0.000224261	-0.000017	0.202591	0.202665	0.202591	0.368206	0.368233538	0.368234	0.571818	0.571818	0.571915	-0.000350
2000	-0.000126755	-0.000282	0.202521	0.202595	0.202521	0.368147	0.368174153	0.368174	0.571958	0.571958	0.572056	-0.000104

Table C.8: Relative Errors for option with $S_0 = 90$, $K = 100$, $T = 1$ in high volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.030261197	0.013604	0.016364	0.016466	0.016364	-0.010294	-0.011378959	-0.011604	-0.064943	-0.064943	-0.064715	-0.036723
400	-0.017642221	0.006667	0.009387	0.009504	0.009387	-0.000199	-0.000630873	-0.000765	-0.037346	-0.037346	-0.037116	-0.015260
500	-0.011409966	0.003145	0.005868	0.005967	0.005868	0.002580	0.002456921	0.002418	-0.024460	-0.024460	-0.024230	-0.014905
600	-0.008021077	0.004166	0.004116	0.004221	0.004116	0.003037	0.003013999	0.003020	-0.017479	-0.017479	-0.017248	-0.009904
700	-0.005996448	0.002733	0.003078	0.003183	0.003078	0.003399	0.003389885	0.003377	-0.013263	-0.013263	-0.013033	-0.005980
800	-0.004596355	0.001707	0.002349	0.002453	0.002349	0.002967	0.002975688	0.002965	-0.010523	-0.010523	-0.010292	-0.005247
900	-0.003654364	0.000998	0.001870	0.001975	0.001870	0.002874	0.002901735	0.002900	-0.008646	-0.008646	-0.008416	-0.003852
1200	-0.002084308	0.001116	0.001065	0.001171	0.001065	0.002599	0.002639996	0.002639	-0.005551	-0.005551	-0.005320	-0.002794
1400	-0.001538110	0.000458	0.000803	0.000909	0.000803	0.002371	0.002408718	0.002406	-0.004494	-0.004494	-0.004263	-0.002235
1600	-0.001193561	-0.000020	0.000622	0.000727	0.000622	0.002328	0.002372291	0.002372	-0.003811	-0.003811	-0.003580	-0.001734
1800	-0.000952962	0.000552	0.000502	0.000607	0.000502	0.002300	0.002348740	0.002350	-0.003341	-0.003341	-0.003110	-0.001446
2000	-0.000777810	0.000187	0.000414	0.000519	0.000414	0.002213	0.002259086	0.002258	-0.003004	-0.003004	-0.002773	-0.001216

Table C.9: Relative Errors for option with $S_0 = 100$, $K = 100$, $T = 3$ in high volatility.

p	$\tilde{v}$	Δ_{SDP}	Δ_{CD}	Δ_{Rich}	Δ_{mean}	Γ_{CD}	$\tilde{\Gamma}$	Γ_{Rich}	Θ_{CD}	Θ_{FD}	Θ_{Rich}	$\mathcal{V}$
300	-0.043033684	0.014189	0.019402	0.019524	0.019402	-0.009649	-0.010038	-0.010050	-0.072524	-0.072524	-0.072306	-0.043397675
400	-0.024981049	0.006161	0.010691	0.010824	0.010691	0.002802	0.002969	0.003068	-0.041513	-0.041513	-0.041294	-0.017902069
500	-0.016091917	0.003097	0.006687	0.006812	0.006687	0.000059	0.000001	0.000003	-0.026945	-0.026945	-0.026725	-0.017911688
600	-0.011291194	0.001651	0.004597	0.004723	0.004597	0.001426	0.001419	0.001420	-0.018995	-0.018995	-0.018774	-0.011579070
700	-0.008425279	0.000835	0.003323	0.003446	0.003323	0.002898	0.002947	0.002953	-0.014185	-0.014185	-0.013963	-0.007035548
800	-0.006452472	0.000401	0.002547	0.002673	0.002547	0.002497	0.002536	0.002543	-0.011052	-0.011052	-0.010830	-0.006137063
900	-0.005124506	0.000117	0.001999	0.002123	0.001999	0.002483	0.002507	0.002499	-0.008904	-0.008904	-0.008682	-0.004463227
1200	-0.002911050	-0.000232	0.001118	0.001244	0.001118	0.002442	0.002489	0.002490	-0.005360	-0.005360	-0.005138	-0.003068265
1400	-0.002146504	-0.000296	0.000826	0.000952	0.000826	0.002476	0.002527	0.002526	-0.004146	-0.004146	-0.003924	-0.002336496
1600	-0.001659509	-0.000326	0.000626	0.000751	0.000626	0.002414	0.002464	0.002462	-0.003363	-0.003363	-0.003141	-0.001785200
1800	-0.001321205	-0.000327	0.000492	0.000617	0.000492	0.002417	0.002470	0.002469	-0.002824	-0.002824	-0.002602	-0.001423075
2000	-0.001075053	-0.000315	0.000398	0.000524	0.000398	0.002359	0.002413	0.002413	-0.002438	-0.002438	-0.002216	-0.001188706

Table C.10: Relative Errors for option with $S_0 = 115$, $K = 100$, $T = 3$ in high volatility.

Appendix D

Appendix for Chapter 3

Figure D.1 provides a visual representation of the relationship between the dependent and independent variable Y and X defined in Equations (3.2.6)-(3.2.7) for the OLS regression performed in Section (3.2.2). For each subsample we present two plots: one that is computed using the initial estimator for volatility σ_0 (panel (a)), and one at the last iteration, computed with the optimized volatility (panel (b)). For simplicity, options with $30 < \text{DTM} < 90$ are labeled A, those with $90 < \text{DTM} < 180$ are B and those with $\text{DTM} > 180$ are C.

