Analysis on the Convergence of the Least Squares Monte Carlo Method and the Finite Differences Method on the Valuation of American Options

By

Teddy K. Waweru

Submitted to the Graduate School in Partial Fulfilment of the Requirements for the Degree of Master of Science in Mathematical Finance and Risk Analytics at Strathmore University

Institute of Mathematical Sciences
Strathmore University
Nairobi, Kenya

June 2025

Declaration and Recommendation

I declare that this work has not been previously submitted and approved for the award of a degree by this or any other University. To the best of my knowledge and belief, the dissertation contains no material previously published or written by another person except where due reference is made in the dissertation itself.

Signature: .... Date: 25 May 2025

Teddy K. Waweru

Approval

The dissertation of Teddy K. Waweru was reviewed and approved by the following:

Dr. Samuel Chege Maina,
Senior Lecturer, School of Humanities and Social Sciences,
Strathmore University

Prof. Godfrey Madigu,
Dean, School of Humanities and Social Sciences,
Strathmore University

Prof. Bernard Shibwabo,
Director of Graduate Studies,
Strathmore University

Abstract

American options, which permit early exercise, pose considerable challenges in financial valuation due to the intricacies involved in identifying optimal exercise strategies. Conventional analytical models are often inadequate for these options, particularly in high-dimensional contexts, thereby necessitating the use of robust numerical techniques such as the Least Squares Monte Carlo (LSM) and Finite Difference (FD) methods.

The LSM method is distinguished by its adaptability in addressing high-dimensional issues, employing a combination of Monte Carlo simulations and regression analysis to estimate continuation values, which renders it particularly effective for complex, path-dependent options. Conversely, the FD method utilizes a grid-based framework to solve the partial differential equations (PDEs) that dictate option pricing, providing stable and dependable solutions, especially in one-dimensional scenarios.

This study compares two numerical methods—Least Squares Monte Carlo and Finite Difference Methods —for pricing American options under a *SABR-based* volatility calibration. LSMC offers flexibility and moderate runtime in higher-dimensional contexts, while FDM provides a deterministic, systematically refined solution. Both approaches are benchmarked against binomial-lumps and Bjerksund–Stensland references for a short-dated American call, with emphasis on computational efficiency, accuracy, and the practical implications of adopting a lognormal SABR model to ensure market-consistent volatility.

We analyze convergence rates, computational runtimes, and pricing accuracy across varying grid resolutions, simulation sizes, and basis selections. Our results delineate the trade-offs between flexibility, dimensional scalability, numerical stability and model run time, offering investors clear guidelines on choosing the optimal technique, LSMC for high-dimensional, path-dependent payoffs and FDM for one-dimensional problems demanding tight error control under SABR-driven market dynamics.

Keywords: American options, Implied Volatility (IV), Least Squares Monte Carlo (LSMC), Finite Differences Methods (FDM), Stochastic Alpha, Beta, Rho (SABR) Model, Control Variate (CV), Antithetic Variate (AV)

Table of Contents

List of Figures

List of Tables

Acknowledgement

I would like to thank my supervisor, Dr. Samuel Chege Maina, for his insightful guidance and unwavering support; the Institute of Mathematical Sciences at Strathmore University for providing access to resources and a stimulating research environment; my classmates and work colleagues for their camaraderie, peer review, and encouragement; and my family for their patience and belief in me. Above all, my greatest appreciation goes to God Almighty for all the favour and peace He has bestowed upon me throughout my life at Strathmore University.

Abbreviations

IV	Implied Volatility
LSMC	Least Squares Monte Carlo
FDM	Finite Differences Method
SABR	Stochastic Alpha, Beta, Rho
CV	Control Variate
AV	Antithetic Variate
BS	Bjerksund–Stensland

Chapter One
Introduction

1.1 Background of the study

In the financial markets, options are widely used derivative instruments that allow holders to benefit from market changes without needing direct exposure to the underlying asset as described by Wilmott (1995). Among the different types of options, American options stand out for their flexibility, as they can be exercised at any time before expiration as shown by Brennan (1977). This feature increases their value for the holder but also complicates the pricing process, requiring the determination of an optimal exercise strategy.

Valuing American options presents significant computational challenges, mainly due to their early exercise feature, which becomes more complex when dealing with high-dimensional or path-dependent contracts. Traditional analytical methods such as the Black Scholes framework work well for European options but fall short for American styles, as they do not consider the dynamic nature of the exercise boundary. As a result, numerical methods, particularly the Least Squares Monte Carlo (LSM) and Finite Difference (FD) techniques, become essential tools for accurately valuing American options, depicted by Longstaff (2001) and Wilmott (1995).

Longstaff (2001) introduced a Monte Carlo simulation framework in conjunction with regression analysis to approximate continuation values, rendering it particularly effective for intricate payoffs characterized by multiple sources of uncertainty. In contrast, the finite-difference method utilizes a grid-based framework to resolve the partial differential equation governing option pricing by discretizing both time and asset-price variables; this approach is especially proficient for one-dimensional scenarios and has been widely applied in option pricing, yielding stable and dependable estimates within controlled environments, Wilmott (1995).

Accurate pricing of American-style options is crucial in financial engineering. Finite Difference Methods have long been employed for single-asset problems but can become computationally expensive at high resolutions or for path-dependent features, Wilmott (1995). In contrast, Monte Carlo approaches such as LSMC can handle more complex payoffs but require careful variance reduction and basis function selection to achieve precision. Meanwhile, the SABR model offers a robust framework for matching implied-volatility smiles, providing a more realistic volatility input than constant-volatility assumptions as depicted by Hagan (2002). This study integrates SABR calibration into both LSMC and FDM, enabling a direct comparison of their performances for American-option pricing.

1.2 Statement of the problem

Valuing American options poses considerable challenges mainly because of the option's early exercise feature, which requires determining an optimal stopping rule. Analytical methods are typically impractical, especially for options that are high-dimensional or path-dependent. Consequently, numerical techniques like Least Squares Monte Carlo (LSM) and Finite Difference (FD) methods have become vital in this field. However, these methods come with certain trade-offs:

1. Least Squares Monte Carlo (LSM): This approach is beneficial due to its flexibility and ability to handle high-dimensional cases; however, it often requires significant computational power. The dependence on regression for estimating continuation values can result in inaccuracies, particularly with complex or non-linear payoff structures.
2. Finite Difference (FD) Method: The FD method is effective for one-dimensional problems and generally produces stable solutions. However, it struggles with high-dimensional options and requires careful management of boundary conditions to maintain both stability and convergence.

Despite the popularity of FDM and LSMC, there is a need for a direct head-to-head comparison under a common volatility model, particularly for short-dated American calls with discrete dividend considerations. Practitioners often face uncertainty about the balance between speed and accuracy when choosing either approach, especially once SABR calibration overhead is included. Hence, a systematic evaluation of LSMC versus FDM under the same SABR-based volatility is warranted.

1.3 Research Objectives

The **main objective** of this study is to **evaluate and compare** the performance of LSMC and FDM in pricing *American-style* options under a *SABR-based* volatility calibration, focusing on accuracy, runtime, and feasibility.

Specific objectives are:

1. To **calibrate** the SABR (lognormal) model to market option quotes, yielding a single volatility for both LSMC and FDM.
2. To **implement** the LSMC algorithm with polynomial or Laguerre basis expansions and, where appropriate, variance-reduction techniques.
3. To **implement** a finite difference lumps approach for American calls, refining spatial and temporal grids to reach desired accuracy.
4. To **benchmark** both methods against binomial-lumps and Bjerksund–Stensland references for a short-dated American call.
5. To **analyze** computational efficiency, highlighting trade-offs between runtime and accuracy in each method.

1.4 Justification of the study

Due to the immense global exposure to derivative instruments and their potential to multiply the losses or profits of investors, the financial sector is fundamentally dependent on precise option pricing models for effective trading, risk management, and hedging strategies. American options are prevalent across various markets, including equities, commodities, and fixed income, due to their early exercise capabilities that enhance flexibility. Nevertheless, the valuation of these options poses significant challenges, necessitating the use of sophisticated numerical methods that can accommodate intricate boundary conditions and diverse sources of uncertainty.

Selecting the right numerical method can greatly affect the speed and reliability of American option pricing. With many institutions adopting SABR for implied volatility consistency, evaluating LSMC and FDM under the same SABR-based model clarifies their relative strengths and limitations. The insights derived enable more informed decisions on resource usage versus required precision, guiding practitioners and researchers toward the most suitable approach for specific American option problems, especially those involving short maturities, discrete dividends, or path-dependent features.

1.5 Expected Output

This research is expected to yield the following outputs:

1. **A comprehensive theoretical and empirical comparison** of the LSM and FD methods for American option valuation, highlighting their convergence rates, computational efficiency, and accuracy under various market scenarios.

2. **Insights into the strengths and limitations of each method**, with practical guidelines on selecting the appropriate method based on option characteristics, dimensionality, and market conditions.

3. **Empirical data and graphical representations** illustrating the performance of LSM and FD methods in terms of stability and accuracy, providing quantitative benchmarks for further research and practical implementation.

4. **Recommendations for financial practitioners** on the best practices for using LSM and FD methods in American option pricing, with considerations for high-dimensional options and complex derivatives.

5. **Identification of potential areas for further research** in option pricing, including the exploration of hybrid methods or enhancements to LSM and FD methods that may improve accuracy and efficiency in high-dimensional settings.

Chapter Two
Literature Review

2.1 Introduction

The valuation of American options, which permit early exercise, poses distinct challenges in mathematical finance due to their dependence on the path taken and their intricate pay off structures. Conventional techniques, such as binomial trees, frequently become impractical for high-dimensional scenarios. The Least Squares Monte Carlo (LSM) method, introduced by Longstaff and Schwartz in 2001, revolutionized this area by integrating regression methods with Monte Carlo simulations, thereby enabling the approximation of the optimal stopping rule for American options.

2.2 Foundational Work

Least Squares Monte Carlo Method

The study by Carriere (1996) is widely acknowledged as one of the first comprehensive dives into the use of regression analysis for pricing early-exercise options. Through the application of non-parametric regression, Carriere effectively estimated the continuation values associated with American options, thereby establishing a foundational framework that would inform future regression-based methodologies. His research not only highlighted the advantages of simulation in capturing the complexities of early exercise but also provided a credible alternative to traditional finite difference methods for the valuation of sophisticated derivatives. The insights gained from Carriere's work have had a lasting influence, inspiring subsequent innovations such as the Least Squares Monte Carlo (LSMC) method, which further demonstrates the efficacy of simulation-based techniques in the context of American option pricing. Overall, Carriere's contributions laid the groundwork for advancements in the understanding and application of regression methods in financial derivatives, marking a pivotal moment in the evolution of option pricing theory.

Broadie and Glasserman's (1997) paper introduction of the stochastic mesh method represents one of the landmark contributions to the pricing of high-dimensional American options through

Monte Carlo techniques. While this method does not directly correspond to the Least Squares Monte Carlo (LSMC) approach, it addresses similar challenges, particularly those associated with the complexities of high-dimensional derivatives. By employing a mesh-based framework, Broadie and Glasserman established critical performance benchmarks for Monte Carlo methodologies, which have been utilized in subsequent research to evaluate the efficiency and accuracy of LSMC. Their findings highlight the significance of alternative simulation strategies, thereby shaping the trajectory of LSMC's development into a more effective and practical solution for high-dimensional pricing challenges.

Longstaff and Schwartz (2001) revolutionized the valuation of American options through the introduction of the Least Squares Monte Carlo (LSM) method. This innovative approach employs least-squares regression on simulated paths to determine the continuation value at each decision juncture, thereby addressing the optimal stopping dilemma inherent in American options. The LSM methodology gained widespread acclaim for its flexibility, computational efficiency, and capacity to manage high-dimensional scenarios. By overcoming the constraints associated with lattice and finite difference techniques, LSM established a new benchmark in the pricing of American options, marking a significant progression in the domain of quantitative finance. In their study, Tsitsiklis and Van Roy (2001) enhanced the theoretical framework for regression-based option pricing by analyzing the implementation of regression techniques in estimating continuation values. Their investigation into the properties of convergence and the criteria for basis function selection offered important insights into the theoretical aspects of the Least Squares Monte Carlo (LSMC) method, demonstrating that regression-based techniques can achieve accurate results for complex American-style options. Their findings contributed to a broader understanding of regression methods and set the stage for future analytical research on the strengths and weaknesses of LSMC.

In their seminal paper, Clement, Lamberton, and Protter (2002) presented the first detailed convergence analysis of the Least Squares Monte Carlo (LSMC) technique, thereby establishing a solid theoretical basis for its application. Their analysis of factors influencing convergence, including the selection of basis functions and the methodology for sample path generation, cemented the robustness of the method while also revealing potential error sources. This research has become an essential reference for grasping the theoretical reliability of LSM, facilitating practitioners in applying the method with greater assurance in financial modelling contexts.

