American option pricing with LSM algorithm and analytic bias correction

Thesis paper

Written by: Balazs Kovacs

Applied mathematics

External Supervisor: Gabor Molnar-Saska

Starts and Modeling, Morgan Stanley

Internal Supervisor: Herczegh Attila

Department of Probability Theory and Statistics

Eotvos Lorand University, Faculty of Science

Eötvös Loránd University

2012

I am grateful to Gabor Molnar-Saska for his endless support and inspiration and thankful to Attila Herczegh for his useful remarks.

Contents

Chapter 1

Introduction

Modern financial markets offer a wide range of derivative products besides traditional assets such as bonds and equities. These instruments are traded in over-the-counter markets by investment banks to hedge funds, commercial banks and government entities. Investment banks hire mathematicians called strategists who are responsible for the pricing and risk management of these products mainly using tools from statistics, probability theory and stochastic calculus. This rapidly growing new branch of mathematics is often referred to as quantitative finance.

The pricing of American style contingent claims has been one of the most popular research topics for the last decades. The objective, particularly for an American option, is to determine the optimal exercise strategy that maximizes the payoff. This in fact is a challenging task given the stochastic nature of the underlying. As a result of the extreme development of information technology, simulation methods are gaining grounds in derivative pricing as well. LSM algorithm introduced by Longstaff and Schwartz is the most well known and widely used Monte Carlo method for calculating American option prices. LSM deals with the problem of re-simulation and provides a flexible and computationally tractable model for pricing. However, it does not consider the embedded foresight bias, which is a product of the backward nature of the calculation along with the Monte Carlo error of simulation.

My goal is to introduce a model based on LSM algorithm, that possesses all the good qualities of the original method and also addresses the problem of the look ahead bias. In the second chapter, I will review the basic financial definitions and the well known result of Black and Scholes on European option pricing, followed by a brief summary of modeling and pricing difficulties of American options. The third chapter begins with the presentation of the modeling framework and LSM algorithm. After a profound analysis of the least squares regression method I will attempt to study the impact of the look ahead, which causes the algorithm to be super-optimal. In the fourth chapter, motivated by Fries's concept, I develop a new six-step analytic approach to approximate the conditional expectation of continuation which also eliminates the foresight bias. Subsequently, I will present sufficient testing that justifies the accuracy of the new method, after which I conclude with areas of further research.

Chapter 2

Valuing options

I am going to start with a brief introduction to the most important definitions and concepts which are crucial to understand in the topic of derivative pricing. Quantitative finance is a relatively new branch of mathematics that is fundamentally based on probability theory, statistics and finance and it is mainly applied in the field of modeling stochastic processes and uncertainty in a financial market environment. This chapter summarizes the essential basics which are used in later chapters and strongly relies on the work of Hull [2011].

A financial derivative instrument defines a contract between two parties for future transactions of securities, assets or payments. Particularly an option specifies a deal on an asset at a reference price. The buyer of the option purchases the right, but not the obligation, to engage in the transaction, but once the option is exercised the seller of the contingent claim is obligated to fulfill the transaction. Based on the common characteristics of different type of payoff functions of the options, they are referred to as European, American, binary or barrier; just to name a few of the prevalent naming conventions.

2.1 European options

A European option provides the right but not the obligation, to buy or sell one unit of the underlying at a predetermined date at a reference price. This one exercise date is referred to as maturity, the underlying is generally a stock and the reference price is called the strike. Whether the option conveys the right to sell or to buy, it is refereed to as:

- European call option: provides the opportunity to buy one unit of the underlying at maturity, thus it's payoff function is $(S(T)-K)^+$ and it's value at time t is denoted by $C(S,t,T,K)$

- European put option: provides the opportunity to sell one unit of the underlying at maturity, thus it's payoff function is $(K - S(T))^+$ and it's value at time t is denoted by $P(S,t,T,K)$

where $S(t)$ is the underlying asset's spot price at t, T is the time of maturity and K is the strike. Call and put options with the same maturity and strike are in close relation on liquid markets, this is represented by the put-call parity equation:

$$P(S,t,T,K) = C(S,t,T,K) + K \cdot B(t,T) - S(t) \tag{2.1.1}$$

where $B(t,T)$ is the value of the bond at time t that matures at T. Furthermore, if the bond interest rate r, is assumed to be constant, then the above relationship simplifies to:

$$P(S,t,T,K) = C(S,t,T,K) + K \cdot e^{-r(T-t)} - S(t)$$

2.2 Valuing European options

Assume that the stock price dynamics is summarized by the following stochastic differential equation:

$$dS(t) = \mu \cdot S(t)dt + \sigma \cdot S(t)dW(t) \tag{2.2.1}$$

which is equivalent to the stock price following a geometric Brownian motion:

$$S(t) = S(0) \cdot exp\left[(\mu - \frac{\sigma^2}{2})t + \sigma W(t)\right] \tag{2.2.2}$$

where μ is the long term return, the drift term of the stock price and σ is the volatility of the underlying. Uncertainty is modeled by the increments of the standard Brownian motion.