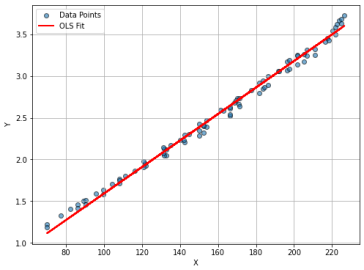

$S0/K < 0.0975$ A (a)

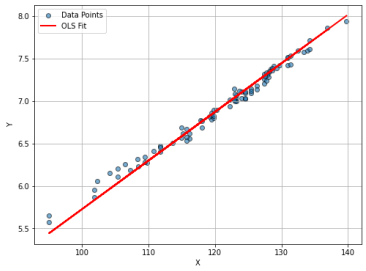

$S0/K < 0.0975$ A (b)

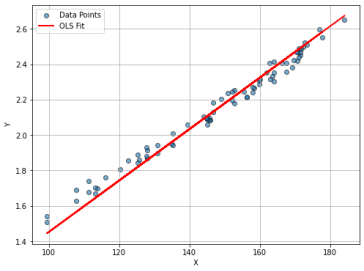

$S0/K < 0.0975$ B (a)

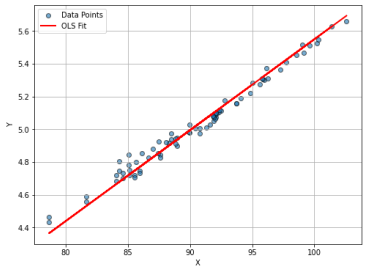

$S0/K < 0.0975$ B (b)

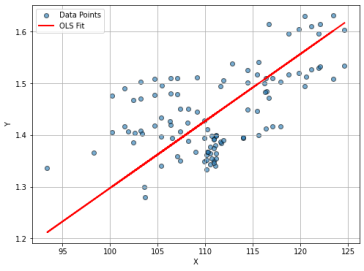

$S0/K < 0.0975$ C - (a)

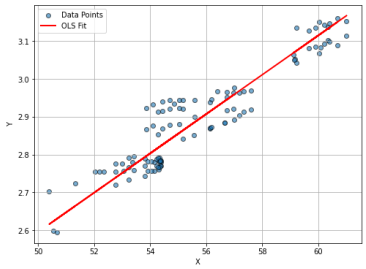

$S0/K < 0.0975$ C - (b)

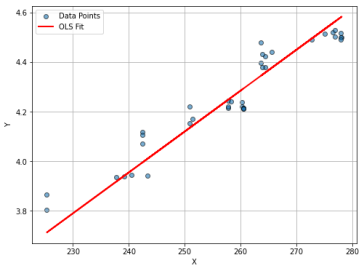

$0.0975 < S0/K < 1$ A - (a)

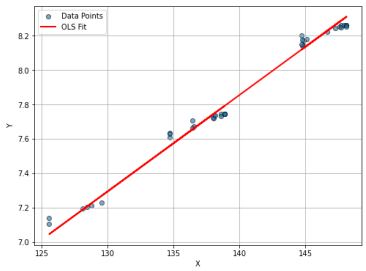

$0.0975 < S0/K < 1$ A - (b)

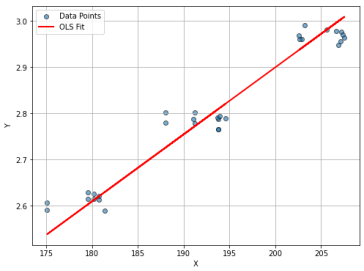

$0.0975 < S0/K < 1$ B - (a)

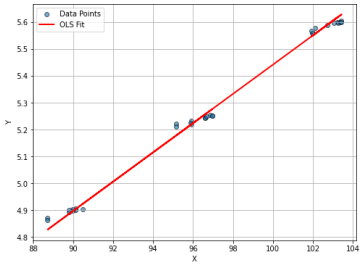

$0.0975 < S0/K < 1$ B - (b)

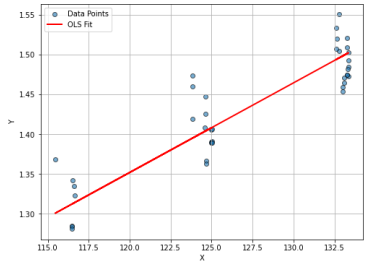

$0.0975 < S0/K < 1$ C - (a)

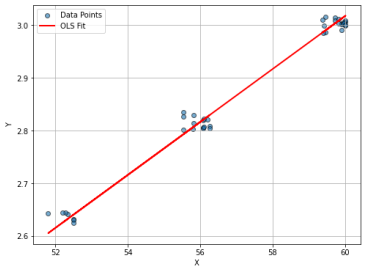

$0.0975 < S0/K < 1$ C - (b)

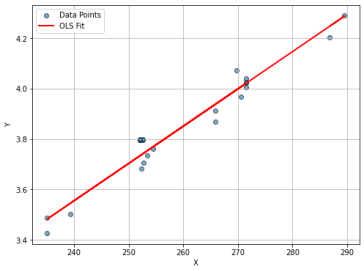

$1 < S0/K < 1.025$ A - (a)

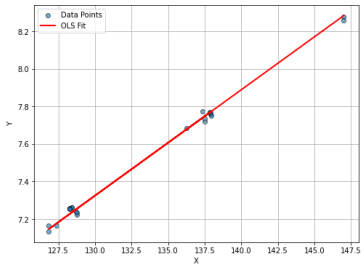

$1 < S0/K < 1.025$ A - (b)

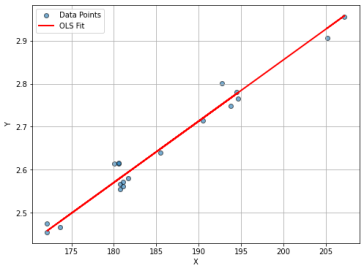

$1 < S0/K < 1.025$ B - (a)

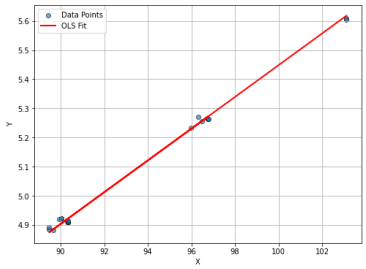

$1 < S0/K < 1.025$ B - (b)

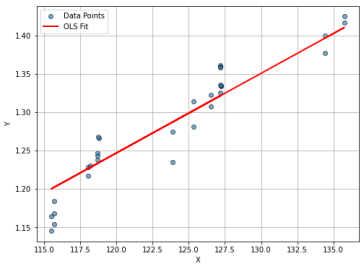

$1 < S0/K < 1.025$ C - (a)

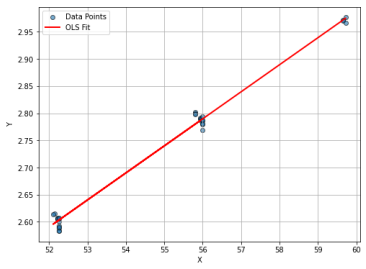

$1 < S0/K < 1.025$ C - (b)