Expanding upon previous theoretical frameworks, Stentoft (2004a) conducted a comprehensive analysis of the convergence characteristics of the Least Squares Monte Carlo method. His research meticulously examined the nuances of LSMC, particularly in the context of valuing options that exhibit complex features or have extended maturities. Stentoft identified specific challenges that practitioners face when applying LSMC to such options, including the intricacies of path dependency where the value of the option is influenced by the entire trajectory of the underlying asset's price and the regression conditions that are crucial for accurately estimating the expected pay off. By providing empirical insights into the various variables that impact the efficacy of LSMC, his work not only elucidated the theoretical underpinnings of the method but also offered practical guidance for its application. This investigation played a pivotal role in enhancing the practical application of LSMC, assisting practitioners in configuring the method to attain optimal outcomes in various financial contexts, thereby bridging the gap between theory and practice in financial modelling. In a complementary paper, Stentoft (2004b) empirically evaluates the accuracy and computational efficiency of LSM, contrasting its performance with established benchmarks. His study underscores the limitations of LSM in certain contexts, notably during periods of high volatility or with long maturities and offers suggestions for practitioners aiming for more accurate valuations. By elucidating LSM's practical strengths and weaknesses, this paper provides vital insights that have contributed to the establishment of LSM as a reliable, yet contextually dependent, resource in quantitative finance.

Andersen and Broadie (2004) paper proposed a primal-dual framework that serves to augment the Least Squares Monte Carlo (LSM) method by offering enhanced bounds for American option valuations. Their approach, grounded in duality principles, results in more stringent price bounds, which is advantageous for multidimensional options that may pose convergence issues for LSM. This significant contribution has paved the way for hybrid approaches that merge LSM with primal-dual strategies, facilitating new avenues for addressing complex pricing scenarios in options. Egloff (2005) enhanced the Least Squares Monte Carlo (LSM) method by incorporating statistical learning techniques that improve its convergence and robustness. He addresses optimal stopping problems, proposing better strategies for selecting basis functions and managing path dependencies. His work is particularly impactful for high-dimensional American options, as his adaptive methods significantly increase LSM's accuracy, making it more relevant for complex derivatives in algorithmic trading and risk management.

Broadie and Cao (2008) developed advanced algorithms aimed at enhancing the lower and upper bounds in the valuation of American options. Their work specifically targets the biases and variability associated with the Longstaff-Schwartz Method (LSM). By improving the robustness of LSM, their methodologies render it more dependable for real-world applications. The findings of this study have gained significant importance for practitioners who are looking to address the shortcomings of LSM, particularly in the realms of risk management and portfolio optimization.

In his paper, Hull (2021) emphasized the pivotal role of implied volatilities in achieving market-consistent pricing, underscoring their value as a cornerstone of practical financial modeling. Hull highlights the necessity of aligning theoretical option pricing frameworks with market data to ensure applicability and reliability. He reinforces the enduring relevance of the Black-Scholes framework, particularly in its adaptability to serve as a foundation for numerical models used in modern financial environments.

Brandimarte (2018) complements this perspective by discussing the implementation challenges of Least Squares Monte Carlo (LSMC) techniques in numerical settings. He stresses the importance of balancing computational efficiency and accuracy, particularly in regression-based Monte Carlo methods. These insights emphasize the need for careful selection of basis functions and numerical parameters to achieve robust pricing outcomes, especially when dealing with complex financial derivatives.

Anurag Sodhi (2018) paper implements the Least Squares Monte Carlo (LSMC) framework of Longstaff and Schwartz (2001) for American put options within the Bakshi, Cao Chen (1997) modelâ€"which features stochastic volatility, stochastic interest rates, and jumpsâ€"and then investigates whether alternative regression methods can improve pricing accuracy.

Finite Differences Method

In their groundbreaking of Brennan and Schwartz (1977) emerged as pioneers in the field of financial modelling by applying finite difference methods specifically to the valuation of American put options. This was a significant achievement, as American options possess the unique characteristic of allowing holders to exercise their options at any point before expiration, which introduces complexities not present in European options. To tackle this early-exercise feature, they framed the problem as a free boundary problem, a sophisticated mathematical approach that involves determining the optimal exercise boundary where the option holder would choose to exercise the option rather than hold it. Their research provided compelling evidence that finite difference grids, which are numerical methods used to approximate solutions to differential equations, could yield both practical and highly accurate solutions for the valuation of American options. This work not only advanced the understanding of option pricing but also solidified finite difference methods as a fundamental and widely accepted approach in the realm of financial modelling, influencing subsequent research and applications in the field. In the subsequent research, Schwartz (1977) showcased the adaptability of finite difference techniques for pricing warrants, which are similar to options. His findings extended the finite difference methodology beyond typical American options, demonstrating its effectiveness in addressing complex derivative structures.

In their seminal work published by Geske and Johnson (1987) introduced a comprehensive analytical framework that adeptly incorporated finite difference methodologies specifically tailored for the valuation of American put options. This innovative research not only provided a robust theoretical foundation for understanding the pricing dynamics of American options but also substantiated the finite difference approach as a precise and dependable instrument in this domain. By effectively bridging the gap between analytical models and computational strategies, their study paved the way for more accurate and efficient pricing mechanisms, enhancing the tool kit available to financial analysts and practitioners in the field of options trading. In their 1990 study, Hull and White made significant strides in the application of finite difference methods by incorporating explicit approaches that optimized computational efficiency. They provided comprehensive guidelines for the practical implementation of these methods in derivative pricing, particularly for American options, which resulted in their widespread adoption in the finance industry.

Elliott and Mckinnon (1995) conducted a thorough examination of the stability and convergence properties of finite difference methods, which are essential for the effective pricing of financial derivatives. Their research focused on refining these methods to better accommodate sophisticated option structures, which often present unique challenges due to their complexity. The contributions made by Elliott and Mckinnon were instrumental in improving the numerical effectiveness of finite difference solutions, thereby enhancing their reliability in practical applications. This work has significantly supported accurate pricing in a variety of market conditions, allowing financial analysts and traders to make more informed decisions based on reliable numerical outputs. Their findings have thus played a crucial role in advancing the field of computational finance.

Gatheral's work (2006) on implied volatility surfaces introduced methodologies for constructing, validating, and smoothing these surfaces, which are critical for market-consistent pricing. His techniques directly influenced the formulation of volatility-dependent coefficients in PDEs, enabling their integration into FDM frameworks for robust and accurate option pricing. The research conducted by Ikonen and Toivanen (2009) unveiled an innovative operator splitting technique that serves to substantially improve the efficiency of finite difference methods specifically tailored for real-time pricing applications. By strategically decoupling the time-stepping mechanism from the management of boundary conditions, their approach achieves significant computational enhancements. This decoupling allows for a more efficient processing of data, which is crucial in high-frequency trading environments where milliseconds can make a difference. As a result, their method not only optimizes the performance of finite difference methods but also expands their applicability in the fast-paced world of financial trading.

Rudiger Seydel's (2015) work examines the complexities of American options, where early exercise creates distinct continuation and stopping regions separated by a dynamically evolving free boundary. To address this, Seydel formulates the problem as a Linear Complementarity Problem (LCP) to incorporate the necessary inequality constraints. The study introduces penalty methods and obstacle formulations to approximate the free boundary without explicitly calculating its location. Seydel places significant emphasis on the stability and accuracy of numerical methods, highlighting the suitability of implicit and Crank-Nicolson schemes for these problems. His detailed exploration of boundary conditions, algorithmic implementation, and numerical stability positions this work as a foundational resource in computational finance.

Brandimarte' (2018) paper addressed the practical challenges of implementing the FDM in computational environments, with a focus on numerical grid design, efficiency optimization, and error minimization. His insights provided guidance on balancing computational cost with accuracy, making FDM applications more feasible for complex financial derivatives. Hull (2021) provided a detailed explanation of how the Black-Scholes partial differential equation (PDE) underpins option pricing models, including adaptations to account for early exercise features inherent in American options. His work laid the foundation for defining boundary conditions and early exercise constraints, essential for accurate numerical implementations such as the Finite Difference Method (FDM).

The study by Cen and Chen (2019) introduces a novel HODIE (high order via differential identity expansion) finite difference scheme for pricing American options. This approach uniquely addresses the linear complementarity problem (LCP) associated with American-style options, which is known for its complexity due to the early exercise feature. The HODIE method combines a piecewise uniform spatial mesh with the implicit Euler method for temporal discretization, resulting in a matrix that is both an M-matrix and maximum-norm stable. This stability is critical, as it ensures the numerical solution remains bounded and physically meaningful under a wide range of parameter settings. The authors demonstrate that their method achieves first-order convergence in the time direction and second-order convergence in the spatial direction, as validated by both theoretical analysis and numerical experiments.

Stochastic Aplha, Beta, Rho (SABR) Model

The *SABR* (Stochastic Alpha, Beta, Rho) model was first introduced by Hagan, Kumar, Lesniewski, and Woodward (2002) in their seminal paper, *Managing Smile Risk*. They proposed a framework in which both the underlying forward price F_t and its instantaneous volatility α_t are driven by correlated stochastic processes. This innovation addressed the persistent challenge of capturing the implied volatility smile/skew in a relatively tractable manner, especially in interest-rate markets. The authors also derived an approximate closed-form formula for the implied volatility—often referred to as the "Hagan formula"—that remains the industry standard for calibrating the SABR parameters $\{\alpha, \rho, \nu, \beta\}$.

Early applications focused heavily on *interest-rate derivatives* (e.g., swaptions), but subsequent research adapted SABR to *equity and FX* contexts by fixing $\beta = 1$ for lognormal dynamics. For instance, West (2005) elaborated on practical calibration steps for SABR, describing how to incorporate weighting schemes for near-the-money points and to handle numerical stability. Over time, practitioners realized that SABR's flexible correlation parameter ρ and vol-of-vol parameter ν provided an intuitive handle on the shape of the skew, prompting further empirical studies—such as in Antonov, Konikov, and Spector (2019)-that compared SABR's smile fit to local-vol and stochastic-vol alternatives.

Subsequent refinements included adjusting the beta parameter (β) to values less than 1 for capturing "normal" or "sub-lognormal" dynamics, particularly in negative-rate environments. In the equity sphere, authors like De Col and Pintaud (2015) explored calibrating SABR with discrete dividends, ensuring the lumpsum-adjusted forward was used in the Hagan formula. More recent literature (e.g., Lund 2023) investigates how the SABR model can be specified and calibrated to capture the full implied volatility smile in interest rate derivatives under varying market conditions, with a practitioner oriented focus on risk management. Nevertheless, the original approximation from Hagan et al. (2002) remains widely employed in both academia and industry, owing to its blend of *analytic tractability* and *empirical success* in matching implied volatility surfaces.

Gaps Identified that pertain to the Study

This study explicitly fills the following gaps in the American option pricing literature:

1. **Application of a market consistent non–constant volatility** Prior work in Monte Carlo and FDM all assumed a constant volatility model. An assumption that presents a significant gap, as it may not accurately capture the effects of volatility skews and term structures observed in real markets. This paper calibrates and calculates a non–constant volatility and applies it in the valuation models of both LSMC and FDMM.

2. **Unified SABR Based Benchmark for LSMC vs. FDM.** Prior work applied Monte Carlo such as Longstaff Schwartz (2001) or PDE methods Hull White (1990) and Elliott McKinnon (1995) in isolation. We calibrate a SABR volatility surface Hagan (2002) and use the identical σ_{SABR} in both LSMC and FDM implementations, enabling a true apples–to–apples comparison.

3. **Realistic Multi Date Discrete Dividends Treatment.** Brennan & Schwartz single lump approach (Brennan & Schwartz (1977)) and Fama & French known–dividend analysis do not extend to arbitrary dividend schedules. We implement and mirror a lumps–shock scheme in both FDM and LSMC for multi–date dividends.

4. **Systematic Convergence and Runtime Comparison.** While Clement, Lamberton & Protter (Clement, Lamberton & Protter (2002)) and Stentoft (Stentoft (2004)) analyze LSMC convergence, and Elliott & McKinnon (Elliott & McKinnon (1995)) assess FDM

stability, no study benchmarks both methods convergence rates, pricing error, and computational cost on the same contract. We fill this gap with detailed error vs runtime analyses.

5. **Variance â€ Reduction and Basis Selection under a SABR Smile.** Foundational LSMC enhancements such as Carriere (1996),Broadie & Glasserman (1997) and recent basis studies Huo (2023) do not consider the impact of a market–calibrated smile. We test polynomial vs. Laguerre bases alongside antithetic and control–variate techniques under our SABR fit, identifying optimal combinations for short–dated calls.

6. **Head to Head Assessment of Multiple FDM Schemes.** Hull & White (1990) and Seydel (2015) describe explicit, implicit, and Crank Nicolson schemes, but do not compare all three under non – constant volatility. We implement and benchmark each, subject to a calculated non - constant market consistent volatility, revealing their respective stability and speed trade offs.

Chapter Three

Research Methodology

3.1 Introduction and Preliminary Set-up

Defining the American Option Valuation Problem

An American option allows early exercise, which makes its valuation a problem of solving a free-boundary partial differential equation (PDE). The optimal stopping problem is commonly formulated for American options. In this section, we define the mathematical model for valuing an American option under both Least Squared Monte Carlo as well as Finite Differences Method, however, in place of the constant volatility, we strictly incorporate an implied volatility framework under the Stochastic Aplha, Beta, Rho (SABR) model, (σ_{SABR}). In both models, we first calculate the σ_{SABR} which then becomes the respective models σ used.