A trading strategy in finance is a set of rules determining when to buy or sell instruments on the market. For instance, a delta neutral strategy aims to reduce the risk associated with the price movements of the underlying stock by maintaining 0 delta for the total portfolio. Delta is the measure of how the price of the derivative, particularly the option price changes as a result of a change in the price of the underlying. Mathematically speaking it equals to the partial derivative of the option's value with respect to the price of the underlying stock. Thus, such portfolio typically consists of an option and the underlying stock such that the positive and negative deltas offset; consequently, the portfolio value is unchanged by small price changes of the underlying. Since the delta hedge portfolio is risk-less, Black and Scholes [1973] argued that the rate of return of this strategy has to equal to the risk free rate otherwise this would provide an opportunity for a risk free profit or equivalently arbitrage. With the aid of Ito [1951] lemma this lead them to the well know partial differential equation published in 1973 referred to as the Black-Scholes equation:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0 \tag{2.2.3}$$

where V, the value of the derivative, has two arguments: time t and spot price of the underlying $S(t)$. Equation (2.2.3) holds true for any European style derivative whose payoff only depends on the value of the underlying at maturity, where $S(t)$ follows (2.2.2), i.e. it is a geometric Brownian motion, with constant drift and volatility. The value of a European option on a non-dividend paying stock is obtained by solving the above partial differential equation with the appropriate boundary conditions, specifically $V(T,S) = C(S,T,T,K) = (S(T) - K)^+$ for a call or $(K - S(T))^+$ for a put. Hence, the Black-Scholes formula for a European call option value is:

$$C(S,t,T,K) = \Phi(d_1)S(t) - \Phi(d_2)Ke^{-r(T-t)} \tag{2.2.4}$$

where $\Phi(\cdot)$ is the cumulative distributive function of the standard normal distribution and d_1 and d_2 are:

$$d_1 = \frac{ln(\frac{S(t)}{K}) + (r + \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}}$$

$$d_2 = \frac{ln(\frac{S(t)}{K}) + (r - \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}} = d_1 - \sigma\sqrt{T-t}$$

Then, I just simply use the put-call parity to derive the value of a European put option:

$$P(S,t,T,K) = C(S,t,T,K) + K \cdot e^{-r(T-t)} - S(t) =$$

$$\Phi(-d_2)Ke^{-r(T-t)} - \Phi(-d_1)S(t)$$

Consequently, based on the Black-Scholes model with constant drift and volatility, it is possible to derive an analytic and closed formula for calculating simple European option prices.

2.3 American options

An American style option defines a contract between two parties for a future transaction on an asset at a reference price. An American call or put option provides the right, but not the obligation, to buy or sell one unit of the underlying any time before or at maturity at a predetermined strike. Hence, as opposed to a European style option, American options have continuous exercise feature. This makes a distinct difference in the valuing process of these type of derivatives. The objective of my thesis is to provide a feasible unbiased and transparent model for pricing American style options that is also computationally tractable.

Since the Black-Scholes formula strongly relies on the assumption that there is one and only one predetermined exercise date, hence it is not feasible for valuing American options. Furthermore, there is no analytic formula for pricing options with continuous callable feature. The binomial option pricing model (BOPM) and simulation are the two most well known and widely used methods for pricing. In practice BOPM becomes less practical once options have several sources of uncertainty and complex features; however, it is a feasible approach for pricing if the risk-free interest rate r and volatility σ are assumed to be constants over time. Even though, the run time of the algorithm is $O(2^n)$, I am going to use the BOPM option price as one of the reference values for evaluating the convergence and accurateness of later introduced methods. On the other hand, Monte Carlo option pricing model has a number of advantages to traditional methods such as simplicity and flexibility, thus it is becoming more and more popular for valuing American style contingent claims.

Simulation methods start with identifying the underlying's dynamics. Let me assume that the stock price follows dynamics (2.2.1) of the Black-Scholes model. $dS(t) = \mu \cdot S(t)dt + \sigma \cdot S(t)dW(t)$ is under the real world physical measure. First fundamental theorem of asset pricing from Shreve [2004] states that a discrete market is arbitrage free if and only if there exists at least one risk neutral measure $\mathcal{Q}$ that is equivalent to the physical measure. Risk neutral or martingale measure is the measure under which the discounted underlying stock price process is a martingale. No arbitrage pricing theory states that the true value of an asset is the expectation of the sum of all future discounted cash flows generated by the asset with respect to the $\mathcal{Q}$ risk neutral measure. Shreve [2004] shows how to change from the real world measure to $\mathcal{Q}$ using Girsanov's theorem. The dynamics of the underlying under the risk neutral measure is:

$$dS(t) = r \cdot S(t)dt + \sigma \cdot S(t)d\tilde{W}(t) \tag{2.3.1}$$

where $\tilde{W}(t)$ denotes the Brownian motion under the martingale measure $\mathcal{Q}$. This is particularly convenient because for the discounted stock price dynamics which is:

$$d\tilde{S}(t) = e^{-rt} \cdot dS(t)$$

it follows that:

$$d\tilde{S}(t) = \sigma \cdot \tilde{S}(t)d\tilde{W}(t) \tag{2.3.2}$$

consequently, $\tilde{S}(t)$ is of course a martingale. From here on, I assume that the existence of an equivalent risk neutral measure $\mathcal{Q}$ and that all dynamics are given under this probability measure. Once the dynamics of the underlying is fixed, its is possible to use Monte Carlo simulation to generate a finite number of realizations of the stock price process. The American option price is then computed by the best possible stopping strategy that maximizes the payoff over all exercise policies. To get an accurate estimate of the continuation value or in other words the option value on a particular path at time t, the underlying's price process should be re-simulated using all the available information at that future trading event, such as the spot price of the underlying, which is of course different from the initial value. The chart below illustrates this concept.

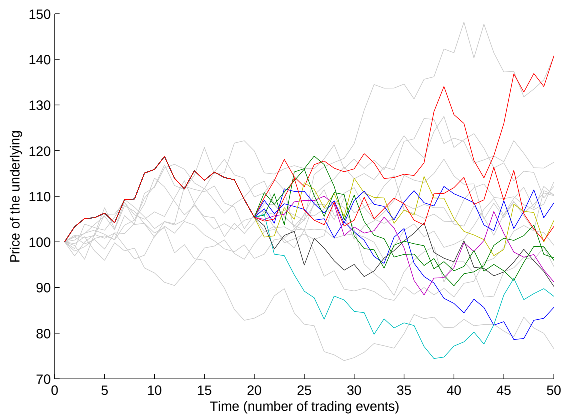

Figure 2.3.1: Brute Force re-simulation

However, this concept leads to exponential number of simulations which is computationally intractable; consequently, a different method has to be used to value American style options. The remainder of the paper is organized as follows. In chapter 3, I am going to introduce a well known approach that addresses the problem of re-simulation. LSM algorithm uses one set of Monte Carlo simulation to estimate the option value; however, this simplicity results in an embedded foresight that causes the algorithm to be biased. In chapter 4, I am going to introduce, analyze and test a new analytic method for bias correction. After pointing out areas of further research, I finish my thesis with concluding remarks.