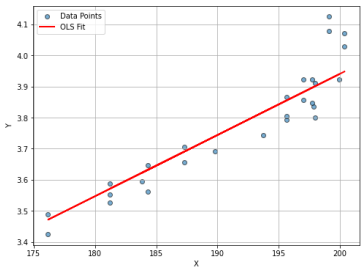

$1.025 < S0/K < 1.05$ A - (a)

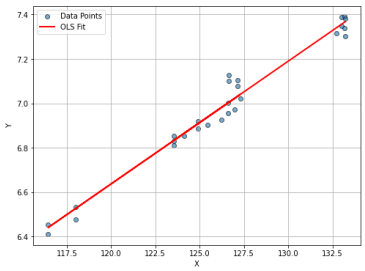

$1.025 < S0/K < 1.05$ A - (b)

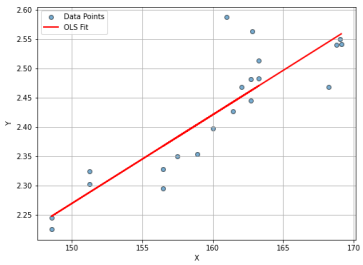

$1.025 < S0/K < 1.05$ B - (a)

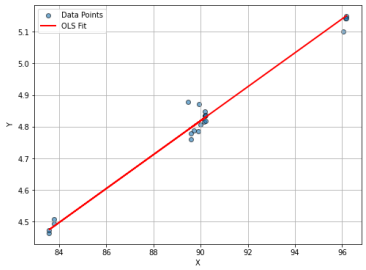

$1.025 < S0/K < 1.05$ B - (b)

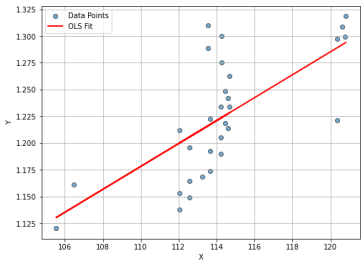

$1.025 < S0/K < 1.05$ C - (a)

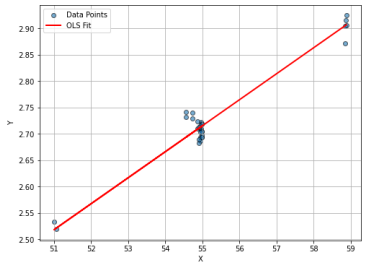

$1.025 < S0/K < 1.05$ C - (b)

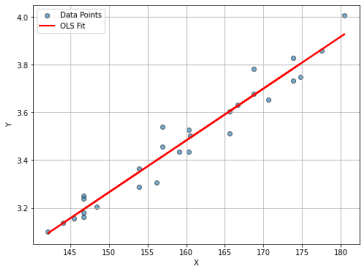

$1.05 < S0/K < 1.075$ A - (a)

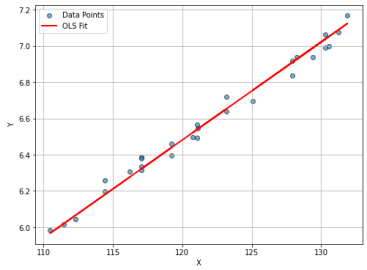

$1.05 < S0/K < 1.075$ A - (b)

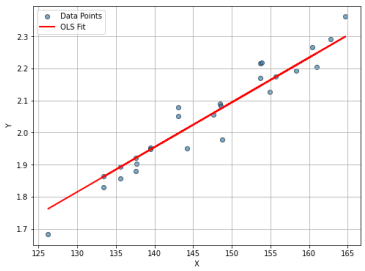

$1.05 < S0/K < 1.075$ B - (a)

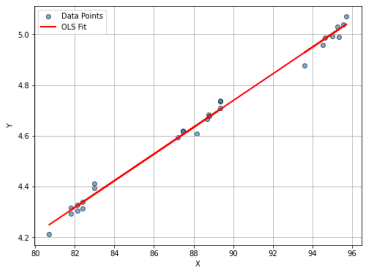

$1.05 < S0/K < 1.075$ B - (b)

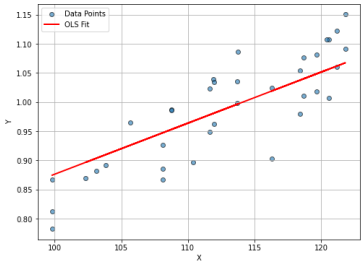

$1.05 < S0/K < 1.075$ C - (a)

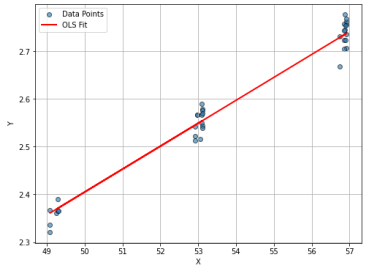

$1.05 < S0/K < 1.075$ C - (b)

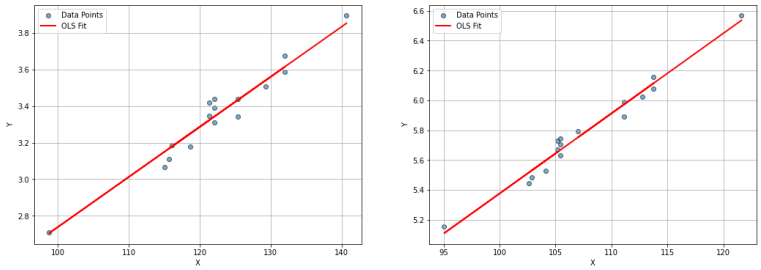

$S0/K > 1.075$ A - (a) $S0/K > 1.075$ A - (b)

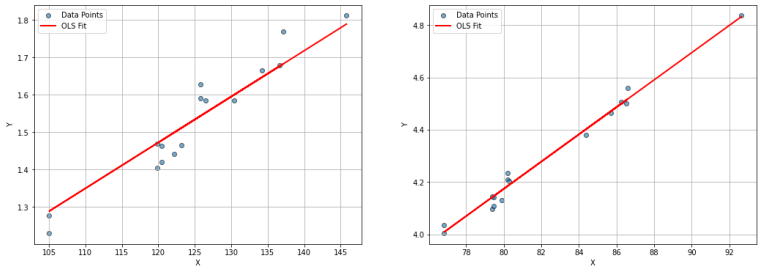

$S0/K > 1.075$ B - (a) $S0/K > 1.075$ B - (b)

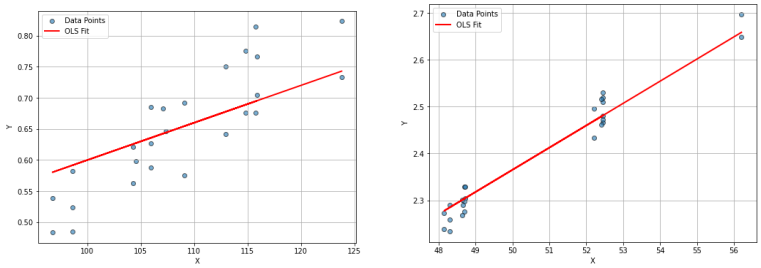

$S0/K > 1.075$ C - (a) $S0/K > 1.075$ C - (b)

Figure D.1: OLS regression fit for call options.