Mathematical Model

Consider an asset price S_t following a Geometric Brownian Motion. The stochastic differential equation (SDE) for S_t is given by:

$$dS_t = \mu S_t \, dt + \sigma S_t \, dW_t, \tag{3.1}$$

where μ is the drift, σ is the volatility, and W_t is a Wiener process.

For an American option, the pay-off function $\phi(S_T)$ at maturity T can be defined as:

$$\phi(S_T) = \max(K - S_T, 0), \tag{3.2}$$

where K is the strike price of the option. This expression represents the payoff of an American put option, which is one example of an American option. The problem is then to find the value $V(t, S_t)$ that maximizes the expected payoff under the possibility of early exercise.

3.2 Least-Squares Monte Carlo (LSMC) for American Options

3.2.1 Overview of American Option Pricing Challenge

An American option permits exercise at any time up to expiry, making it more complex than a European option. Standard closed-form solutions (e.g., Black–Scholes) do not apply directly to American calls on dividend-paying stocks. Instead, numerical methods like binomial trees or finite differences can handle early exercise, but can be computationally expensive for high-dimensional or path-dependent problems.

The *Least-Squares Monte Carlo* (LSMC) method, introduced by Longstaff and Schwartz (2001), is a flexible approach that uses regression techniques on Monte Carlo paths to approximate the continuation value at each step, thereby deciding whether to exercise early on each path.

3.3 Incorporation of SABR Calibrated Implied Volatility (σ_{SABR})

We consider an American option on an asset whose price S_t follows a Geometric Brownian Motion. The discrete time grid is defined as $t_0, t_1, \ldots, t_N$ over the lifespan $[0, T]$ of the option, where T is the expiration time. We simulate N sample paths for the asset price $S_t^{(i)}$, $i = 1, 2, \ldots, N$, using Monte Carlo sampling of the underlying SDE.

3.3.1 Background on the SABR Framework

The SABR model (Stochastic Alpha, Beta, Rho), introduced by Hagan et al. (2002), is a popular approach for capturing the implied volatility smile or skew observed in equity, FX, and interest rate markets. In its general form, the SABR model posits that the forward price F_t of the underlying and its volatility α_t evolve under the risk-neutral measure as follows:

$$\begin{cases} dF_t \;=\; \alpha_t \, F_t^{\beta} \, dW_t^{(1)}, \\ d\alpha_t \;=\; \nu \, \alpha_t \, dW_t^{(2)}, \end{cases} \tag{3.3}$$

with

$$\mathrm{Corr}\big(dW_t^{(1)}, dW_t^{(2)}\big) \;=\; \rho. \tag{3.4}$$

Here,

- F_t would be a forward price of the underlying,
- α_t is the instantaneous volatility of F_t,
- β is a parameter controlling the dependence of the volatility on the level of F_t,
- ν is the volatility-of-volatility parameter,
- ρ is the correlation between the Brownian increments driving F_t and α_t.

Lognormal Case ($\beta = 1$) and Hagan's Approximation for Implied Volatility

In many equity or FX markets, a common choice is $\beta = 1$. This implies that

$$dF_t = \alpha_t\, F_t\, dW_t^{(1)}, \quad d\alpha_t = \nu\, \alpha_t\, dW_t^{(2)}, \quad \mathrm{Corr}(W_t^{(1)}, W_t^{(2)}) = \rho. \tag{3.5}$$

The forward price F_t thus evolves roughly like a *lognormal* process (when α_t is slowly varying). Hagan et al. derived a closed-form approximation for the implied volatility of a European option in this setting, which has become standard in practice.

For a given strike K and maturity T, Hagan et al. (2002) showed that the SABR-implied volatility, denoted $\sigma_{\mathrm{SABR}}(K)$, can be approximated by

$$\sigma_{\mathrm{SABR}}(K) \;=\; \alpha\, \frac{z}{\chi(z)}, \tag{3.6}$$

where

$$z \;=\; \frac{\nu}{\alpha}\, \ln\!\Big(\frac{F}{K}\Big), \quad \chi(z) \;=\; \ln\!\Big(\frac{\sqrt{1 - 2\rho z + z^2} + z - \rho}{1-\rho}\Big), \tag{3.7}$$

and:

- α is the *base volatility* (sometimes referred to as the initial level of α_t),
- ν is the *volatility-of-volatility*,
- ρ is the correlation between the Brownian motions driving the forward and the volatility,
- F is the lumpsum-forward of the underlying,
- K is the option strike.

This formula is often called the *Hagan SABR formula* for $\beta = 1$.

Parameter Constraints and Domains

In order to ensure numerical stability and meaningful parameters:

$$\alpha \; > \; 0, \quad \nu \; > \; 0, \quad |\rho| \; < \; 1.$$

These conditions guarantee that the approximate solution remains well-defined and that the forward price dynamics avoid pathological behaviors.

Weighted Least-Squares Calibration

To calibrate $\{\alpha, \rho, \nu\}$ to market data, we use a weighted least-squares objective function:

$$\text{SSE} \; = \; \sum_{i=1}^{N} \, w_i \, \Big(\sigma_{\text{SABR}}(K_i) - \sigma_{\text{market},i} \Big)^2, \tag{3.8}$$

where:

$$w_i \; = \; \frac{1}{1 + (K_i - F_i)^2}, \tag{3.9}$$

and:

$$\sigma_{\text{SABR}}(K_i)$$

is computed via the Hagan formula for each strike K_i. The lumpsum-forward F_i is obtained for each maturity T_i by subtracting discrete dividends (as described in the lumpsum-forward section).

ATM Emphasis via Weighting. The factor $w_i = 1 / \big(1 + (K_i - F_i)^2\big)$ places heavier emphasis on near-the-money options $\big(K_i \approx F_i\big)$. This is often desired because ATM implied volatilities typically carry the greatest market liquidity and are crucial for many trading and risk-management strategies.

Minimization Technique

- **Algorithm:** We employ the L-BFGS-B method, which handles bound constraints on α, ρ, ν. This strictly ensures: $\alpha, \nu > 10^{-4}, \quad |\rho| < 0.999$, preventing degenerate parameter values.

  Resulting SABR Surface. Once $\{\alpha, \rho, \nu\}$ are obtained, one can generate a full *implied volatility surface* over a range of strikes and maturities by evaluating the Hagan formula at each (K, T). This surface is then used in subsequent American option pricing (via a constant-vol approximation at a chosen strike/maturity) or for other risk-management tasks.

  This is done to make the LSMC method market-consistent.

3.3.2 LSMC in American Option Pricing

1. **Monte Carlo Path Simulation:**
   - We simulate N independent paths $\{S_t^{(i)}\}_{i=1}^N$ of the underlying under the risk-neutral measure.
   - Each path has M time steps, so the time increment is $\Delta t = \frac{T}{M}$, where T is the total time to expiry.
   - If the underlying follows a geometric Brownian motion with constant volatility σ, the update rule is:

$$S_{(t+\Delta t)}^{(i)} = S_t^{(i)} \exp\Big(\big(r - \tfrac{1}{2}\sigma^2\big)\,\Delta t + \sigma\sqrt{\Delta t}\; Z_{t,i}\Big), \tag{3.10}$$

     where $Z_{t,i} \sim N(0,1)$ are independent standard normals.

2. **Backward Induction for Early Exercise:**
   - At final time $t = M\Delta t = T$, the payoff of an American call is $\max(S_T^{(i)} - K, 0)$ for each path i.

– For earlier times $t = (M-1)\Delta t, \ldots, \Delta t$, we proceed backward. At each time step:

1. Identify *in-the-money* paths, i.e. those where $S_t^{(i)} > K$.
2. For those paths, we have two potential payoffs:

$$\text{Immediate exercise} \;=\; \max\bigl(S_t^{(i)} - K, 0\bigr),$$

and

$$\text{Continuation value} \;\approx\; \widehat{C}\bigl(S_t^{(i)}\bigr),$$

which must be estimated via regression.

3. If $\max(S_t^{(i)} - K, 0) \;>\; \widehat{C}\bigl(S_t^{(i)}\bigr)$, then early exercise is optimal on that path at time t.
4. Otherwise, the path continues, and the payoff remains the discounted payoff from future steps.

– By iterating this procedure backward through all time steps, we determine whether each path exercises early at some time or continues to the next step.

3. **Regression & Estimation of the Continuation Value**

– The crux of LSMC is to *estimate* the continuation value $\widehat{C}(S_t)$ using a regression on simulated paths. We estimate the continuation value by regression on a set of basis functions. The continuation value represents the expected pay-off if the option is held beyond the current time step.

– Suppose we have in-the-money paths at time t with states $\bigl\{S_t^{(i)}\bigr\}i \in I_t$. We already know their discounted future payoffs $Y_i = \text{Payoff}t + \Delta t^{(i)}\, e^{-r\,\Delta t}$ from step $t + \Delta t$ onward (after deciding exercise at $t + \Delta t$ or later).

– We fit a function $C_\theta(x)$ (often a linear combination of basis functions) to the pairs (x_i, Y_i) where $x_i = S_t^{(i)}$:

$$\min_\theta \sum_{i \in I_t} \bigl(Y_i - C_\theta(x_i)\bigr)^2. \tag{3.11}$$

– The fitted function $\widehat{C}(x) = C_{\widehat{\theta}}(x)$ is the *continuation value* estimate at state x.

– At each time step t_n and for each path $S_t^{(i)}$, we compute the discounted cash flows and approximate the continuation value $C_n(S_t^{(i)})$ by fitting a regression model. We minimize the squared difference between the actual cash flows and the estimated continuation value. The approximation takes the form:

$$C_n(S_t^{(i)}) \approx \sum_{j=0}^{P} \beta_j \phi_j(S_t^{(i)}), \tag{3.12}$$

where β_j are the regression coefficients that minimize the following error:

$$\text{Error} = \sum_{i=1}^{M} \left(C_n(S_t^{(i)}) - \sum_{j=0}^{P} \beta_j \phi_j(S_t^{(i)}) \right)^2 . \tag{3.13}$$

4. **Exercise & Exercise Condition vs. Continue Decision**

For each time step t_n and each path i, we calculate the option pay-off if exercised, $\phi(S_t^{(i)})$, and compare it with the estimated continuation value $C_n(S_t^{(i)})$ to determine the optimal decision.

Define the pay-off of the option at time t_n as $\phi(S_t^{(i)})$, which for a put option is:

$$\phi(S_t^{(i)}) = \max(K - S_t^{(i)}, 0). \tag{3.14}$$

We then set the value $V(t_n, S_t^{(i)})$ at each point in time according to:

$$V(t_n, S_t^{(i)}) = \max\left(\phi(S_t^{(i)}), C_n(S_t^{(i)})\right). \tag{3.15}$$

If $\phi(S_t^{(i)}) \geq C_n(S_t^{(i)})$, the optimal decision is to exercise the option at t_n; otherwise, the option is held, and the continuation value is applied.

3.3.3 Choice of Basis Functions

LSMC requires a finite set of basis functions $\{\phi_1(x), \phi_2(x), \ldots, \phi_d(x)\}$ to approximate the continuation value $C_\theta(x)$:

$$C_\theta(x) \;=\; \theta_1\,\phi_1(x) \;+\; \theta_2\,\phi_2(x) \;+\; \cdots + \; \theta_d\,\phi_d(x). \tag{3.16}$$

Two common choices:

1. **Polynomial Basis**:

$$\phi_j(x) = x^j, \quad j = 0, \ldots, p. \tag{3.17}$$

For instance, $\{1, x, x^2, \ldots, x^p\}$.

2. **Laguerre Basis**:

$$L_0(x) = 1, \quad L_1(x) = 1 - x, \quad L_2(x) = 1 - 2x + \tfrac{x^2}{2}, \quad \dots \tag{3.18}$$

These are truncated Laguerre polynomials often used in LSMC.

In practice, p is kept small (e.g. 2–4) to avoid overfitting. We solve for θ via ordinary least squares:

$$\widehat{\theta} \;=\; \arg\min_{\theta} \sum_{i \in I_t} \big(Y_i - C_\theta(x_i)\big)^2. \tag{3.19}$$

3.3.4 Algorithmic Steps of LSMC

1. **Simulate Paths:** Generate N sample paths $\{S_t^{(i)}\}$, $t = 0, \Delta t, \dots, T$, under the risk-neutral measure.
2. **Compute Terminal Payoffs:** For each path i, $\text{Payoff}_T^{(i)} = \max\big(S_T^{(i)} - K,\, 0\big)$.
3. **Backward Induction:** For $t = T - \Delta t, \dots, \Delta t$:
   (a) Identify in-the-money paths $I_t = \{i \mid S_t^{(i)} > K\}$.
   (b) Let Y_i be the discounted payoff from time $t + \Delta t$ onward (already computed at the next step).
   (c) Fit a regression $Y_i \approx C_\theta(S_t^{(i)})$ using the chosen basis.
   (d) For each $i \in I_t$, compare $\text{exercise}_t^{(i)} = \max(S_t^{(i)} - K, 0)$ to $\widehat{C}\big(S_t^{(i)}\big)$.
   (e) If $\text{exercise}_t^{(i)} > \widehat{C}\big(S_t^{(i)}\big)$, set $\text{Payoff}_t^{(i)} = \text{exercise}_t^{(i)}$, and treat the path as exercised (no future payoff).
   (f) Else, continue the path with payoff Y_i.
4. **Final Price:** The time-zero LSMC price is the average of $\text{Payoff}_0^{(i)}$ (discounted if needed) across all paths:

$$\text{Price}_{\text{LSMC}} \;=\; e^{-r\Delta t}\, \frac{1}{N} \sum_{i=1}^{N} \text{Payoff}_0^{(i)}. \tag{3.20}$$

3.4 Finite Difference Method (FDM) Methodology

3.4.1 Overview of the FDM Method

The Finite Difference Method (FDM) is a widely used numerical technique to approximate solutions to partial differential equations (PDEs), making it suitable for the valuation of American options. By discretizing time and space, FDM allows us to approximate the option value function, which satisfies the Black-Scholes PDE with an early exercise feature. As provided for by Morton Mayers (2005) and further refined in Brandimarte's 2018 Hull's 2021 papers in implementation in computational settings, this section will describe a take on FDM in valuing American Options with non constant volatility as incorporated in Gatheral's 2006 paper.