Chapter 3

Analysis of LSM algorithm

Valuing American Options by Simulation: A Simple Least-Squares Method (LSM) was published by Longstaff and Schwartz [2001]. They present a simple yet powerful algorithm to provide a path-wise approximation of the optimal exercise rule. The objective is to price an American-style derivative security by maximizing its random cash-flows occurring in a finite time frame. This simulation approach has many advantages compared to alternative valuing methods. First, simulation allows the underlying to follow more complex stochastic process such as geometric Brownian motion, jump diffusion or Levy process. Second, the method provides the flexibility to use multiple factors or regressors to determine the value of the option. Third, the algorithm can easily handle complex payoffs and path dependent derivatives and it can also be used for sensitivity tests for risk management purposes. Last, simulation is well suited for parallel computing, thus makes it computationally attractive as well.

3.1 Modeling framework

Let $(\Omega, \mathcal{F}, \mathcal{P})$ be a complete probability space, where Ω is the set of all possible realizations of the stochastic stock price process on the finite time horizon of $[0, T]$. Let $\mathcal{F}(\mathrm{t})$ be the $\sigma - algebra$ filtration generated by the different events until time t and $\mathcal{F}$ denotes the final stage of information $\mathcal{F}(\mathrm{T})$, where $\mathcal{F}(0) \subseteq \mathcal{F}(1) \subseteq \ldots \subseteq \mathcal{F}(\mathrm{t}) \subseteq \ldots \subseteq \mathcal{F}(\mathrm{T})$. Furthermore $\mathcal{P}$ is the probability measure defined on ω elements of Ω. Based on the no-arbitrage pricing theory I assume the existence of an equivalent martingale measure $\mathcal{Q}$. The algorithm is suitable for derivatives with payoffs that are elements of the space of square-integrable or finite-variance functions of $\mathcal{L}^2(\Omega, \mathcal{F}, \mathcal{Q})$. Consequently, LSM is feasible to price most derivatives with callable or cancel-able features; however, from here on I am going to restrict my attention to American put and call options. The objective of the algorithm is to approximate the optimal path-wise stopping rule that maximizes the value of the American option. For modeling and implementational purposes, I assume that the option is only exercisable at N discrete times $0 = t_0 < t_1 < t_2 < \ldots < t_N = T$, this might seems like a Bermudian setup, but selecting N large enough makes this approach suitable to value continuously exercisable American options as well.

At the final exercise date, the strategy of the buyer is straightforward, either exercise the option if it is in-the-money which means positive intrinsic value or allow it to expire if out-of-money which implies no intrinsic value. Prior to expiry, at each t_j exercise date the investor has to decide whether or not to exercise the option. This is done by comparing the immediate exercise value, known at t_j, with the expected value of continuation, which is the sum of discounted future cash flows, and then exercise if immediate exercise is equal or more valuable. Therefore, it is fundamental to have an accurate estimate of the expected value of continuation, which is of course a random variable at t_k. No-arbitrage principle for derivative pricing states, that the value of continuation is given by the expectation of all future discounted cash flows with respect to the equivalent risk-neutral pricing measure Q. Let $CF(\omega, t_j; t_k, T)$ denote the different future cash flows generated by the security by following the optimal stopping rule at the different t_j exercise times, where $t_k < t_j \leq t_N = T$, and conditional on that the option is not exercised at or before t_k. Consequently, the path-wise value of the option assuming that it can only be exercised after t_k or equivalently, the value of continuation at t_k is given by:

$$F(\omega; t_k) = E_Q \left[\sum_{j=k+1}^{N} exp(-\int_{t_k}^{t_j} r(\omega, s) ds) \cdot CF(\omega, t_j; t_k, T) | \mathcal{F}(\mathrm{t_k}) \right] \tag{3.1.1}$$

where $r(\omega, t)$ is the risk-less discount rate and the expectation conditioned on $\mathcal{F}(\mathrm{t_k})$, the set of information obtained until t_k, with respect to the risk-neutral measure. Let me note that in the specific case of American options, at most there is one none zero $CF(\omega, t_j; t_k, T)$ value, where $t_k < t_j \leq t_N = T$ for any t_k, this is because the option is only callable once over its lifespan. The price of the American option is then calculated by averaging $F(\omega, 0)$ over all ω paths. In the following section, I will introduce a well known theoretical result that justifies this approach.

3.2 Snell-envelope

Let τ_{t_k} denote all stopping times, where each $\zeta \in \tau_{t_k}$ is an element of $\{t_k, t_{k+1}, \ldots, t_N\}$. The following standard result is from Bensoussan [1984] and Karatzas [1988]:

3.2.1 Theorem(Snell-envelope): Let $t_0 < t_1 < \ldots < t_k < \ldots < t_N$ be a set of discrete trading events, $X(t_k)$ denote the value of an American option's immediate payoff at time t_k and $Z(t_k)$ be a discrete stochastic process, where $Z(t_N) = X(t_N)$ and for $\forall t_k < t_N$, $Z(t_k)$ is defined recursively such that:

$$Z(t_k) = max\left[X(t_k), E\left(Z(t_{k+1})|\mathcal{F}(\mathrm{t_k})\right)\right]$$

Furthermore, let ν_{t_k} denote the following stopping rule:

$$\nu_{t_k} = min\left[t_j \geq t_k : X(t_j) = Z(t_j)\right]$$

Then the following the propositions hold true:

1. $\nu_{t_k} \in \tau_{t_k}$
2. $Z(t_k) = E\left(X(\nu_{t_k})|\mathcal{F}(\mathrm{t_k})\right) = sup_{\zeta \in \tau_{t_k}} E\left(X(\zeta)|\mathcal{F}(\mathrm{t_k})\right)$
3. $E(X(\nu_{t_k})) = E(Z(t_k)) = sup_{\zeta \in \tau_{t_k}} E(X(\zeta))$

Practically, ν_{t_k} maximizes the option value conditioned on the fact that it is not exercised before t_k, based on the no-arbitrage paradigm, $Z(t_k)$ equals to the value of the option at time t_k. For $t_0 = 0$ special case of the above theorem, $E(X(\nu_0)) = E(Z(0)) = sup_{\zeta \in \tau_0} E(X(\zeta))$. This means that ν_0 is the optimal stopping rule, which maximizes the value of the American option out of all stopping rules on $\{t_1, \ldots, t_N\}$. Furthermore; based on ν_{t_k}'s definition the optimal exercise policy is the first time when $X(t_k) = Z(t_k)$ or equivalently the first time when $X(t_k) \geq E\left(Z(t_{k+1})|\mathcal{F}(\mathrm{t_k})\right)$, in other words the option has to be exercised the first time when the immediate exercise is equal or greater than the conditional expectation of continuation.

3.3 LSM algorithm

LSM algorithm uses the cross-sectional information in the simulated paths to approximate the conditional expectation function. This is carried out by regressing the sum of all future discounted cash flows on a set of basis functions of relevant state variables. The fitted value of this regression is an accurate estimate of the conditional expectation function, which provides the path-wise continuation value. Thus, it can be used to compute the optimal stopping strategy which is the objective of the algorithm.

If the conditional expectation function is an element of the $\mathcal{L}^2(\Omega, \mathcal{F}, \mathcal{Q})$ square-integrable functions, then since $\mathcal{L}^2$ is a Hilbert space, it has a countable orthonormal basis and thus $F(\omega, t_k)$ the conditional expectation function can be represented as the linear combination of the basis, a countable set of $\mathcal{F}(\mathrm{t_k})$-measurable functions. Typical types of basis functions include the Laguerre, Hermite, Legrendre, Chebysew and Jacobi polynomials. In practice often times it is necessary to use several indicators to sufficiently describe the current state of a more complex derivative, hence multiple number of state variables are needed for the approximation. For sake of simplicity let me assume, a two dimensional setup, where x and z are the state variables, then the set of basis functions should include terms in each variable and cross products of them as well. As a result of this specification, $F(\omega, t_k)$ can be represented as:

$$F(\omega, t_k) = \sum_{i=0}^{\infty} \beta_i B_i(x, z)$$

where B_i denotes the different basis functions and β_i coefficients are constants, this is of course path dependent as it is indicated on the left hand side, x and z are also functions of ω as they may differ from path to path. Assuming that the underlying follows a Markov process, hence past realizations are irrelevant towards determining the future path of the asset; therefore, the spot price alone is a sufficient state variable for American options. In addition, Judd [1998] showed that the number of basis functions need not grow exponentially, in fact numerical tests suggest increasing their number only polynomially with the dimension of the problem is sufficient to obtain convergence even in higher degree cases. Furthermore, in practice even simple powers of state variables give accurate results. Thus, $B_i(x) = x^i$ is a possible choice of basis functions, specifying the following fitted value :

$$\hat{F}_M(\omega, t_k) = \hat{\beta}_0 1 + \hat{\beta}_1 x + \hat{\beta}_2 x^2 + \ldots + \hat{\beta}_{M-1} x^{M-1} \tag{3.3.1}$$

where M denotes the number of basis functions. Given the backwards nature of the algorithm, at any given t_k time, the expectation of $CF(\omega, t_j; t_k, T)$ is known for each path. $F(\omega, t_k)$ is then estimated by regressing the discounted values of $CF(\omega, t_j; t_k, T)$ on the set of basis functions, hence ultimately by $F_M(\omega, t_{N-1})$. Since for out of the money paths it is never optimal to exercise the option, no exercise decision has to be made; consequently, LSM only uses in-the-money paths for the regression. This limits the region over which the conditional expectation function has to be determined, thus yielding more accurate results with even fewer number of basis functions. Since the values of basis functions are independently and identically distributed across all paths, the result of White [1984] states that the fitted value of the projection $\hat{F}_M(\omega, t_k)$ converges in mean square and in probability to $F_M(\omega, t_k)$ if the number of in-the-money paths go to infinity. Furthermore, Amemiya [1985] implies that $\hat{F}_M(\omega, t_k)$ is the best linear unbiased estimator of $F_M(\omega, t_k)$ based on the mean-squared metric.

The algorithm works as follows: once the conditional expectation function at time t_{N-1} is determined, it is straightforward to determine the optimal stopping strategy for all the in-the-money paths as well, simply stopping at each ω path, where the immediate exercise is equal or greater than the fitted value of continuation $\hat{F}_M(\omega, t_k)$. Now that the cash flows for t_{N-1} are determined, after appropriate discounting, these values can be regressed on a set of basis functions of state variables of time t_{N-2}, this provides an accurate estimate of the continuation function at t_{N-2}, repeating this procedure until the stopping rule for all exercise times over all paths are determined. The American option value is then calculated by finding the first stopping time on each path and discounting the indicated cash flow back to time zero and then taking the average over all ω paths.

3.4 Convergence results

This part presents two results from Longstaff and Schwartz [2001] on the theoretical convergence of the algorithm; however, the best test of the performance of the algorithm is in practice with realistic number of paths and basis functions.