Figure D.2 illustrates the RMSE for the L-BFGS-B method method outlined in Section (3.2.1) across different samples. The figure is organized into six rows, each corresponding to a distinct moneyness interval and three panels, labeled (a), (b), and (c), which represent different ranges of time to maturity. Panels (a) correspond to samples with $30 < \text{DTM} < 90$, panels (b) corresponds to samples with $90 < \text{DTM} < 180$, and panels (c) corresponds to samples with $\text{DTM} > 180$.

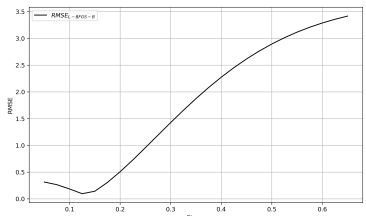

$S_0/K < 0.975$ - (a)

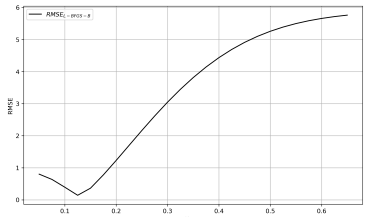

$S_0/K < 0.975$ - (b)

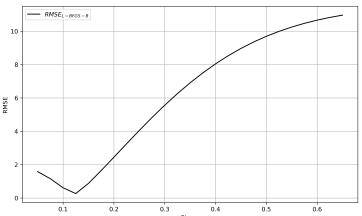

$S_0/K < 0.975$ - (c)

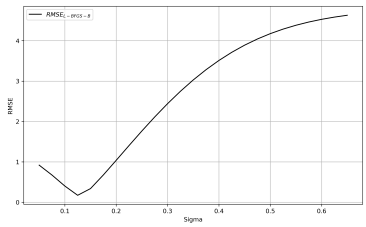

$0.975 < S_0/K < 1$ - (a)

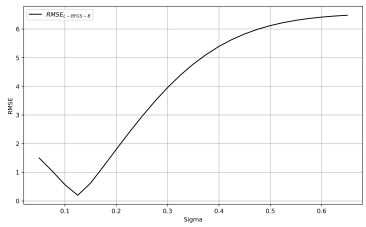

$0.975 < S_0/K < 1$ - (b)

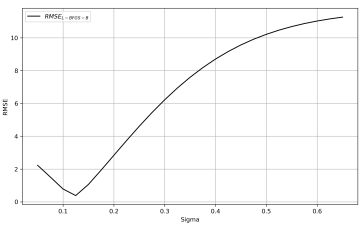

$0.975 < S_0/K < 1$ - (c)

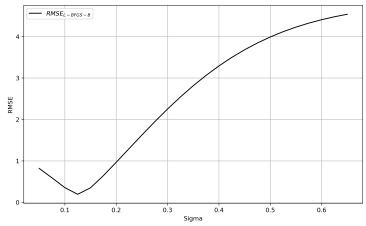

$1 < S_0/K < 1.025$ - (a)

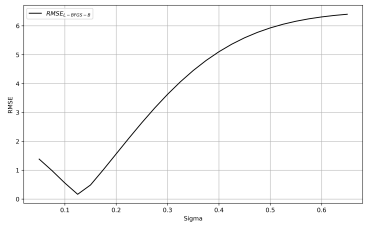

$1 < S_0/K < 1.025$ - (b)

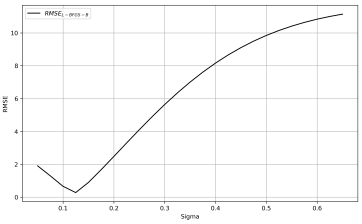

$1 < S_0/K < 1.025$ - (c)

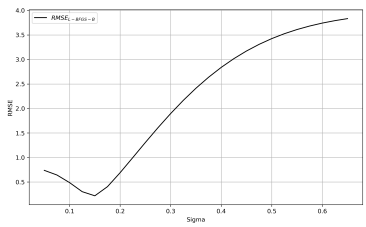

$1.025 < S_0/K < 1.05$ - (a)

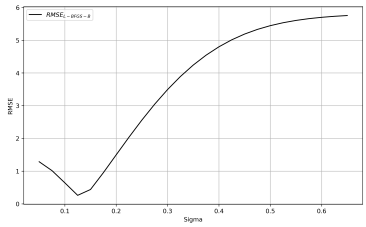

$1.025 < S_0/K < 1.05$ - (b)

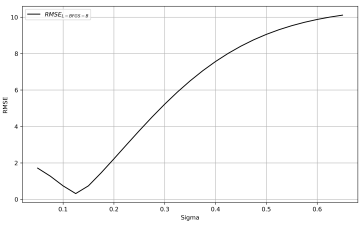

$1.025 < S_0/K < 1.05$ - (c)

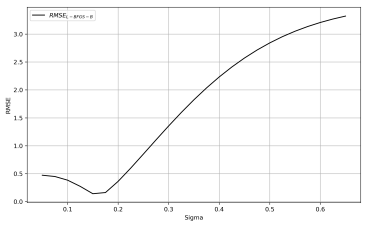

$1.05 < S_0/K < 1.075$ - (a)

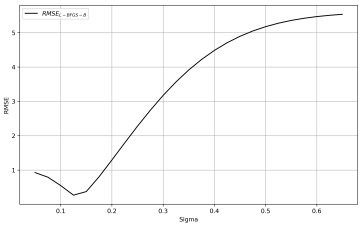

$1.05 < S_0/K < 1.075$ - (b)

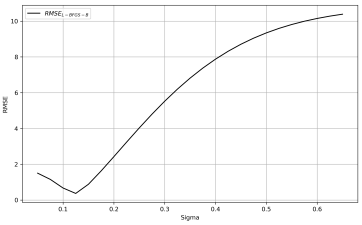

$1.05 < S_0/K < 1.075$ - (c)