3.4.2 SABR Calibration (Lognormal, $\beta = 1$)

In the lognormal (i.e. $\beta = 1$) SABR model, we have

$$\begin{aligned} dS_t &= \alpha_t\, S_t\, dW_{1t}, \\ d\alpha_t &= \nu\, \alpha_t\, dW_{2t}, \quad \text{corr}(dW_1, dW_2) \;=\; \rho. \end{aligned} \tag{3.21}$$

For a European option, Hagan et al. derived an approximate formula for the implied volatility $\sigma_{\text{SABR}}(F, K)$. In the special case $\beta = 1$, that formula often takes the form

$$\sigma_{\text{SABR}} \;\approx\; \alpha\, \frac{z}{\chi(z)}, \quad \text{where } z \;=\; \frac{\nu}{\alpha}\, \ln\!\Big(\frac{F}{K}\Big), \quad \chi(z) \;=\; \ln\!\Big(\frac{\sqrt{1 - 2\,\rho\, z + z^2} + z - \rho}{1 - \rho}\Big). \tag{3.22}$$

Here,

- α is the initial volatility parameter,
- ν is the vol-of-vol,
- ρ is the correlation between Brownian increments,
- F is the forward price for maturity T,
- K is the option strike.

3.4.3 Discrete Dividends $\rightarrow$ Lumpsum Forward and Weighted Objective for SABR

Because the underlying pays *discrete* dividends, the forward F for time T is not simply $S_0\, e^{rT}$. Instead, we do a lumpsum approach:

$$F \;=\; S_0\, e^{rT} \;-\; \Big(\text{sum of lumps}\Big). \tag{3.23}$$

In simpler code, if there are n lumps each of size $\mathrm{DivAmount}$, then

$$F \;=\; S_0\, e^{rT} \;-\; n\; \times \mathrm{DivAmount}.$$

(This is an approximation, ignoring the precise discounting of each lumpsum.)

To calibrate (α, ρ, ν), we define an objective function over all market quotes:

$$\mathrm{Objective}(\alpha, \rho, \nu) \;=\; \sum_{\text{rows}} \; w(K, F)\, \Big[\sigma_{\mathrm{SABR}}(F, K) - \sigma_{\mathrm{market}}\Big]^2, \tag{3.24}$$

where σ_{market} is the market implied vol, and σ_{SABR} is given by Hagan's formula above. The weight $w(K, F) = 1/[\,1 + (K - F)^2\,]$ can emphasize near-ATM quotes. We run a numerical minimizer (e.g. `scipy.optimize.minimize`) to find α, ρ, ν that minimize this objective.

3.4.4 Extracting a Single σ

After calibrating, for a chosen row $(K_{\text{sel}}, T_{\text{sel}}, r_{\text{sel}})$, we compute lumpsum forward F and plug into the SABR formula:

$$\sigma_{\mathrm{SABR}} \;=\; \sigma_{\mathrm{SABR}}(F, K_{\text{sel}};\; \alpha, \rho, \nu), \tag{3.25}$$

which yields a single "constant" volatility used below in the PDE. This ensures the PDE's σ is consistent with real data.

3.5 PDE for an American Call with Lumps

3.5.1 PDE Setup

We consider the PDE for an American call (with no continuous dividend):

$$\frac{\partial V}{\partial t} + \tfrac{1}{2}\,\sigma^2\, S^2\, \frac{\partial^2 V}{\partial S^2} + r\, S\, \frac{\partial V}{\partial S} - r\, V = 0, \tag{3.26}$$

under the terminal (final) condition

$$V(T, S) \;=\; \max(S - K, 0), \tag{3.27}$$

and the American constraint

$$V(t, S) \;\ge\; (S - K)^{+}. \tag{3.28}$$

3.5.2 Discrete Dividends (Lumps "Shock" Approach)

If the underlying pays lumps at times $t_{\text{div1}}, t_{\text{div2}}, \ldots$, the PDE solution is standard between those times. But at each ex-div instant, the underlying is reduced by DivAmount. The code handles this by:

- *Re-interpolating* the PDE solution in S, shifting $S_{\text{old}} \mapsto S_{\text{old}} + \text{DivAmount}$,
- Then storing the new solution array on the original S-grid.

Hence, each lumps reduces S in the PDE by a fixed amount, capturing discrete dividends.

3.5.3 Spatial Grid

We define a domain $S \in [0, S_{\max}]$, where $S_{\max}$ is e.g. $2.5\, S_0, 3.0\, S_0$, or $4.0\, S_0$, depending on T. Let M be the number of steps, so

$$\Delta S \;=\; \frac{S_{\max}}{M}. \tag{3.29}$$

We store the solution in an array $V[i],\ i = 0, \ldots, M$.

3.5.4 Time Stepping

We define N steps in time from $t = T$ down to $t = 0$, so

$$\Delta t \;=\; \frac{T}{N}. \tag{3.30}$$

We initialize $V_{\text{final}}[i] = \max(S_{\text{grid}}[i] - K, 0)$. Then at each step $n = N, \ldots, 1$, we compute V_{n-1} from V_n, possibly applying a lumps shock if crossing an ex-div date.

3.5.5 Finite Difference Schemes

Explicit Scheme. We write

$$V_{n-1}[i] \;=\; V_n[i] \;-\; \Delta t \left[\tfrac{1}{2}\, \sigma^2\, S_i^2\, V_n''(i) \;+\; r\, S_i\, V_n'(i) \;-\; r\, V_n[i]\right], \tag{3.31}$$

where

$$V_n''(i) \;\approx\; \frac{V_n[i+1] - 2\, V_n[i] + V_n[i-1]}{\Delta S^2}, \quad V_n'(i) \;\approx\; \frac{V_n[i+1] - V_n[i-1]}{2\, \Delta S}. \tag{3.32}$$

The explicit scheme is conditionally stable and requires:

$$\Delta t \leq \frac{1}{2} \left(\frac{\Delta S^2}{\sigma(K,T)^2 S^2} \right). \tag{3.33}$$

Implicit Scheme. The implicit scheme solves for V_i^{n+1} by setting up a tridiagonal system of equations:

$$-a_i V_{i-1}^{n+1} + (1 + b_i)V_i^{n+1} - c_i V_{i+1}^{n+1} = V_i^n. \tag{3.34}$$

This system is solved efficiently using the Thomas algorithm, which operates in $O(M)$ time.

Crank-Nicolson Scheme. The Crank-Nicolson scheme averages the explicit and implicit schemes:

$$-\frac{a_i}{2}V_{i-1}^{n+1} + \left(1 + \frac{b_i}{2}\right)V_i^{n+1} - \frac{c_i}{2}V_{i+1}^{n+1} = \tag{3.35}$$

$$\frac{a_i}{2}V_{i-1}^{n} + \left(1 - \frac{b_i}{2}\right)V_i^{n} + \frac{c_i}{2}V_{i+1}^{n}. \tag{3.36}$$

This scheme achieves second-order accuracy in both time and space.

3.5.6 Boundary and Early Exercise Conditions

Early Exercise: Enforce the early exercise condition at each time step:

$$V_i^n = \max(V_i^n, \phi(S_i)). \tag{3.37}$$

Boundary Conditions:

- At $S = 0$: $V(t, 0) = Ke^{-r(T-t)}$, reflecting intrinsic value for a put option.
- At $S = S_{\text{max}}$: $V(t, S_{\text{max}}) \to 0$, as the option value approaches zero for very high asset prices.

3.5.7 American Constraint

After updating V_{n-1} from V_n, we enforce

$$V_{n-1}[i] \;\leftarrow\; \max\bigl(V_{n-1}[i],\; S_{\text{grid}}[i] - K\bigr). \tag{3.38}$$

This ensures the PDE solution respects the possibility of early exercise.

3.5.8 Lumps "Shock" Implementation

If a lumps δ is paid at t_{div}, we shift $S \mapsto S - \delta$. In practice, we re-interpolate:

$$\text{old coords} = (s_{\text{grid}} + \delta), \quad \text{new coords} = s_{\text{grid}}. \tag{3.39}$$

We use a spline or linear interpolation to map the old solution to the new coords. At $t = 0$, we have $V_0[i]$. We then interpolate in S to get

$$\text{Price} \;=\; \operatorname{interp}\bigl(s_{\text{grid}},\, V_0,\, S_0\bigr).$$

3.5.9 Comparison to Binomial & BFS

The code compares:

1. *Binomial lumps:* A simpler lumps approach in a binomial tree for American calls,
2. *Bjerksund–Stensland (BFS):* An analytic approximation for an American call with continuous dividend yield $q \approx \text{annual\_div}/S_0$.

We see how close the PDE lumps price is to those references, for a sanity check.

3.6 Convergence Analysis of LSMC and FDM Methods

To verify the accuracy and stability of both the Least Squares Monte Carlo (LSMC) and Finite Difference (FDM) methods, we analyze the convergence behavior of each method by gradually refining their key parameters and observing the stability of the option price estimates.

3.6.1 Convergence of the LSMC Method

The convergence of the LSMC method is assessed by increasing the number of simulated paths, testing two basis functions, introducing variance reduction techniques and evaluating the effect on the estimated option price. The paper will focus on these three aspects.

1. **Path Simulation Refinement:** Adjusting the number of time steps M and the number of simulated paths N (and possibly using advanced variance reduction methods).
2. **Basis Refinement:** Improving the approximation of the continuation value by increasing polynomial order or switching to richer basis functions (e.g. Laguerre polynomials).
3. **Variance Analysis:** Repeatedly running the Monte Carlo (with different seeds) to observe mean and standard deviation of the estimated American option price, thereby quantifying numerical stability.

In this section, we expand on these aspects, showing the relevant LaTeX-friendly equations and methodology.

3.6.2 Path Simulation Refinement

Time Discretization and Number of Paths

We consider an American call on a stock with constant risk-free rate r and time to maturity T. The LSMC method simulates N paths $\{S_t^{(i)}\}_{i=1}^N$ in M discrete time steps $\Delta t = T/M$. For each path i and step $t_m = m\Delta t$, we update

$$S_{t_{m+1}}^{(i)} = S_{t_m}^{(i)} \exp\Big(\big(r - \tfrac{1}{2}\sigma^2\big)\Delta t + \sigma\sqrt{\Delta t}\; z_{m,i}\Big), \tag{3.40}$$

where $z_{m,i} \sim N(0,1)$ are independent draws (unless variance reduction modifies them). The *refinement* typically involves:

– **Increasing M:** More time steps yield a finer temporal resolution, capturing early-exercise decisions more accurately.

– **Increasing N:** More simulated paths reduce the Monte Carlo standard error, stabilizing the LSMC price.

Eulerâ€“Maruyama Scheme: The above exponential update is a discrete form of a geometric Brownian motion (or other processes) under the risk-neutral measure. As $M \to \infty$, this converges to the continuous model, but in practice, moderate M would suffice if early-exercise points are well approximated.

3.6.3 Basis Refinement for Continuation Value Regression

A central step in LSMC is the *backward induction* regression for the *continuation value* at each time step. Let t_m be a backward time step, and suppose we have *discounted* payoffs from future times, Y_i, for all in-the-money paths i. We regress Y_i against a set of basis functions of $S_{t_m}^{(i)}$:

$$Y_i \;\approx\; \sum_{j=0}^{p} \beta_j \, \phi_j\big(S_{t_m}^{(i)}\big). \tag{3.41}$$

Solving for β_j via least squares yields an estimate of the *continuation value* at state $S_{t_m}^{(i)}$.

Polynomial Basis. A polynomial basis of order p might be:

$$\phi_j(x) = x^j, \quad j = 0, 1, \ldots, p. \tag{3.42}$$

Hence,

$$C_\theta(x) \;=\; \theta_0 + \theta_1 \, x + \theta_2 \, x^2 + \cdots + \theta_p \, x^p. \tag{3.43}$$

Laguerre Basis. Alternatively, truncated *Laguerre polynomials* can be used, e.g.:

$$L_0(x) = 1, \quad L_1(x) = 1 - x, \quad L_2(x) = 1 - 2x + \tfrac{x^2}{2}, \ \ldots$$

One typically regresses Y_i onto $\{L_0(X_i),\, L_1(X_i),\, \ldots\}$. For $p = 2$, say, that includes L_0, L_1, L_2.