3.4.1 Theorem: For any finite choice of M, N and $\beta \in \mathbb{R}^{\mathrm{M}\times\mathrm{N}}$ representing the coefficients for the M basis functions at each of the N exercise dates, let $LSM(\omega, M, N)$ denote the discounted cash flow resulting from following the LSM rule of exercising when the immediate exercise value is positive and greater than or equal to $\hat{F}_M(\omega, t_k)$ as defined by β. Then the following inequality holds almost surely:

$$V(x) \geq \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} LSM(\omega_i, M, N)$$

where ω_i denotes the ith trajectory. $V(x)$ represents the true value of the American option, this is calculated with the stopping rule that maximizes $V(x)$ out of all stopping rules. Heuristically this result means that if the number of paths go to infinity, the American option value implied by LSM algorithm is less than or equal to that implied by the optimal stopping rule. This provides an objective criterion for convergence and it is particularly useful since it provides guidance for determining the number of basis functions, increase M until the value implied by LSM increases.

3.4.2 Theorem: Assume that the value of an American option depends on a single variable x with support on $(0, \infty)$ which follows a Markov process. Assume further that the option can only be exercised at time t_1 and t_2, and that the conditional expectation function $F(\omega, t_1)$ is absolutely continuous and

$$\int_0^\infty e^{-x} F^2(\omega, t_1) dx < \infty$$

$$\int_0^\infty e^{-x} F_x^2(\omega, t_1) dx < \infty$$

Then for any $\epsilon > 0$, there exists an $M < \infty$ such that:

$$\lim_{n\to\infty} P\left[|V(x) - \frac{1}{n} \sum_{i=1}^{n} LSM(\omega_i, M, N)| > \epsilon \right] = 0$$

The intuition for this result is that by selecting M large enough and letting $n \to \infty$, LSM results in a value for the American option within ϵ of the true value if only two trading events are assumed. Thus, for this particular set up LSM converges to any desired accuracy since ϵ is arbitrary. The fundamental reason behind is that the convergence of $F_M(\omega, t_1)$ to $F(\omega, t_1)$ is uniform on $(0, \infty)$ when the above integral conditions are met. This bounds the maximum error in estimating the conditional expectation and consequently the maximum pricing error as well. From a technical perspective, the number of basis functions needed to obtain a desired level of accuracy need not go to infinity as $n \to \infty$. Even though, this result is a one dimensional setup, similar result can be achieved for higher dimensional problems by meeting the uniform convergence condition.

3.5 Least squares regression

Least squares regression method is a approach used to approximate the solution of overdetermined systems such as $\sum_{j=1}^{M} X_{ij}\beta_j = y_i$ for each $i = 1, \ldots, n$ where $n > M$. This is done by determining the best-fitting curve to the set of points by minimizing the sum of the squares of the offsets. The offset r, commonly referred to as residual is the difference between the observed sample and the fitted value of the regression, so in matrix form it is:

$$X\hat{\beta} + r = y \tag{3.5.1}$$

where $\hat{\beta}$ contains the fitted values determined by the regression. This projection value is of course generally different from the theoretical solution of the overdetermined system, thus the following holds true

$$X\beta + \epsilon = y \tag{3.5.2}$$

where ϵ is the unobservable error, the difference between the sample and the theoretical, unknown solution of the system. To clarify on the notation β is the vector of the unknown theoretical parameters whereas $\hat{\beta}$ is the vector of the fitted parameters, furthermore let me introduce the following notation, let $Y = X\beta$ and $\hat{Y} = X\hat{\beta}$.

Generally the method starts with obtaining a set data points such as (x_i, y_i), $i = 1, \ldots, n$, where x_i is an independent variable and y_i is the dependent variable, this is done by observation. In the case of LSM x_i is equivalent to the spot price on the ith trajectory. Suppose (3.3.1) is the set of basis functions being used, then x is the second column of the X matrix, which is different for each t_k trading event. On the other hand, y_i is the corresponding discounted cash flow value at t_{k+1}, conditioned on the fact that the option is not exercised at or before t_k, which equals to $e^{-r(t_j - t_k)} \cdot CF(\omega, t_j; t_k, T)$, the discounted intrinsic value of the first indicated exercise time after t_k. These figures are of course known at t_k, thus they are observable. The earlier introduced $\hat{F}_M(\omega, t_k)$ regressor function's coefficients, the M adjustable parameters determined by the regression are held in the $\hat{\beta}$ vector. The objective is to find the $\hat{F}_M(\omega, t_k)$ curve defined by $\hat{\beta}$, that minimizes the sum of the squared residuals, specifically:

$$R \doteq \sum_{i=1}^{n} r_i^2 = \sum_{i=1}^{n} (y_i - \hat{F}_M(\omega, t_k))^2 \to min \tag{3.5.3}$$

The minimum is found by setting the gradient to zero.

$$\frac{\partial R}{\partial \beta_j} = 2 \cdot \sum_{i=1}^{n} r_i \frac{\partial r_i}{\partial \beta_j} = 0$$

for each $j = 1, \ldots, M$. The objective is that $\hat{\beta}$ minimizes R, hence the coefficients of $\hat{\beta}$ has to be chosen as follows. Since the value of the residuals are $r_i = y_i - \sum_{j=1}^{M} \hat{\beta}_j X_{ij}$, for the derivatives it follows that:

$$\frac{\partial r_i}{\partial \hat{\beta}_j} = -X_{ij}$$

thus the equation becomes:

$$2 \cdot \sum_{i=1}^{n} \left(y_i - \sum_{k=1}^{M} \hat{\beta}_k X_{ik} \right) (-X_{ik}) = 0$$

Rearranging the equation, we obtain:

$$\sum_{i=1}^{n} \sum_{k=1}^{M} X_{ij} X_{ik} \hat{\beta}_j = \sum_{i=1}^{n} X_{ij} y_i$$