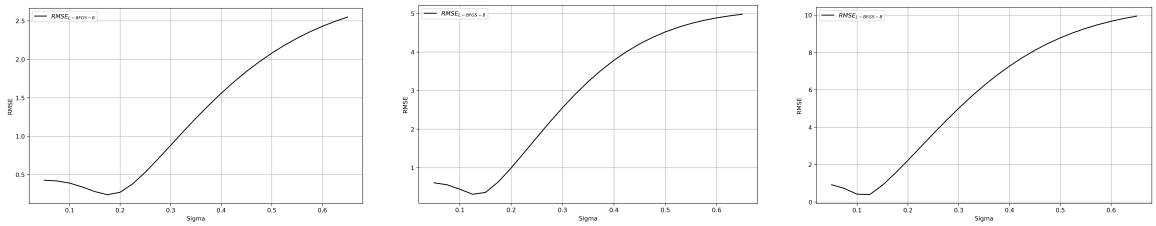

$S_0/K > 1.075$ - (a) $S_0/K > 1.075$ - (b) $S_0/K > 1.075$ - (c)

Figure D.2: RMSE across volatility values for L-BFGS-B method.

Figure D.3 illustrates the RMSE for the FK method method outlined in Section (3.2.2) across different samples. The figure is organized into six rows, each corresponding to a distinct moneyness interval and three panels, labeled (a), (b), and (c), which represent different ranges of time to maturity. Panels (a) correspond to samples with $30 < \text{DTM} < 90$, panels (b) corresponds to samples with $90 < \text{DTM} < 180$, and panels (c) corresponds to samples with $\text{DTM} > 180$.

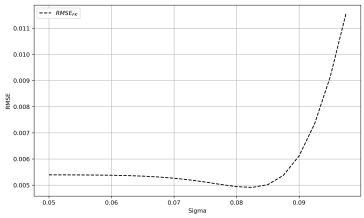

$S_0/K < 0.975$ - (a)

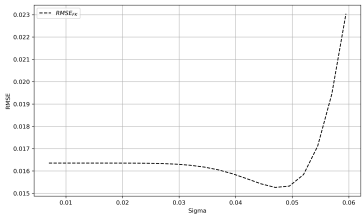

$S_0/K < 0.975$ - (b)

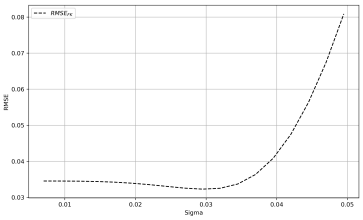

$S_0/K < 0.975$ - (c)

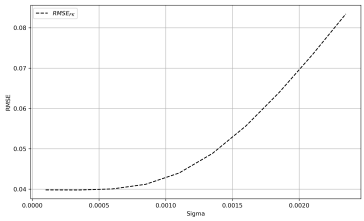

$0.975 < S_0/K < 1$ - (a)

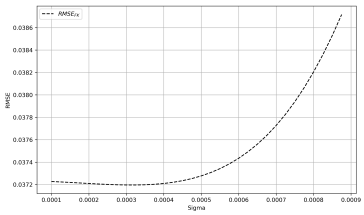

$0.975 < S_0/K < 1$ - (b)

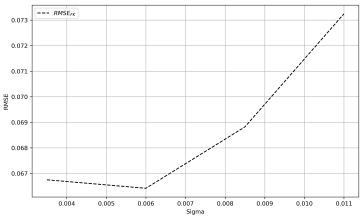

$0.975 < S_0/K < 1$ - (c)

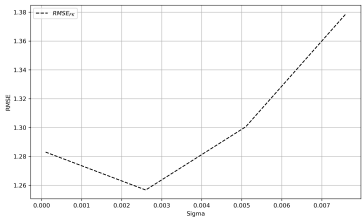

$1 < S_0/K < 1.025$ - (a)

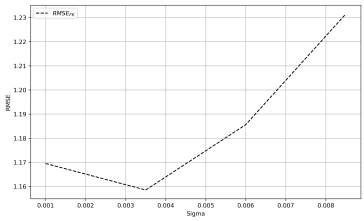

$1 < S_0/K < 1.025$ - (b)

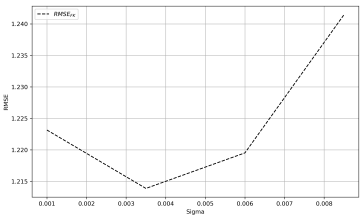

$1 < S_0/K < 1.025$ - (c)

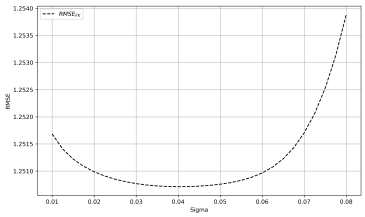

$1.025 < S_0/K < 1.05$ - (a)

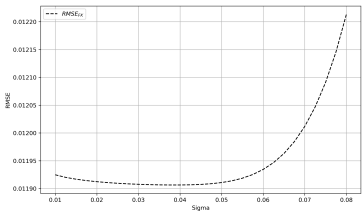

$1.025 < S_0/K < 1.05$ - (b)

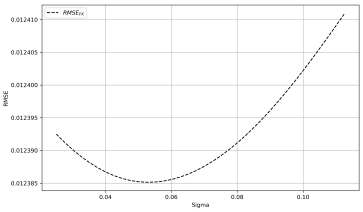

$1.025 < S_0/K < 1.05$ - (c)

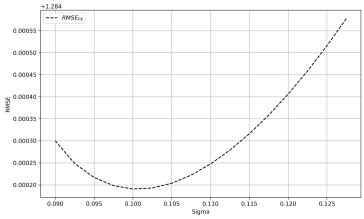

$1.05 < S_0/K < 1.075$ - (a)

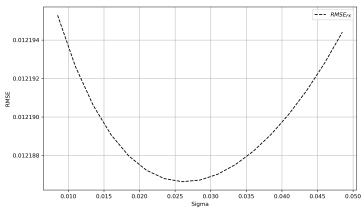

$1.05 < S_0/K < 1.075$ - (b)