3.6.4 Variance Analysis

Even with advanced path simulation and a refined basis, LSMC remains a Monte Carlo method subject to sampling noise. A robust practice is to:

1. **Fix a Scenario:** e.g. choose $(p = 2,\ N = 2000,\ \text{VR} = \text{antithetic},\ \text{CV} = \text{False})$.
2. **Re-run LSMC Multiple Times:** use different seeds $\text{seed}k$ for $k = 1, \ldots, \text{num\_runs}$ to generate different random draws zt, i.
3. **Collect LSMC Prices:** let $\{X_k\}_{k=1}^{\text{num\_runs}}$ be the resulting American call prices from each run.
4. **Compute Mean and Std. Dev.:**

$$\overline{X} = \frac{1}{\text{num\_runs}} \sum_{k=1}^{\text{num\_runs}} X_k, \quad \text{StdDev}(X) = \sqrt{\frac{1}{\text{num\_runs} - 1} \sum_{k=1}^{\text{num\_runs}} \left(X_k - \overline{X}\right)^2}. \tag{3.44}$$

5. **Interpretation:**
   - If the standard deviation is large, it indicates the methodâ€™s variance is high for that scenario (more paths or stronger variance reduction would be needed).
   - If the standard deviation is small, the method is stable for that $(p,\ N)$ choice.

Combining With Benchmarks. If a benchmark (like a binomial or Bjerksund–Stensland price) is available or can be calculated, we can also measure the *bias* or *absolute error*:

$$\left|\overline{X} - X_{\text{benchmark}}\right|,$$

alongside the standard deviation. A small bias plus a small std. dev. suggests the LSMC method is both *accurate* and *low variance*.

Hence, *refining* the path simulation (increasing N and M), *refining* the basis (increasing polynomial order or switching basis), and *analyzing variance* via multiple runs all form an integrated approach to ensure the LSMC method converges to an accurate, low-variance estimate of the American option price.

Convergence is achieved when further increases in N result in negligible changes in $V_{LSM}(N)$, indicating that the LSM method has stabilized.

3.6.5 Convergence of the FDM

The convergence of the FD method is evaluated by refining the grid resolution in both time Δt and asset price ΔS.

1. **Grid Refinement**: We start with a coarse grid and refine it by decreasing ΔS and Δt. Let $V_{FD}(\Delta S, \Delta t)$ denote the option price for a given grid size. For example, we can define:

$$M \in \{50, 100, 200, 400, 800\}, \quad N \in \{50, 100, 200, 400, 800\}.$$

2. **Error Analysis**: We run the PDE for each pair (M, N), obtaining a price $P_{\text{PDE}}(M, N)$. We compare $P_{\text{PDE}}(M, N)$ to a reference (e.g. a binomial lumps price or an analytic approximation) and define an *error*:

$$\text{Error}(M, N) = \left| P_{\text{PDE}}(M, N) - P_{\text{ref}} \right|. \tag{3.45}$$

   As we increase M and N, we expect $\text{Error}(M, N) \to 0$, indicating convergence. This refinement step is crucial to ensure the PDE solution is not heavily dependent on a coarse grid. By plotting or tabulating $\text{Error}(M, N)$ against M or N, we can see if the solution is stable and accurate. As the grid is refined, this error should decrease if the method is converging.

   Convergence is achieved when further grid refinement results in negligible changes in $V_{FD}(\Delta S, \Delta t)$, indicating that the FD method has stabilized.

3.7 Computational Efficiency Analysis

Compare the computational efficiency of LSMC and FDM based on:

- **Computation Time:** Measure runtime for different N, P, Δt, and ΔS.
- **Accuracy vs. Runtime:** Analyze the trade-off between accuracy and computation time for both methods.

Chapter Four

Model Fitting, Sensitivity Analysis and Numerical Simulations

4.1 LSMC under SABR Volatility

4.1.1 SABR Model Calibration Analysis

The SABR model is used to capture the volatility smile/smirk observed in market data. Its dynamics are:

$$dF_t = \sigma_t F_t^{\beta} dW_t^{(1)} \quad \text{and} \quad d\sigma_t = \nu \sigma_t dW_t^{(2)}, \tag{4.1}$$

with correlation ρ between $dW_t^{(1)}$ and $dW_t^{(2)}$. In our calibration, we fix $\beta = 1$ (i.e., a lognormal model).

We calibrate the model by minimizing the weighted squared error:

$$\text{SSE} = \sum_{i=1}^{n} w_i \Big(\sigma_{\text{SABR}}(K_i, F, T_i; \alpha, \rho, \nu) - \sigma_{\text{market},i} \Big)^2,$$

where the weight w_i gives more importance to near-the-money options. Our calibration yields:

$$\alpha = 0.292724, \quad \beta = 1.0, \quad \rho = -0.200011, \quad \nu = 0.300019,$$

with an objective function value of approximately 0.00251006. This very low error indicates a very close match between the SABR-implied volatilities and market data.

Financial Interpretation A negative ρ is common in equity markets since lower strikes (often associated with put protection) exhibit higher implied volatilities. The calibrated α and ν values are within a realistic range, ensuring that our volatility surface is consistent with observed market behavior. The precise calibration is essential for subsequent pricing steps because any misfit here could propagate errors throughout the model.

4.1.2 Chosen Option and Benchmark Pricing

For the selected option, AAPL 3/21/25 C210, the key output values are:

$$\begin{aligned} \text{Strike } (K) &= 210.0, \\ \text{Maturity } (T) &\approx 0.0192 \text{ years (about 7 days)}, \\ \text{Market Implied Volatility (IVM\_dec)} &\approx 0.341216, \\ \text{SABR Implied Volatility (SABR\_IV\_fit)} &\approx 0.293232, \\ \text{IV Error} &= -0.047984. \end{aligned}$$

The corresponding LSMC input parameters are:

$$S_0 = 213.49, \quad K = 210, \quad T = 0.0192, \quad r \approx 0.0005, \quad \sigma_{\text{SABR}} \approx 0.2932.$$

The SABR-implied volatility is computed using the calibrated parameters, and the resulting difference (IV Error) highlights calibration nuances that can occur in short-dated options. Short maturities are inherently more sensitive to parameter estimation errors. The slight discrepancy between the market-observed and SABR-implied volatilities is acceptable in this regime. Moreover, the selection of an ATM (or nearly ATM) option ensures that the pricing is relevant for highly liquid and actively traded options, which is critical for risk management and hedging.

Benchmark Pricing The model calculates benchmarks using two independent methods:

- **Binomial Tree Method:** A discrete-time model that naturally incorporates early exercise features.
- **Bjerksund–Stensland Approximation:** An analytical method designed to approximate American option prices.

The output benchmarks are:

$$\text{Binomial Price} = 5.4551, \quad \text{Bjerksund–Stensland Price} = 5.4550.$$

The near-identical values from two independent approaches validate the consistency of our pricing framework. The binomial method, with a large number of steps (5000 in our case), converges to the correct value, while the Bjerksund–Stensland method is known to perform well for American options.

4.1.3 Convergence Analysis

The convergence analysis explores how the LSMC price converges as the number of simulation paths (N) increases, across various configurations:

- Different basis orders ($p = 1, 2, 3$),
- Variance reduction settings (none vs. antithetic),
- With and without control variates.

Monte Carlo methods have a convergence rate proportional to $1/\sqrt{N}$. The table shows that as N increases, the LSMC prices (both using polynomial and Laguerre basis functions) converge towards the benchmark price of approximately 5.455. The application of antithetic variates substantially reduces the error at lower N, confirming the effectiveness of variance reduction techniques.

For practitioners, a convergent pricing model is essential to ensure accurate risk management. The convergence table demonstrates that with a sufficient number of paths and proper variance reduction, the pricing error becomes negligible. This reliability is critical when such models are used to hedge portfolios or to price exotic derivatives.

Below is the full convergence table for basis order $p = 1, 2, 3$

Table 4.1: Convergence Table for $P = 1, 2, 3$ (across various N, VR settings, and CV flags)

p	N	VR	CV	Price_Poly	Price_Lag	Error_Poly_Bin	Error_Lag_Bin	Error_Poly_BS
1	500	none	False	5.800150	5.800150	0.345041	0.345041	0.345161
1	500	none	True	5.800150	5.800150	0.345041	0.345041	0.345161
1	500	antithetic	False	5.458010	5.458010	0.002901	0.002901	0.003021
1	500	antithetic	True	5.458010	5.458010	0.002901	0.002901	0.003021
1	1000	none	False	5.667089	5.667089	0.211980	0.211980	0.212100
1	1000	none	True	5.667089	5.667089	0.211980	0.211980	0.212100
1	1000	antithetic	False	5.172469	5.172469	0.282640	0.282640	0.282519
1	1000	antithetic	True	5.172469	5.172469	0.282640	0.282640	0.282519
1	2000	none	False	5.535782	5.535782	0.080673	0.080673	0.080794
1	2000	none	True	5.535782	5.535782	0.080673	0.080673	0.080794
1	2000	antithetic	False	5.305983	5.305983	0.149126	0.149126	0.149005
1	2000	antithetic	True	5.305983	5.305983	0.149126	0.149126	0.149005
1	4000	none	False	5.508515	5.508515	0.053406	0.053406	0.053527
1	4000	none	True	5.508515	5.508515	0.053406	0.053406	0.053527
1	4000	antithetic	False	5.364332	5.364332	0.090777	0.090777	0.090656
1	4000	antithetic	True	5.364332	5.364332	0.090777	0.090777	0.090656

Continued on next page

Table 4.1 – Continued from previous page

p	N	VR	CV	Price_Poly	Price_Lag	Error_Poly_Bin	Error_Lag_Bin	Error_Poly_BS
1	8000	none	False	5.493262	5.493262	0.038154	0.038154	0.038274
1	8000	none	True	5.493262	5.493262	0.038154	0.038154	0.038274
1	8000	antithetic	False	5.380572	5.380572	0.074537	0.074537	0.074416
1	8000	antithetic	True	5.380572	5.380572	0.074537	0.074537	0.074416
1	10000	none	False	5.433227	5.433227	0.021882	0.021882	0.021761
1	10000	none	True	5.433227	5.433227	0.021882	0.021882	0.021761
1	10000	antithetic	False	5.368765	5.368765	0.086344	0.086344	0.086224
1	10000	antithetic	True	5.368765	5.368765	0.086344	0.086344	0.086224
1	15000	none	False	5.407934	5.407934	0.047175	0.047175	0.047054
1	15000	none	True	5.407934	5.407934	0.047175	0.047175	0.047054
1	15000	antithetic	False	5.401326	5.401326	0.053783	0.053783	0.053662
1	15000	antithetic	True	5.401326	5.401326	0.053783	0.053783	0.053662
1	20000	none	False	5.388904	5.388904	0.066205	0.066205	0.066084
1	20000	none	True	5.388904	5.388904	0.066205	0.066205	0.066084
1	20000	antithetic	False	5.428502	5.428502	0.026607	0.026607	0.026487
1	20000	antithetic	True	5.428502	5.428502	0.026607	0.026607	0.026487
2	500	none	False	5.937554	5.937554	0.482445	0.482445	0.482565

Continued on next page

Table 4.1 – Continued from previous page

p	N	VR	CV	Price_Poly	Price_Lag	Error_Poly_Bin	Error_Lag_Bin	Error_Poly_BS
2	500	none	True	5.937554	5.937554	0.482445	0.482445	0.482565
2	500	antithetic	False	5.660774	5.660774	0.205665	0.205665	0.205786
2	500	antithetic	True	5.660774	5.660774	0.205665	0.205665	0.205786
2	1000	none	False	5.778853	5.778853	0.323744	0.323744	0.323865
2	1000	none	True	5.778853	5.778853	0.323744	0.323744	0.323865
2	1000	antithetic	False	5.668847	5.668847	0.213739	0.213739	0.213859
2	1000	antithetic	True	5.668847	5.668847	0.213739	0.213739	0.213859
2	2000	none	False	5.561925	5.561925	0.106816	0.106816	0.106937
2	2000	none	True	5.561925	5.561925	0.106816	0.106816	0.106937
2	2000	antithetic	False	5.473853	5.473853	0.018744	0.018744	0.018865
2	2000	antithetic	True	5.473853	5.473853	0.018744	0.018744	0.018865
2	4000	none	False	5.517374	5.517374	0.062266	0.062266	0.062386
2	4000	none	True	5.517374	5.517374	0.062266	0.062266	0.062386
2	4000	antithetic	False	5.398687	5.398687	0.056422	0.056422	0.056302
2	4000	antithetic	True	5.398687	5.398687	0.056422	0.056422	0.056302
2	8000	none	False	5.514329	5.514329	0.059220	0.059220	0.059341
2	8000	none	True	5.514329	5.514329	0.059220	0.059220	0.059341