Which is equivalent with the following matrix notation:

$$(X^T X)\hat{\beta} = X^T y$$

Thus, the $\hat{\beta}$ solution that minimizes (3.5.3) is:

$$\hat{\beta} = (X^T X)^{-1} X^T y \tag{3.5.4}$$

The linear regression model is summarized by $y = X\beta + \epsilon$, where ϵ is unknown random error that follows a normal distribution. Furthermore each ϵ_i, ϵ_j $(i \neq j)$ pair is independent and identically distributed with the following property:

$$\epsilon \sim N(0, \sigma^2 I) \tag{3.5.5}$$

Since the errors are unobservable, their analysis must be done indirectly using residuals:

$$r = y - X\hat{\beta} = y - X(X^T X)^{-1} X^T y$$

Let V denote the $X(X^T X)^{-1} X^T$ matrix and I am going to use the $y = X\beta + \epsilon$ substitution:

$$r = y - Vy = (I - V)y = (I - V)(X\beta + \epsilon) =$$

$$X\beta - X(X^T X)^{-1} X^T X\beta + (I - V)\epsilon = X\beta - X\beta + (I - V)\epsilon = (I - V)\epsilon$$

Consequently, the relationship between the errors and residuals only depend on V, usually referred to as the hat matrix, this is summarized by the following:

$$r = (I - V)\epsilon \tag{3.5.6}$$

$$r_i = \epsilon_i - \sum_{j=1}^{n} v_{ij}\epsilon_j$$

The hat matrix is symmetric and idempotent, these special properties simply follow from linear algebraic transformations.

$$V^T = (X(X^TX)^{-1}X^T)^T = X(X^TX)^{-1}X^T = V$$

$$V^2 = X(X^TX)^{-1}X^TX(X^TX)^{-1}X^T = X(X^TX)^{-1}X^T = V$$

When calculating the distribution of the residuals from the distribution of the errors, the matrix transformation does not change the distribution family. Thus the residuals, similarly to the ϵ errors, will also follow a normal distribution. Using (3.5.6) to determine the specific distribution properties, the expected value is $E(r) = (I - V)E(\epsilon) = 0$ and the variance is calculated using V's symmetry and idempotency:

$$var(r) = (I - V)\sigma^2(I - V)^T = \sigma^2(I - V)(I - V) = \sigma^2(I - V - V + V^2) = \sigma^2(I - V)$$

As a result, the distribution of r is:

$$r \sim N(0, \sigma^2(I - V)) \tag{3.5.7}$$

Furthermore, it is possible to show that the expected value of r_i is:

$$E(r_i) = E(\epsilon_i - \sum_{j=1}^{n} v_{ij}\epsilon_j) = E(\epsilon_i) - \sum_{j=1}^{n} v_{ij}E(\epsilon_j) = 0$$

Since ϵ_i is independent from ϵ_j if $i \neq j$, for the variance of r_i is holds that:

$$E(r_i - Er_i)^2 = E(r_i^2) = E(\epsilon_i - \sum_{j=1}^{n} v_{ij}\epsilon_j)^2 = E(\epsilon_i^2 - 2\epsilon_i \sum_{j=1}^{n} v_{ij}\epsilon_j + \sum_{j=1}^{n} v_{ij}^2\epsilon_j^2) =$$

$$E(\epsilon_i^2 - 2v_{ii}\epsilon_i^2 + \sum_{j=1}^{n} v_{ij}^2\epsilon_j^2) = E(\epsilon_i^2) - 2v_{ii}E(\epsilon_i^2) + \sum_{j=1}^{n} v_{ij}^2E(\epsilon_j^2) =$$

$$\sigma^2 - 2v_{ii}\sigma^2 + \sum_{j=1}^{n} v_{ij}^2\sigma^2 = \sigma^2(1 - 2v_{ii} + \sum_{j=1}^{n} v_{ij}^2) \tag{3.5.8}$$

Using the symmetry and idempotency of V, it follows that $v_{ii} = \sum_{j=1}^{n} v_{ij}v_{ji} = \sum_{j=1}^{n} v_{ij}^2$, substituting this into (3.5.8):

$$\sigma^2(1 - 2v_{ii} + v_{ii}) = \sigma^2(1 - v_{ii})$$

Summarizing the above calculations, the residuals follow a normal distribution and the variance of each individual r_i is:

$$var(r_i) = \sigma^2(1 - v_{ii}) \tag{3.5.9}$$

Now that the variance of the residuals are known, hence it is possible to examine the variance of the fitted values. Using the definition of the residuals (3.5.1) and errors (3.5.2):

$$var\left(\hat{Y}_i\right) = var(y_i - r_i) = var(Y_i + \epsilon_i - r_i)$$

Since Y_i is a theoretical constant, it's variance is zero:

$$var(\hat{Y}_i) = var(\epsilon_i - r_i) = var(\epsilon_i) + var(r_i) - 2cov(\epsilon_i, r_i) =$$

$$\sigma^2 + \sigma^2(1 - v_{ii}) - 2E\left(\epsilon_i r_i - E(\epsilon_i)E(r_i)\right) = \sigma^2 + \sigma^2(1 - v_{ii}) - 2E(\epsilon_i r_i)$$

Using the earlier proved relationship (3.5.6) between the errors and the residuals and the fact that the expected value of the cross product of the different ϵ_i and ϵ_j variables are zero, since if $i \neq j$ then these random variables are independent, it follows:

$$\sigma^2 + \sigma^2(1 - v_{ii}) - 2E\left(\epsilon_i(\epsilon_i - \sum_{j=1}^{n} v_{ij}\epsilon_j)\right) = \sigma^2 + \sigma^2(1 - v_{ii}) - 2E(\epsilon_i^2 - v_{ii}\epsilon_i^2) =$$

$$\sigma^2 + \sigma^2(1 - v_{ii}) - 2\sigma^2(1 - v_{ii}) = v_{ii}\sigma^2$$

Hence, the variance of the fitted value is:

$$var(\hat{Y}_i) = v_{ii}\sigma^2 \tag{3.5.10}$$

Since v_{ii} ranges on $[0, 1]$ then if v_{ii} is small, the variance of the fitted value is small and the variance of the residual is large and vice versa. Consequently, better understanding of the hat matrix V and especially the magnitude of the diagonal elements is essential to evaluate the accuracy of the regression. For this purpose, Cook and Weisberg [1982] suggest to divide X into the sum of two projections such that $X = (X_1, X_2)$, where X_1 is an $n \times q$ matrix and its rank is q. Furthermore, let U be $X_1(X_1^T X_1)^{-1} X_1^T$ and let X_2^*

equal to $(I-U)X_2$, the component of X_2 that is orthogonal to X_1. In this setup:

$$T^* = X_2^*(X_2^{*T}X_2^*)^{-1}X_2^{*T} = (I-U)X_2(X_2^T(I-U)X_2)^{-1}X_2^T(I-U)$$

is the operator that projects onto the subspace of X_2 and thus V can be calculated as:

$$V = U + T^* \tag{3.5.11}$$

The first column of the X matrix in LSM is the 0 degree polynomial of the chosen kind or equivalently constant 1. Consequently, to make practical use of the above calculations, let X_1 be a vector of ones, hence from (3.5.11) it follows that

$$V = \mathbf{1}/n + X_2^*(X_2^{*T}X_2^*)^{-1}X_2^{*T}$$

where $\mathbf{1}$ denotes the vector of ones and thus for v_{ii}:

$$v_{ii} = \frac{1}{n} + x_i^T(X_2^{*T}X_2^*)^{-1}x_i$$

where x_i^T is the ith row of X_2^*. Let $\mu_1 \geq \mu_2 \geq \ldots \geq \mu_n$ denote the eigenvalues of $X_2^{*T}X_2^*$ and let $p_1, \ldots, p_n$ denote the corresponding eigenvectors. Cook and Weisberg [1980] previously showed that, assuming the intercept is included in the model, then using the spectral decomposition of the corrected cross product matrix, for v_{ii} it follows:

$$v_{ii} = \frac{1}{n} + \sum_{l=1}^{p}\left(\frac{p_l^T x_i}{\sqrt{\mu_l}}\right)^2$$

Further letting θ_{li} denote the angle between p_l and x_i, then

$$cos(\theta_{li}) = \frac{p_l^T x_i}{(x_i^T x_i)^{1/2}}$$

Thus, for v_{ii}:

$$v_{ii} = \frac{1}{n} + x_i^T x_i \sum_{l=1}^{p} \frac{cos^2(\theta_{li})}{\mu_l}$$

Hence, v_{ii} is large if:

1. $x_i^T x_i$ is large, which is equivalent with, x_i is well removed from the bulk of the cases or

2. x_i has similar direction as an eigenvector corresponding to a small eigenvalue of $X_2^{*T}X_2^*$.

From this it follows that fitted values at remote places will have relatively large variances whereas the corresponding residuals will have small variances. This is intuitively justifiable since at remote places the number of samples are much lower, thus the regression will fit these points better, resulting in small residual variances for these cases. However, this implies that these fitted points will have relatively larger errors compared to the bulk of the cases.

3.6 The bias of LSM

Let me apply the results of the previous section for LSM algorithm for the specific case of valuing an American call option. Since the discounted stock price process is a martingale under the equivalent risk natural measure $\mathcal{Q}$, cases remote in the factor space correspond to the paths where the underlying obtains relatively large or small values. Let me examine an extreme path that realizes one of the largest values at maturity.

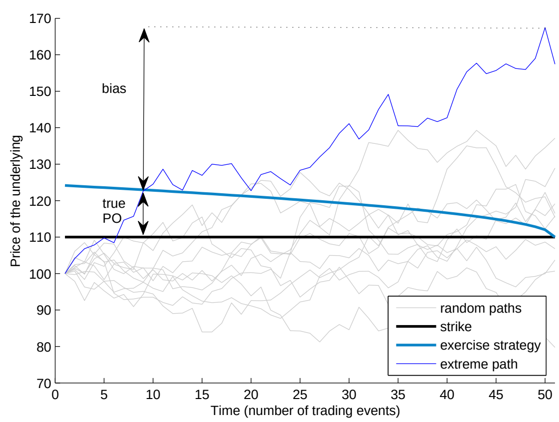

Figure 3.6.1: Extreme path

The bold black line indicates the strike whereas the bold blue curve is a possible selection for an exercise strategy, namely exercise the first time when the stock price hits the curve. Since LSM works backwards, as the algorithm starts it immediately obtains the huge payoff value of exercise at maturity. In the hypothetical case of the above example, the blue path is well removed from the bulk of the cases, hence the regression will have relatively small variance and will fit the independent observed variable very well. Hence, the approximation of the conditional expectation will be biased as this extreme case dominates the region, resulting the algorithm to benefit from unavailable information in a real life setup. Thus in general, working backwards, on these extreme paths the algorithm will always suggest to keep going and not to exercise the option since the conditional expectation of the future cash flows are larger than the value of current exercise. For example, the first time the above extreme path hits the exercise boundary, there are no other paths in the $(110, 170]$ range so the regression will indicate much higher payoff for continuation based on the future path of this particular realization. Consequently, the exercise strategy implied by LSM is biased by the look ahead of the algorithm. This problem does not arise in the bulk of the cases since there are sufficient number of paths to average out individual properties typically belonging to one particular realization of the underlying. As a result of this foresight bias, for finite number of paths, LSM results in a value greater than the true value of the American option.

The chart below plots the American option value determined by LSM as a function of the number of paths used in the simulation, whereas the yellow shaded area is the variance of the simulation. The downward sloping curve implies that as the number of paths increases the above described bias decreases making the value of the option calculated by the algorithm decrease as well. For detailed information on the parameters of the simulations presented, please refer to the Appendix.