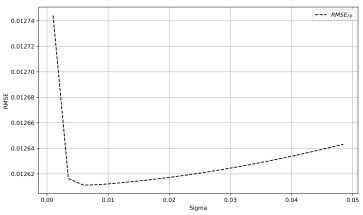

$1.05 < S_0/K < 1.075$ - (c)

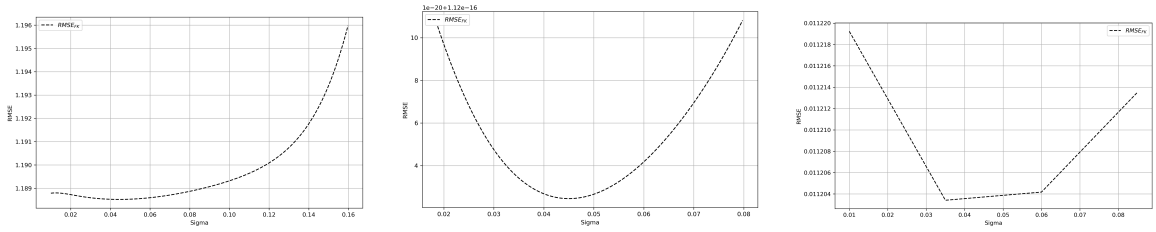

$S_0/K > 1.075$ - (a) $S_0/K > 1.075$ - (b) $S_0/K > 1.075$ - (c)

Figure D.3: RMSE across volatility values for FK method.

Figure D.4 provides the implied Gamma surfaces using the initial volatility approximation (panel (a)) defined in Section (3.1), as well as the optimal IV obtained with the L-BFGS-B method (panel (b)) and the FK method (panel (c)) described both in Section (3.2).

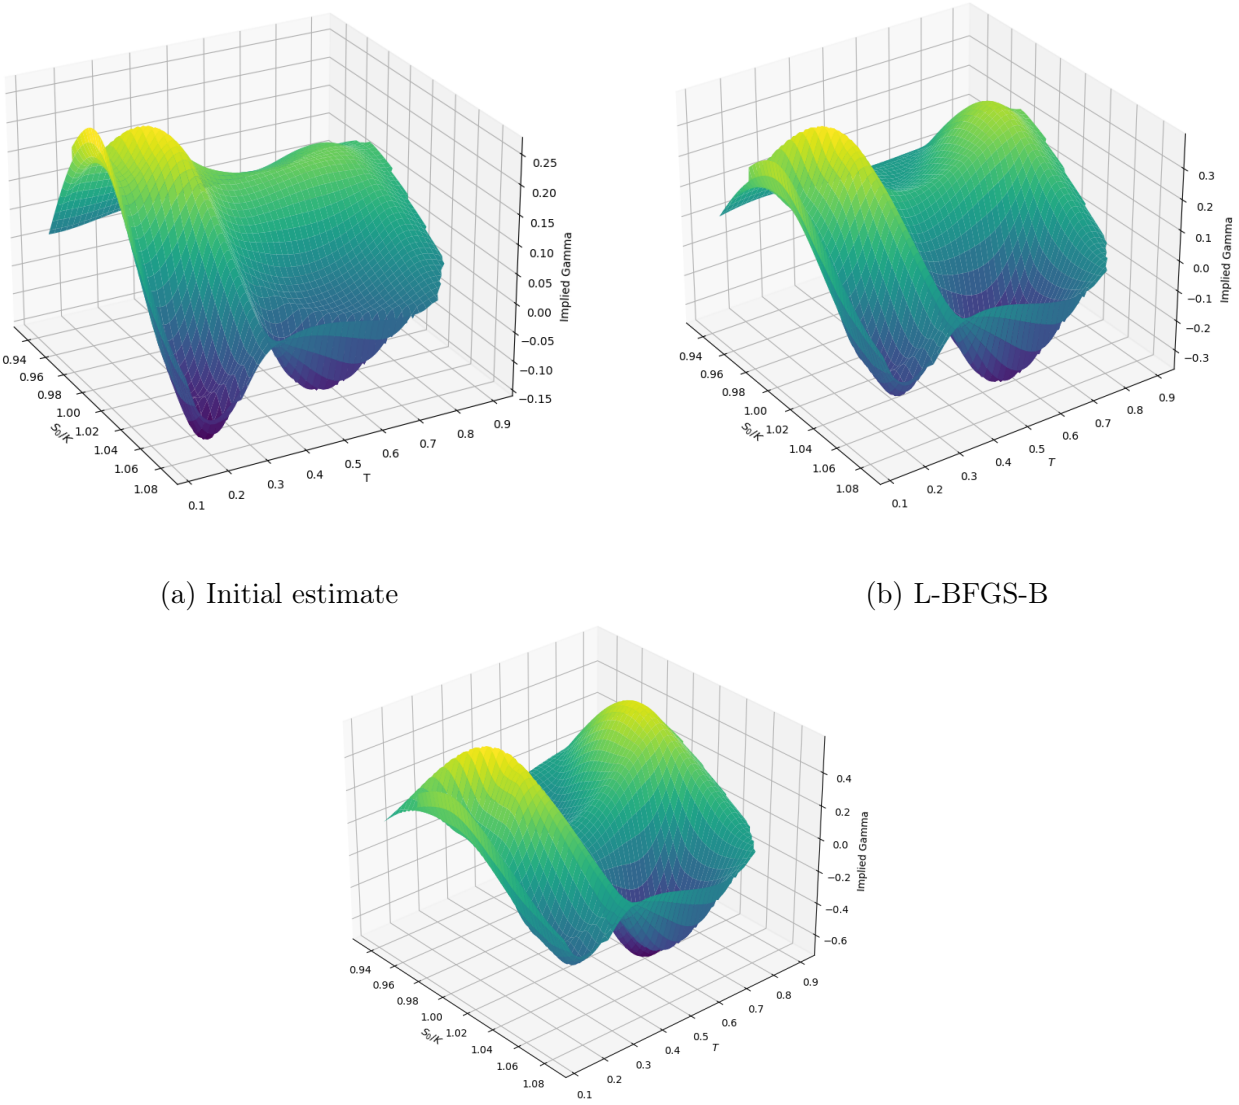

(a) Initial estimate

(b) L-BFGS-B

(c) FK

Figure D.4: Implied Gamma surfaces.

Figure D.5 provides the implied Theta surface using the initial volatility approximation (panel (a)) defined in Section (3.1), as well as the optimal IV obtained with the L-BFGS-B method (panel (b)) and the FK method (panel (c)) described both in Section (3.2).