Continued on next page

Table 4.1 – Continued from previous page

p	N	VR	CV	Price_Poly	Price_Lag	Error_Poly_Bin	Error_Lag_Bin	Error_Poly_BS
2	8000	antithetic	False	5.406941	5.406941	0.048168	0.048168	0.048047
2	8000	antithetic	True	5.406941	5.406941	0.048168	0.048168	0.048047
2	10000	none	False	5.437386	5.437386	0.017723	0.017723	0.017603
2	10000	none	True	5.437386	5.437386	0.017723	0.017723	0.017603
2	10000	antithetic	False	5.409193	5.409193	0.045916	0.045916	0.045795
2	10000	antithetic	True	5.409193	5.409193	0.045916	0.045916	0.045795
2	15000	none	False	5.453157	5.453157	0.001952	0.001952	0.001832
2	15000	none	True	5.453157	5.453157	0.001952	0.001952	0.001832
2	15000	antithetic	False	5.414657	5.414657	0.040452	0.040452	0.040331
2	15000	antithetic	True	5.414657	5.414657	0.040452	0.040452	0.040331
2	20000	none	False	5.451439	5.451439	0.003670	0.003670	0.003549
2	20000	none	True	5.451439	5.451439	0.003670	0.003670	0.003549
2	20000	antithetic	False	5.439510	5.439510	0.015599	0.015599	0.015478
2	20000	antithetic	True	5.439510	5.439510	0.015599	0.015599	0.015478
3	500	none	False	5.977854	5.977854	0.522745	0.522745	0.522866
3	500	none	True	5.977854	5.977854	0.522745	0.522745	0.522866
3	500	antithetic	False	5.847140	5.906694	0.392032	0.451585	0.392152

Continued on next page

Table 4.1 – Continued from previous page

p	N	VR	CV	Price_Poly	Price_Lag	Error_Poly_Bin	Error_Lag_Bin	Error_Poly_BS
3	500	antithetic	True	5.847140	5.906694	0.392032	0.451585	0.392152
3	1000	none	False	5.807325	5.807325	0.352217	0.352217	0.352337
3	1000	none	True	5.807325	5.807325	0.352217	0.352217	0.352337
3	1000	antithetic	False	5.714261	5.734303	0.259152	0.279195	0.259273
3	1000	antithetic	True	5.714261	5.734303	0.259152	0.279195	0.259273
3	2000	none	False	5.603846	5.607842	0.148737	0.152733	0.148858
3	2000	none	True	5.603846	5.607842	0.148737	0.152733	0.148858
3	2000	antithetic	False	5.506520	5.518849	0.051411	0.063740	0.051531
3	2000	antithetic	True	5.506520	5.518849	0.051411	0.063740	0.051531
3	4000	none	False	5.646509	5.646509	0.191400	0.191400	0.191521
3	4000	none	True	5.646509	5.646509	0.191400	0.191400	0.191521
3	4000	antithetic	False	5.519806	5.527227	0.064698	0.064698	0.064818
3	4000	antithetic	True	5.519806	5.527227	0.064698	0.064698	0.064818
3	8000	none	False	5.557444	5.557444	0.102335	0.102335	0.102455
3	8000	none	True	5.557444	5.557444	0.102335	0.102335	0.102455
3	8000	antithetic	False	5.426837	5.426837	0.028272	0.028272	0.028151
3	8000	antithetic	True	5.426837	5.426837	0.028272	0.028272	0.028151

Continued on next page

Table 4.1 – Continued from previous page

p	N	VR	CV	Price_Poly	Price_Lag	Error_Poly_Bin	Error_Lag_Bin	Error_Poly_BS
3	10000	none	False	5.467846	5.467846	0.012737	0.012737	0.012858
3	10000	none	True	5.467846	5.467846	0.012737	0.012737	0.012858
3	10000	antithetic	False	5.434479	5.434479	0.020630	0.020630	0.020509
3	10000	antithetic	True	5.434479	5.434479	0.020630	0.020630	0.020509
3	15000	none	False	5.462796	5.462796	0.007687	0.007687	0.007807
3	15000	none	True	5.462796	5.462796	0.007687	0.007687	0.007807
3	15000	antithetic	False	5.439034	5.439034	0.016075	0.016075	0.015955
3	15000	antithetic	True	5.439034	5.439034	0.016075	0.016075	0.015955
3	20000	none	False	5.455985	5.455985	0.000876	0.000876	0.000996
3	20000	none	True	5.455985	5.455985	0.000876	0.000876	0.000996
3	20000	antithetic	False	5.432541	5.432541	0.022568	0.022568	0.022447
3	20000	antithetic	True	5.432541	5.432541	0.022568	0.022568	0.022447

The convergence table clearly shows that as the number of simulation paths N increases, the LSMC price estimates converge towards the benchmark price of approximately 5.455. Notably, the use of antithetic variates substantially reduces the error at lower N, in line with the theoretical $1/\sqrt{N}$ convergence rate of Monte Carlo methods. Additionally, the results across different basis orders ($p = 1, 2, 3$) and control variate settings exhibit minimal differences, demonstrating that the regression approximations (whether using polynomial or Laguerre bases) are robust. Overall, the table confirms both the mathematical consistency of the LSMC approach and its practical reliability for pricing American options.

4.1.4 Simulated Asset Price Paths

Before running the LSMC, a set of asset paths was simulated using the risk-neutral dynamics:

$$S_{t+1} = S_t \exp\Big((r - \tfrac{1}{2}\sigma_{\text{SABR}}^2)\Delta t + \sigma_{\text{SABR}}\sqrt{\Delta t}\, z\Big), \tag{4.2}$$

where $z \sim \mathcal{N}(0, 1)$. The graph below (Figure 4.1) displays 50 simulated paths over 50 time steps. Each path is generated using the risk-neutral dynamics. This visualization helps confirm that the random evolution of prices adheres to the expected stochastic process.

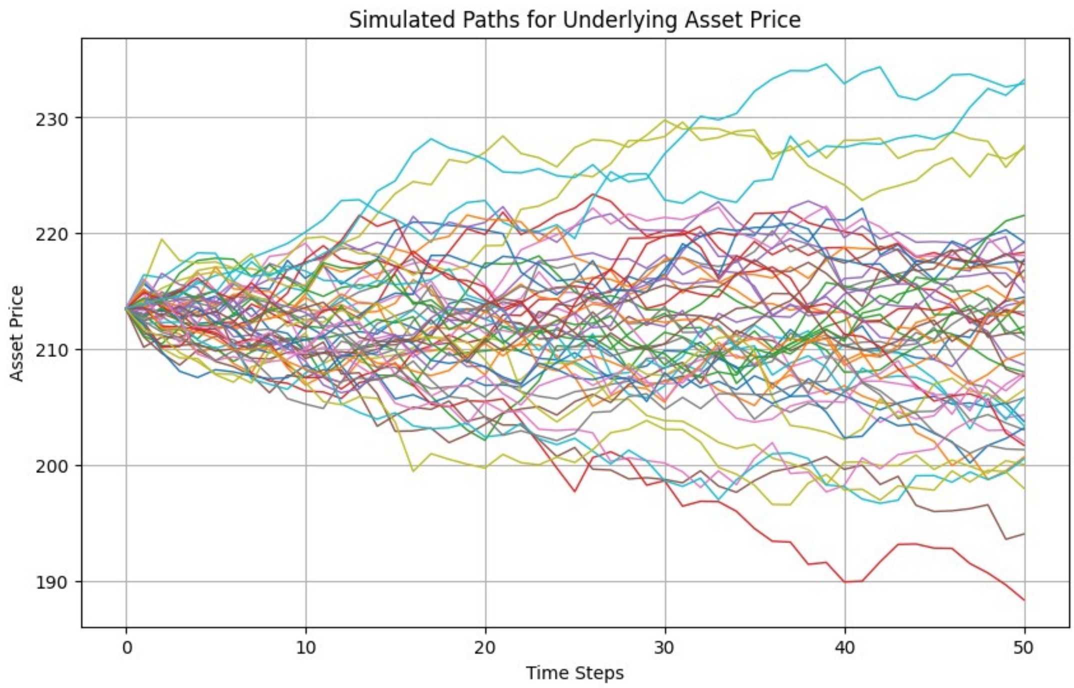

Figure 4.1: Graph of Simulated paths for the underlying asset price.

4.1.5 Convergence Graphs

The following figures present two graphs: one for the LSMC Price (Polynomial basis) versus the number of simulation paths N, and the other for the Absolute Error (difference between LSMC price and Binomial benchmark) versus N. This graph plots the LSMC price (using the polynomial basis) as a function of the number of simulation paths N. The plot shows that at lower N (e.g., 500 or 1000), the estimated price is higher (or lower) than the benchmark, but as N increases, the estimated price converges toward the benchmark value of approximately 5.455.

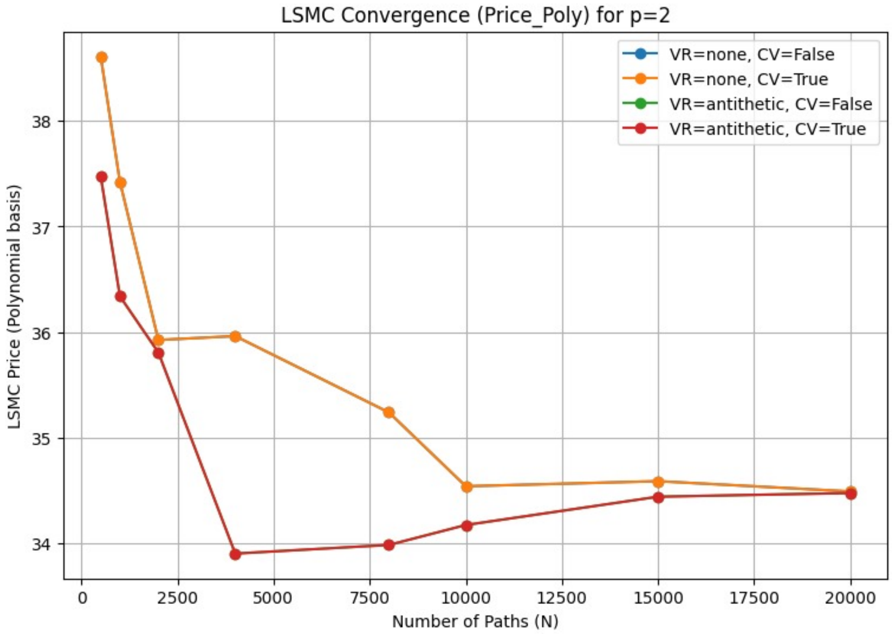

Figure 4.2: Graph depicting the convergence of the LSMC price under Polynomial basis towards the benchmark value as N increases.

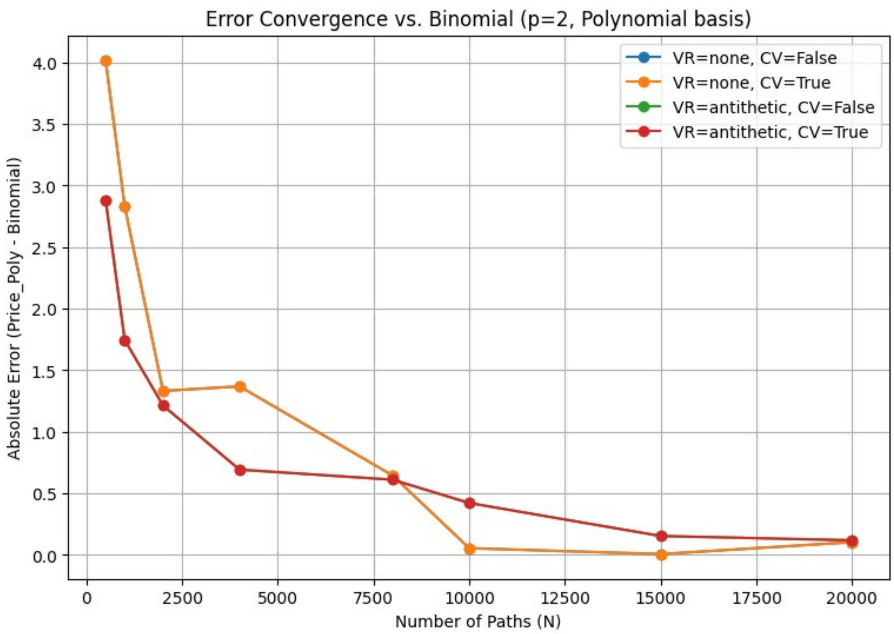

Figure 4.3: Graph of absolute error reductions as N increases.

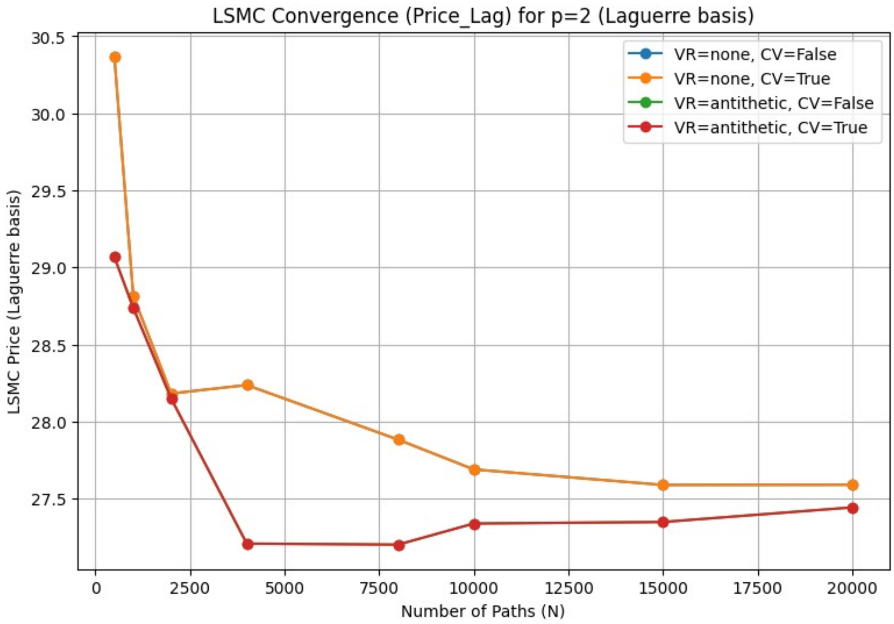

Figure 4.4: Graph showing the convergence of the LSMC price under the Laguerre basis towards the benchmark value as the number of simulation paths N increases.