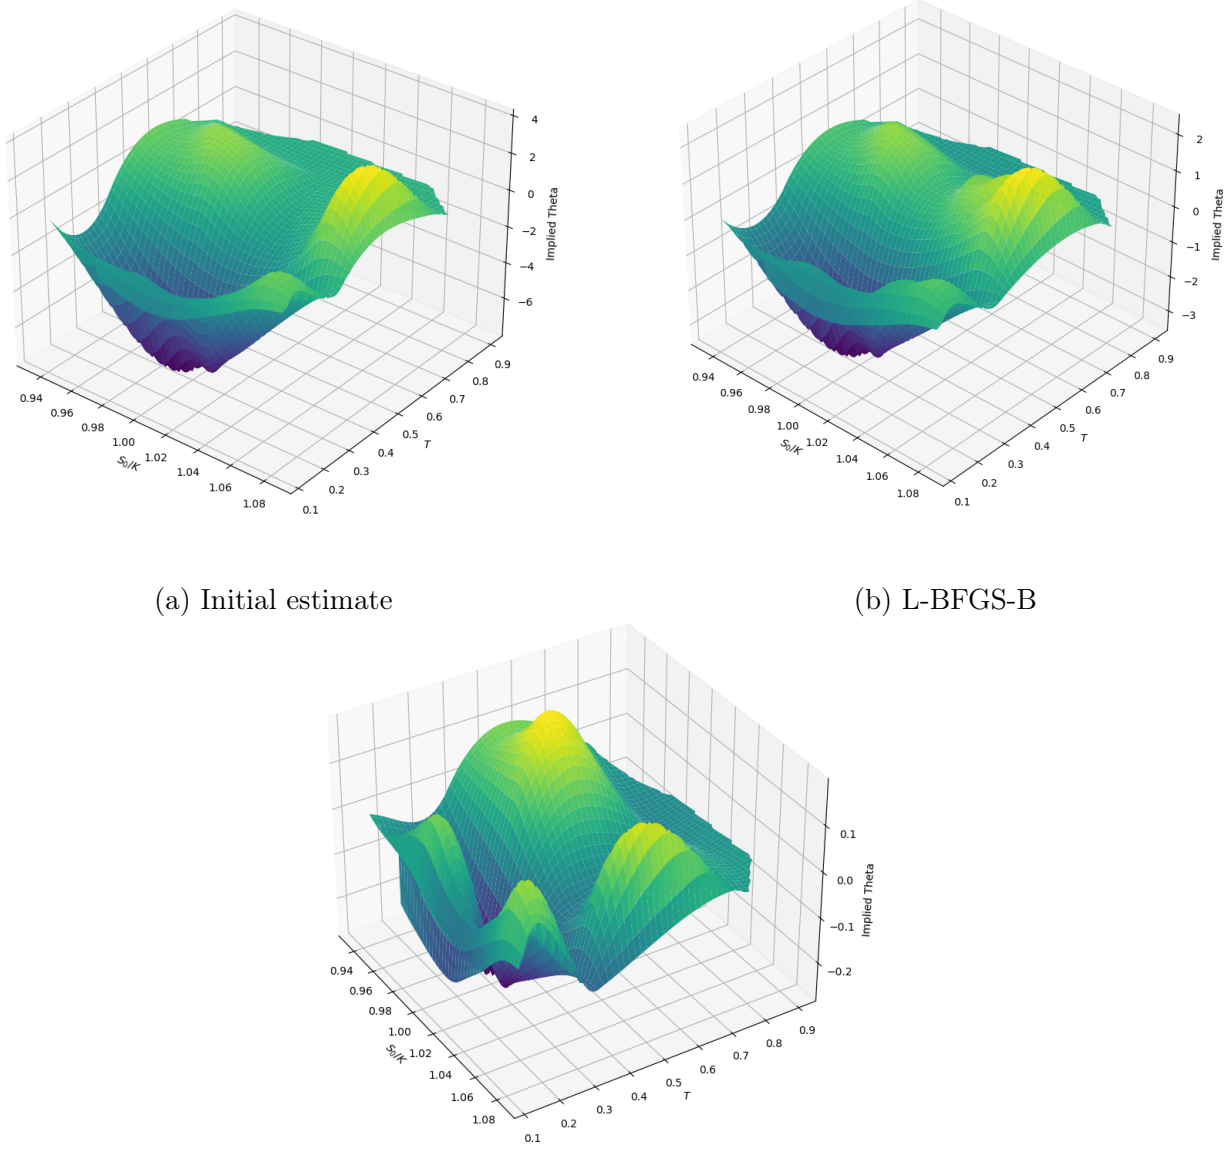

(a) Initial estimate

(b) L-BFGS-B

(c) FK

Figure D.5: Implied Theta surfaces.

Table D.3 provides the p-values for the BP test conducted to test for the presence of heteroskedasticity in the samples defined in Table 3.1. The BP test is a statistical test used to check for heteroskedasticity in a regression model. It examines whether the variance of residuals depends systematically on the independent variable(s).

The BP test is based on the idea that if heteroskedasticity is present, the variance of residuals should not be constant but should instead follow a function of the explanatory variables.

- H_0 : (Null Hypothesis - Homoskedasticity): The variance of residuals is constant, meaning it doesn't depend on X: $\mathbb{V}(\varepsilon|X) = \sigma^2$.
- H_1: (Alternative Hypothesis - Heteroskedasticity): The variance of residuals if a function of X: $\mathbb{V}(\varepsilon|X) = \sigma^2(X)$.

To perform the test, we first estimate the original regression model $Y = \beta_1 \times X + \varepsilon$, and obtain the residuals $\hat{\varepsilon}$. Since heteroskedasticity affects the variance of the errors, we compute the squared residuals $\hat{\varepsilon}^2$ and run an auxiliary regression where we regress $\hat{\varepsilon}^2$ on the explanatory variable X: $\hat{\varepsilon}^2 = \alpha_1 \times X + u$. From this regression, we extract the coefficient of determination R^2 and calculate the BP statistic $LM = nR^2$ where n is the number of observations. Under the null hypothesis of homoskedasticity, the test statistic follows a chi-square distribution $LM \sim \chi^2_k$, where k represents the number of independent variables in the auxiliary regression. To determine whether heteroskedasticity is present, we compute the p-value associated with the test statistic.

Moneyness	DTM	p-value
$S_0/K < 0.975$	$30 < \text{DTM} < 90$	8.041517×10^{-6}
	$90 < \text{DTM} < 180$	1.208104×10^{-5}
	$\text{DTM} > 180$	3.619665×10^{-10}
$0.975 < S_0/K < 1$	$30 < \text{DTM} < 90$	3.748659×10^{-5}
	$90 < \text{DTM} < 180$	5.531928×10^{-4}
	$\text{DTM} > 180$	2.309685×10^{-4}
$1 < S_0/K < 1.025$	$30 < \text{DTM} < 90$	3.863082×10^{-4}
	$90 < \text{DTM} < 180$	2.827093×10^{-3}
	$\text{DTM} > 180$	4.155405×10^{-4}
$1.025 < S_0/K < 1.05$	$30 < \text{DTM} < 90$	8.196686×10^{-3}
	$90 < \text{DTM} < 180$	4.098168×10^{-2}
	$\text{DTM} > 180$	1.555547×10^{-3}
$1.05 < S_0/K < 1.075$	$30 < \text{DTM} < 90$	3.110016×10^{-4}
	$90 < \text{DTM} < 180$	1.951333×10^{-3}
	$\text{DTM} > 180$	1.509491×10^{-4}
$S_0/K > 1.075$	$30 < \text{DTM} < 90$	1.883280×10^{-3}
	$90 < \text{DTM} < 180$	9.525762×10^{-3}
	$\text{DTM} > 180$	3.201966×10^{-4}

Table D.3: Breusch-Pegan test p-values.

Table D.3 reveals that in all subsamples $p < 0.05$, we reject H_0 and conclude that heteroskedasticity is present.

Bibliography

Abramowitz, M. and Stegun, I. A. (1948). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, volume 55. US Government printing office.

Aitken, A. C. (1936). On least squares and linear combination of observations. *Proceedings of the Royal Society of Edinburgh*, 55:42–48.

Bakshi, G., Cao, C., and Chen, Z. (1997). Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5):2003–2049.

Ball, C. A. and Roma, A. (1994). Stochastic volatility option pricing. *Journal of Financial and Quantitative Analysis*, 29(4):589–607.

Barone-Adesi, G. and Whaley, R. E. (1987). Efficient analytic approximation of American option values. *The Journal of Finance*, 42(2):301–320.

Bayer, C., Eigel, M., Sallandt, L., and Trunschke, P. (2023). Pricing high-dimensional bermudan options with hierarchical tensor formats. *SIAM Journal on Financial Mathematics*, 14(2):383–406.

Bellman, R. (1957). A markovian decision process. *Journal of Mathematics and Mechanics*, pages 679–684.

Ben-Abdellatif, M., Ben-Ameur, H., Chérif, R., and Rémillard, B. (2024). Dynamic programming for designing and valuing two-dimensional financial derivatives. *Risks*, 12(12):183.

Ben-Ameur, H., Breton, M., and L'Écuyer, P. (2002). A dynamic programming procedure for pricing American-style asian options. *Management Science*, 48(5):625–643.

Ben-Ameur, H., de Frutos, J., Fakhfakh, T., and Diaby, V. (2013). Upper and lower bounds for convex value functions of derivative contracts. *Economic Modelling*, 34:69–75.

Biagini, F., Hu, Y., Øksendal, B., and Zhang, T. (2008). *Stochastic calculus for fractional Brownian motion and applications*. Springer Science & Business Media.

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

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity (GARCH). *Journal of Econometrics*, 31(3):307–327.

Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics*, 4(3):323–338.

Brennan, M. J. and Schwartz, E. S. (1978). Finite difference methods and jump processes arising in the pricing of contingent claims: A Synthesis. *Journal of Financial and Quantitative Analysis*, 13(3):461–474.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.

Chen, A. H. (1970). A model of warrant pricing in a dynamic market. *The Journal of Finance*, 25(5):1041–1059.

Christoffersen, P., Heston, S., and Jacobs, K. (2009). The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Management Science*, 55(12):1914–1932.

Cosma, A., Galluccio, S., Pederzoli, P., and Scaillet, O. (2020). Early exercise decision in American options with dividends, stochastic volatility, and jumps. *Journal of Financial and Quantitative Analysis*, 55(1):331–356.

Cox, J. C., Ingersoll Jr, J. E., and Ross, S. A. (2005). A theory of the term structure of interest rates. In *Theory of valuation*, pages 129–164. World Scientific.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263.

Geske, R. and Johnson, H. E. (1984). The American put option valued analytically. *The Journal of Finance*, 39(5):1511–1524.

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

Huber, P. J. (1992). Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pages 492–518. Springer.

Hull, J. et al. (2009). *Options, Futures and Other Derivatives/John C. Hull*. Prentice Hall.

Hull, J. and White, A. (1990). Valuing derivative securities using the explicit finite difference method. *Journal of Financial and Quantitative Analysis*, 25(1):87–100.

Joshi, M. (2001). Log-type models, homogeneity of option prices and convexity. *QUARC, Royal Bank of Scotland Working Paper*.

Karatzas, I. and Shreve, S. E. (1998). *Methods of Mathematical Finance*, volume 39. Springer.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101.

Létourneau, P. and Stentoft, L. (2022). Efficient pricing of large panels of options. *Available at SSRN 4094237*.

Longstaff, F. A. and Schwartz, E. S. (2001). Valuing American options by simulation: a simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147.

MacMillan, L. W. (1986). Analytic approximation for the American put option. *Advances in Futures and Options Research*, 1(1):119–139.

Madan, D. B., Carr, P. P., and Chang, E. C. (1998). The variance gamma process and option pricing. *Review of Finance*, 2(1):79–105.

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

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144.

Oksendal, B. (2013). *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media.

Richardson, L. F. (1911). IX. the approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 210(459-470):307–357.

Saad, Y. (2003). *Iterative methods for sparse linear systems*. SIAM.

Tilley, J. A. (1993). Valuing American options in a path simulation model. *Transactions of the Society of Actuaries*, 45:499–519.

Tukey, J. W. et al. (1977). *Exploratory data analysis*, volume 2. Springer.

Wilmott, P., Howison, S., and Dewynne, J. (1995). *The Mathematics of Financial Derivatives: A student introduction*. Cambridge University Press.

Wooldridge, J. M. (2016). *Introductory Econometrics: A Modern Approach*. Cengage Learning, 6th edition.

Young, D. (1954). Iterative methods for solving partial difference equations of elliptic type. *Transactions of the American Mathematical Society*, 76(1):92–111.