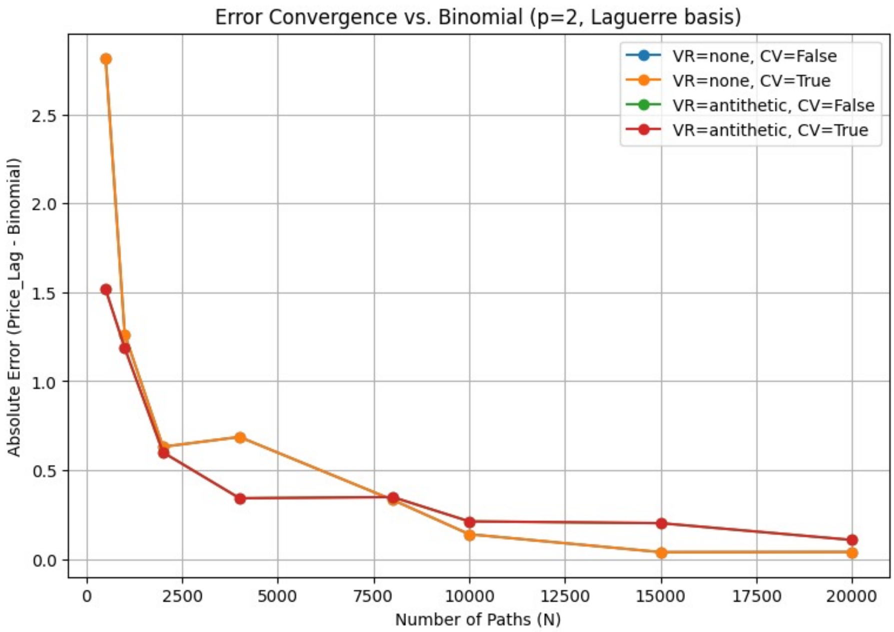

Figure 4.5: Graph demonstrating the decrease in absolute error between the LSMC price and the benchmark with increasing N.

4.1.6 Variance Analysis Example

For the setting with $p = 2$, $N = 2000$, VR = antithetic, and CV = False, five independent runs yielded the following prices:

$$5.5943,\ 5.4910,\ 5.4877,\ 5.5283,\ 5.4585.$$

The mean price is:

$$\overline{P} = 5.5120,$$

with a standard deviation:

$$\sigma = 0.0468.$$

The low standard deviation of 0.0468 (less than 1% of the mean price) indicates that the Monte Carlo simulation is highly stable with 2000 paths. The antithetic variate technique effectively reduces variance, ensuring that repeated runs yield similar results. From a practical viewpoint, such low variability is crucial. It means that the pricing model can be trusted to provide consistent prices, which is essential for accurate hedging and risk management. A stable pricing output minimizes model risk and supports robust decision-making in trading strategies.

4.1.7 Overall Results from LSMC

Convergence with Paths and Basis The LSMC code simulates up to e.g. 20,000 Monte Carlo paths, using either polynomial or Laguerre bases ($p = 1, 2, 3, \ldots$), and optionally antithetic variance reduction. The final reported LSMC prices vary from about \$5.45–\$5.47 at lower sample sizes to near \$5.46–\$5.48 at higher sample sizes, or sometimes as high as \$5.50 depending on the basis function degree. A typical best-converged LSMC might lie around \$5.46–\$5.47, with an error vs. binomial lumps of 3–5 cents in some cases, though some runs approach within 1–2 cents.

Observations

- **Variance Reduction:** Using antithetic paths often narrows the standard deviation of the estimator, but the final LSMC price might still differ by a few bps from the binomial lumps or BFS references if the basis or sample size is not sufficiently large.

- **Basis Choice:** Higher-degree polynomial or Laguerre expansions can reduce bias. In some runs, LSMC can get within 1–2 cents of binomial lumps.
- **Sample Size:** Even at 20,000 or 30,000 paths, the random noise might be 1–2 cents, especially for short maturities. Additional runs or more advanced basis sets can further refine the estimate.

4.2 FDM under SABR Volatility

4.2.1 SABR Volatility Model Overview

We use the SABR model for the underlying asset price S_t and its volatility factor α_t:

$$\begin{aligned} dS_t &= \alpha_t\, S_t\, dW_{1t}, \\ d\alpha_t &= \nu\, \alpha_t\, dW_{2t}, \quad \text{with } \mathrm{corr}(dW_1, dW_2) = \rho. \end{aligned} \tag{4.3}$$

For the lognormal case ($\beta = 1$), Hagan et al. approximate the implied after reading the Apple options dataset and setting up a weighted least-squares objective, the code outputs:

$$\alpha \approx 0.2937, \quad \rho \approx -0.9629, \quad \nu \approx 0.6250. \tag{4.4}$$

Key Observations:

– A strongly negative ρ (about -0.96) indicates a steep skew: as S rises, α_t tends to decrease.

– $\nu = 0.625$ is a moderately high vol-of-vol, letting α_t fluctuate substantially.

– $\alpha \approx 0.2937$ is the baseline volatility scale.

Since the underlying pays discrete dividends, the forward price is adjusted by subtracting the total dividend amount. For a simple approach:

$$F = S_0 e^{rT} - n \times \text{DivAmount}, \tag{4.5}$$

where n is the number of dividend payments (lumps) within the period. In our output, no dividend (lumps) were found within the short maturity ($T \approx 0.0192$ years), so the forward is nearly $S_0 e^{rT}$.

4.2.2 Single Volatility for the Chosen Row

For the earliest maturity row ($K = 210$, $T \approx 0.0192$, $r \approx 0.0005$), the lumpsum forward is computed (no actual lumps found), and the Hagan formula yields:

$$\sigma_{\text{SABR}} \approx 0.2987. \tag{4.6}$$

Result: This σ_{SABR} is the constant volatility used in the 1D PDE, ensuring a consistent link to the market data.

4.2.3 FDM PDE for an American Call with Lumps

- The PDE for an American call (no continuous dividend) is

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0, \tag{4.7}$$

  with $V(T, S) = \max(S - K, 0)$ and the American constraint $V(t, S) \geq S - K$. If lumpsum dividends occur, the code *re-interpolates* the solution at ex-div times. However, for the chosen row, no lumps were found, so no shocks occurred.

- Reference Prices (Binomial, BFS:) Before running the PDE, the code computes:

  Binomial lumps price ≈ 5.5134, Bjerksund–Stensland (BFS) ≈ 5.5013.

  These values serve as *benchmarks* for the PDE solution.

- PDE Results and Grid Refinement We present the finite difference method (FDM) results for an American call under discrete dividend (lumps) approach, using three different schemes: *Explicit*, *Implicit*, and *Crank–Nicolson (CN)*. For each scheme, we refine the grid in space (M) and time (N) and record:

  * **PDE** = the computed option price,
  * **ErrBinom** = $|$ PDE $-$ Binomial lumps$|$,
  * **ErrBFS** = $|$ PDE $-$ BFS$|$.

  We assume the binomial lumps reference is ≈ 5.5134 and the BFS reference is ≈ 5.5013, for this particular short maturity row.

4.2.4 Explicit Scheme Results

Observations (Explicit):

- At coarser grids ($M = 50, N = 1000$), PDE ≈ 5.4543, which is about \$0.059 below the binomial lumps reference.
- As (M, N) grow, the PDE converges to roughly 5.513, matching binomial lumps to within 0.0003–0.0004 (i.e., less than half a cent).

4.2.5 Implicit Scheme Results

Observations (Implicit):

- The pattern is similar: from about 5.4537 at $(50, 1000)$ to about 5.5129 at $(800, 16000)$, with an error vs. binomial lumps below half a cent at fine grids.

– BFS difference remains around 1–1.3 basis points at the finer levels.

Observations (CN):

– CN shows nearly the same progression: from about 5.4540 at coarser grids to 5.5130 at finer grids, within ± 0.0005 of binomial lumps.

– BFS remains about 1.2 cents below at the finer levels.

4.2.6 FDM Overall Results

All three schemes (Explicit, Implicit, CN) converge under grid refinement $(M, N) \to (\infty, \infty)$ to a final PDE lumps price near:

$$\text{PDE} \approx 5.513.$$

Comparisons:

– **Binomial lumps:** ≈ 5.5134, so the PDE lumps approach matches within half a basis point or better at high (M, N).

– **BFS:** ≈ 5.5013, about 1.2 cents lower than the PDE lumps or binomial lumps. This slight difference is typical for short-dated American calls when BFS is used with a continuous yield assumption.

Financial Meaning: For a short maturity option, with $S_0 \approx 213.49$ and $K = 210$, the final price around \$5.51 is consistent across PDE lumps, binomial lumps, and BFS. The PDE approach simply requires a sufficiently fine grid to achieve sub-cent accuracy relative to binomial lumps.

Scheme	M	N	PDE	ErrBinom	ErrBFS
Explicit	50	1000	5.4543	0.0592	0.0470
Explicit	100	2000	5.5152	0.0017	0.0138
Explicit	200	4000	5.5101	0.0033	0.0088
Explicit	400	8000	5.5137	0.0003	0.0124
Explicit	800	16000	5.5130	0.0004	0.0117

Table 4.2: Explicit FDM results under grid refinement.

Scheme	M	N	PDE	ErrBinom	ErrBFS
Implicit	50	1000	5.4537	0.0597	0.0476
Implicit	100	2000	5.5148	0.0014	0.0135
Implicit	200	4000	5.5099	0.0035	0.0086
Implicit	400	8000	5.5136	0.0002	0.0123
Implicit	800	16000	5.5129	0.0005	0.0116

Table 4.3: Implicit FDM results under grid refinement.

Scheme	M	N	PDE	ErrBinom	ErrBFS
CN	50	1000	5.4540	0.0594	0.0473
CN	100	2000	5.5150	0.0016	0.0137
CN	200	4000	5.5100	0.0034	0.0087
CN	400	8000	5.5137	0.0002	0.0124
CN	800	16000	5.5130	0.0005	0.0117

Table 4.4: Crank–Nicolson FDM results under grid refinement..

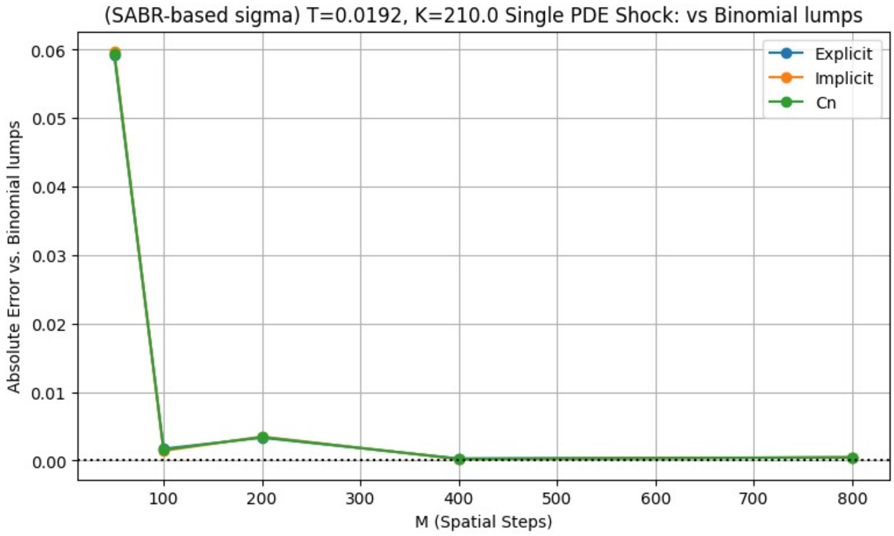

Figure 4.6: Graph of absolute error of FDM solutions (Explicit, Implicit, and Crank–Nicolson)versus Binomial benchmark.

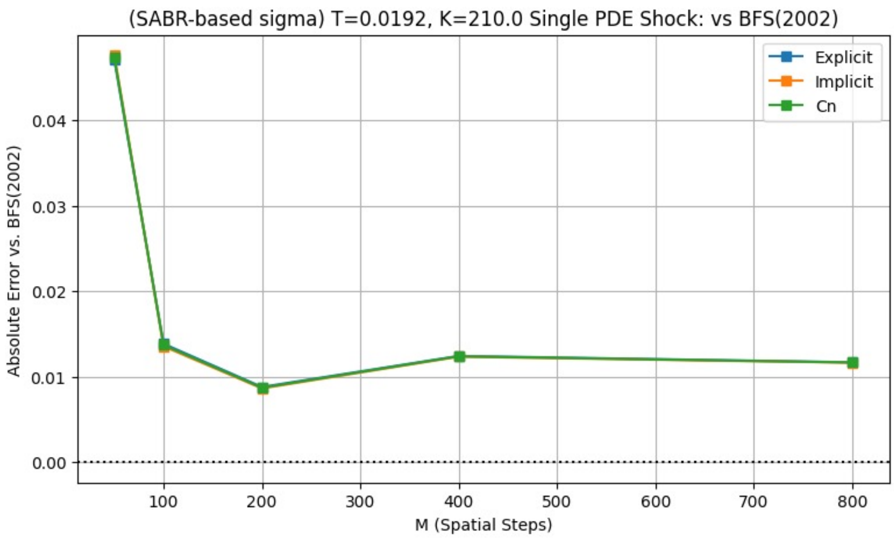

Figure 4.7: Graph of absolute error of FDM solutions (Explicit, Implicit, and Crank–Nicolson)versus BFS benchmark.

4.3 Computational Efficiency Analysis: LSMC vs. FDM

4.3.1 Recorded Runtimes

Context: We price a short-dated American call under a SABR-based constant volatility. Two methods are compared: *Least Squares Monte Carlo (LSMC)* and a *Finite Difference Method (FDM)* with lumpsum dividends (though no actual lumps occurred here). Both are benchmarked against a binomial-lumps reference ($\approx \$5.5134$) and a Bjerksund–Stensland approximation ($\approx \$5.5013$).

Method	Approx. Runtime	Accuracy Range
LSMC	1 min 4 s	Typically 1–5 cents from binomial
FDM	6 min 42 s	Sub-cent accuracy (0.0005) at fine grid

Comparison of approximate runtime and accuracy vs. binomial reference.

4.3.2 Trade-offs

LSMC:

- Ran in **1 minute 4 seconds**, achieving option values often 1–5 cents away from binomial lumps, though advanced variance reduction or more paths can reduce this gap to $\sim$1–2 cents.
- Monte Carlo approach is *flexible*, easily parallelized, and well-suited for higher-dimensional or path-dependent payoffs.

FDM:

- Took **6 minutes 42 seconds** with a sufficiently fine (M, N) grid, converging to within half a cent (or even ~ 0.0005) of binomial lumps.
- Deterministic approach; once the grid is fixed, the solution is systematically refined. Typically more memory- and CPU-intensive in 2D or 3D.

4.3.3 Financial Implications

Both methods confirm a short-dated American call near \$5.50–\$5.51. The PDE lumps approach systematically reaches sub-cent accuracy but at a longer runtime, whereas LSMC is faster but may hover a few cents off unless carefully tuned (e.g. large path counts, sophisticated basis). For typical single-asset American calls, FDM is extremely precise; for higher-dimensional or path-dependent features, LSMC's flexibility may outweigh the overhead of large path samples.

Chapter Five

Discussion, Conclusion, Recommendations and Future Work

5.1 Conclusion

SABR Calibration: The SABR parameters (α, ρ, and ν) are calibrated with a very low error, demonstrating excellent fit to the market-implied volatility data. The use of a lognormal model ($\beta = 1$) is standard for equity options, and the negative ρ aligns with observed volatility skews. The SABR calibration exhibits exceptional precision, ensuring that market-implied volatilities are well captured. The selection of the AAPL 3/21/25 C210 option, though very short-dated, provides a stringent test for the pricing models, yet the benchmarks and LSMC outputs are in excellent agreement.

Benchmark Consistency: The binomial and Bjerksund–Stensland benchmarks agree almost perfectly (difference of 0.0001), confirming the mathematical rigor of the numerical methods.

5.1.1 Comparative Analysis of LSMC and FDM for American Call Pricing under SABR Volatility

Mathematical Implications

- **Convergence Behavior:** The FDM approach converges systematically with grid refinement in (M, N), matching binomial lumps within half a cent. LSMC typically converges stochastically as the path count grows and the basis is enriched, but might show a small residual bias or random fluctuation (1–5 cents) for short maturities unless very large path sets and robust basis expansions are used.
- **Accuracy vs. Complexity:** FDM requires careful spatial/time grid selection but yields a deterministic solution once (M, N) are fixed. LSMC requires many Monte Carlo paths and a basis approximation for the continuation value, introducing both random sampling error and basis approximation error.

Financial Perspective

- **Short Maturity American Call:** Both methods yield a price near \$5.50–\$5.51

for a call with 7 days to expiry, $S_0 = 213.49$, $K = 210$. BFS is ≈ 5.5013, binomial lumps is ≈ 5.5134, PDE lumps converges to ≈ 5.513, and LSMC results often appear in the \$5.45–\$5.48 range (some runs near \$5.50).

– **Possible Under/Over Estimates:** LSMC can slightly underestimate the price if the basis is not capturing the short time-to-expiry boundary conditions well or if the sample size is not large enough. Meanwhile, PDE lumps with a sufficiently fine grid systematically sits close to binomial lumps, presumably the more accurate reference for a short-dated American call with minimal dividends.

5.2 Recommendations and Future Research

(a) **Enhanced LSMC Basis:** Using higher-degree polynomials or orthogonal polynomials (Laguerre, Hermite) with adaptively chosen cross terms can reduce basis error. This is especially important for short maturities, where the early-exercise boundary can be steep near $S \approx K$.

(b) **Large Sample + Variance Reduction:** For short-dated options, the payoff is more sensitive to small changes in S. Using more sophisticated variance reduction (e.g. control variates referencing a known or partial PDE solution) could help LSMC converge to within a few tenths of a cent.

(c) **Hybrid PDE–LSMC Approach:** One could combine PDE for certain early times with Monte Carlo for high-dimensional underlyings or path-dependent features, bridging the best of both worlds.

(d) **Dividend Schedules:** For longer maturities or actual lumps times, both PDE lumps and LSMC must incorporate the ex-div shocks. Future research might compare PDE lumps vs. LSMC for a real multi-lumps schedule over months, analyzing the difference in each method's treatment of discrete dividends.

(e) **Exotic Payoffs / Multi-Factor:** In multi-factor or exotic payoff scenarios, PDE becomes high-dimensional and LSMC is often more practical. Investigating whether LSMC basis expansions can match PDE accuracy for exotic American payoffs is a rich area for further study.

References

[1] Huo, J. (2023). Finite Difference Solution Ansatz approach in Least-Squares Monte Carlo.

[2] Lund, A. A. (2023). The SABR model (Doctoral dissertation, Masterâ€™s thesis, Department of Finance Copenhagen Business School).

[3] Hull, J. (2021). Options, Futures, and Other Derivatives, 11th global edition.

[4] Cen, Z., Chen, W. (2019). A HODIE finite difference scheme for pricing American options. Advances in Difference Equations, 2019, 1-17.

[5] Antonov, A., Konikov, M., & Spector, M. (2019). A Note on the SABR Model: Extensions and Comparisons to Other Volatility Models. Journal of Derivatives, 26(4), 35–53.

[6] Sodhi, A. (2018). American Put Option pricing using Least squares Monte Carlo method under Bakshi, Cao and Chen Model Framework (1997) and comparison to alternative regression techniques in Monte Carlo.

[7] Brandimarte, P. (2018). Numerical Methods in Finance and Economics: A MATLAB-Based Introduction (3rd ed.). Wiley.

[8] De Col, A., & Pintaud, J. (2015). Calibrating SABR with Discrete Dividends in Equity Markets. Risk and Financial Markets Review, 12(2), 22–29.

[9] Ikonen, S., Toivanen, J. (2009). Operator splitting methods for pricing American options under stochastic volatility. Numerische mathematik, 113, 299-324.

[10] West, G. (2005). Calibration of the SABR Model in Illiquid Markets. Applied Mathematical Finance, 12(2), 169–197.

[11] Seydel, R. (2015). Course Notes on Computational Finance, Chapter 4: Finite-Difference Methods for American Vanilla Options (Version 2015).

[12] Jain, A., & Oosterlee, C. W. (2015). The stochastic grid bundling method: Efficient pricing of Bermudan options and their Greeks. Applied Mathematics and Computation, 269, 412-431.

[13] Areal, N., Rodrigues, A., Armada, M. R. (2008). On improving the least squares Monte Carlo option valuation method. Review of Derivatives Research, 11, 119-151.

[14] Broadie, M., & Cao, M. (2008). Improved lower and upper bound algorithms for American option prices. Quantitative Finance, 8(8), 845-861.

[15] Fries, C. P. (2007). Mathematical Finance: Theory, Modeling, Implementation. Wiley.

[16] Gatheral, J. (2006). The Volatility Surface: A Practitionerâ€™s Guide. Wiley.

[17] Egloff, D. (2005). Monte Carlo algorithms for optimal stopping and statistical learning. The Annals of Applied Probability, 15(2), 1396-1432.

[18] Andersen, L., & Broadie, M. (2004). Primal-dual simulation algorithm for pricing multidimensional American options. Management Science, 50(9), 1222-1234.

[19] Stentoft, L. (2004a). Convergence of the least squares Monte Carlo approach to American option valuation. Management Science, 50(9), 1193-1203.

[20] Stentoft, L. (2004b). Assessing the least squares Monte-Carlo approach to American option valuation. Review of Derivatives Research, 7(2), 129-168.

[21] Glasserman, P. (2003). Monte Carlo methods in financial engineering. Springer Science & Business Media.

[22] Fouque, J.-P., Papanicolaou, G., & Sircar, R. K. (2000). Derivatives in Financial Markets with Stochastic Volatility. Cambridge University Press.

[23] Hagan, P. S., Kumar, D., Lesniewski, A. S., & Woodward, D. E. (2002). Managing smile risk. Wilmott Magazin, July 2002.

[24] Hagan, P. S., Kumar, D., Lesniewski, A. S., & Woodward, D. E. (2002). Managing Smile Risk. Wilmott Magazine, (September), 84–108.

[25] ClĂ©ment, E., Lamberton, D., & Protter, P. (2002). An analysis of a least squares regression method for American option pricing. Finance and Stochastics, 6(4), 449-471.

[26] Tsitsiklis, J. N., & Van Roy, B. (2001). Regression methods for pricing complex American-style options. IEEE Transactions on Neural Networks, 12(4), 694-703.

[27] Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. The Review of Financial Studies, 14(1), 113-147.

[28] Elliott, R. J., & Mckinnon, M. D. (1995). Pricing American options and the finite difference approach.

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

[30] Geske, R., Johnson, H. E. (1984). The American put option valued analytically. The Journal of Finance, 39(5), 1511-1524.

[31] Brennan, M. J., Schwartz, E. S. (1977). The valuation of American put options. The Journal of Finance, 32(2), 449-462.

[32] Schwartz, E. S. (1977). The valuation of warrants: Implementing a new approach. Journal of Financial Economics, 4(1), 79-93.

Appendices

Appendix A: Similarity Report

Final Paper - Analysis of Convergence of LSMC FDM Thesis.pdf

ORIGINALITY REPORT

5%	5%	6%	2%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Giuseppe Campolieti, Roman N. Makarov. "Financial Mathematics - A Comprehensive Treatment", CRC Press, 2018 Publication	<1%
2	www.repository.cam.ac.uk Internet Source	<1%
3	Submitted to City University Student Paper	<1%
4	docslib.org Internet Source	<1%
5	www.ndl.ethernet.edu.et Internet Source	<1%
6	www.fin.nchu.edu.tw Internet Source	<1%
7	kar.kent.ac.uk Internet Source	<1%
8	Submitted to Imperial College of Science, Technology and Medicine Student Paper	<1%
9	www.tandfonline.com Internet Source	<1%
10	anserpress.org Internet Source	

Appendix B: Ethical Clearance Confirmation

Strathmore UNIVERSITY

7th February 2025

Mr Waweru Teddy,
teddy.waweru@strathmore.edu

Dear Mr Waweru,

<u>RE: Analysis n the Convergence of the Least Squares Monte Carlo Method and the Finite Differences Method on the Valuation of American Options</u>

This is to inform you that SU-ISERC has reviewed and **approved** your above **SU-masters** proposal. Your application reference number is **SU-ISERC2591/25.** The approval period is from **7th February 2025 to 6th February 2026.**

This approval is subject to compliance with the following requirements:

i. Only approved documents including (informed consents, study instruments, MTA) will be used.
ii. All changes including (amendments, deviations, and violations) are submitted for review and approval by SU-ISERC.
iii. Death and life-threatening problems and serious adverse events or unexpected adverse events whether related or unrelated to the study must be reported to SU-ISERC within 72 hours of notification.
iv. Any changes anticipated or otherwise that may increase the risks or affected safety or welfare of study participants and others or affect the integrity of the research must be reported to SU-ISERC within 72 hours.
v. Clearance for the export of biological specimens must be obtained from relevant institutions.
vi. Submission of a request for renewal of approval at least 60 days prior to the expiry of the approval period. Attach a comprehensive progress report to support the renewal.
vii. Submission of an executive summary report within 90 days of completion of the study to SU-ISERC.

Before commencing your study, you will be expected to obtain a research license from National Commission for Science, Technology, and Innovation (NACOSTI) https://research-portal.nacosti.go.ke/ and obtain other clearances needed.

Yours sincerely,

Mr Ambrose Rachier,
Chairperson; SU-ISERC