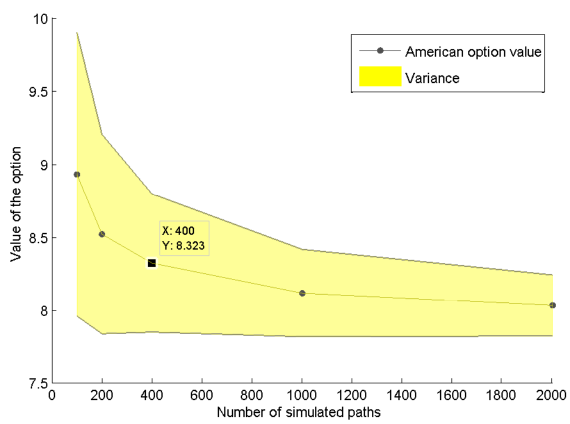

Figure 3.6.2: LSM American option value

There are a number of different solutions to this problem. One alternative approach is to simulate a set of paths to determine the optimal stopping rule and then to use this strategy on an independent set of new paths. Thus, eliminating all extra biased information from the stopping rule. In practice this and slight modifications of this method is the most widely used approach to eliminate the look ahead of the algorithm. The chart below shows the same numerical test as above for the independent path method. Here the curve is upward sloping, the fundamental behind this phenomenon is that the stopping rule, studied on an independent set of realizations, gets more and more accurate and relevant as the number of paths increases. In addition to the fact that Monte Carlo simulation method also yields more reliable results as the number of paths increases and variance decreases.

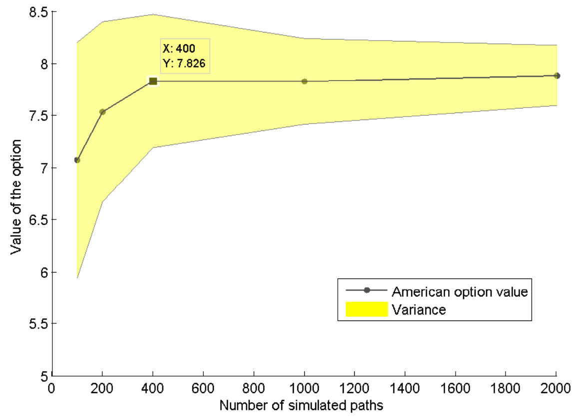

Figure 3.6.3: Independent paths method option value

From a theoretical point of view another approach is increasing the number of paths to infinity. This will ease the problem because even at remote places there are going to be enough number of samples to have an unbiased estimate of the conditional expectation function of continuation. Even though, this approach is obviously not feasible due to computational limitations, testing the algorithm for increasing number of paths does justify the conjecture regarding the bias of the algorithm. As the below chart shows the implied price by LSM algorithm and independent path method and also compares these with the benchmark option value. Both LSM and independent path method converges to the theoretical option price and the errors are well within the variance of the simulation. The yellow shaded area is the variance of the independent path method.

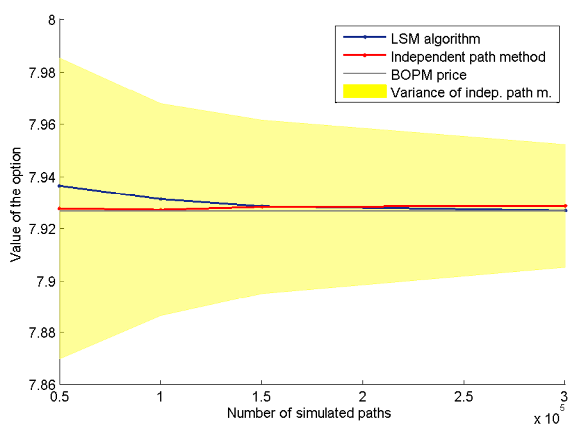

Figure 3.6.4: Convergence of the methods

Monte Carlo simulation is very time and memory consuming, thus in practice one tries to avoid Monte Carlo method for very large number of paths and re-simulation as much as possible. In the following chapter, I am going to introduce an analytic approach to correct this foresight of the algorithm.

Chapter 4

Analytic bias correction

In this chapter I am going to introduce a third alternative approach to generating two independent set of paths and Brute-Force algorithm to determine the unbiased American option price. The analytic bias correction method is an extension of LSM algorithm and similarly to the approach introduced by Fries [2005], it begins with determining a closed form of the actual bias of the model. Using this formula it is possible to analytically clean the inherited foresight bias from the conditional expectation function. Once the investor uses this numerically corrected value for comparison to create the exercise decision, the resulting stopping rule and the implied American option price will be bias-free as well. Let me start with first determining the bias term.

4.1 How biased is it?

LSM algorithm uses least squares regression model at each trading event to approximate the expected value of continuation. For finite number of paths this value is different from the theoretical estimate of continuation. With the aid of the distributional properties of the residuals and the errors it is possible to calculate an unbiased approximation of the current value of the option. The degree of the foresight is different from path to path; consequently, the aim is to derive the bias on each trajectory. In this chapter, according to the previous notations the lower index for S_i, Y_i, $\hat{Y}_i$, y_i, r_i and ϵ_i will denote the value on a particular ω_i trajectory or equivalently the corresponding ith row of the following vectors S, Y, $\hat{Y}$, y, r and ϵ respectively. Specifically for a put option at a given and set time t_k, the objective, as Fries [2005] suggests, is to path-wise determine the following:

$$E\left(max((K - S_i)^+, Y_i)|\mathcal{F}(\mathrm{t_k})\right) \tag{4.1.1}$$

The objective is to find out how much different is the value implied by (4.1.1) from $\hat{Y}_i$, which is the approximation used by LSM algorithm. The relationship between Y_i and $\hat{Y}_i$ is summarized by:

$$Y_i = y_i - \epsilon_i = \hat{Y}_i + r_i - \epsilon_i$$

Since $r_i - \epsilon_i$ will prove to be a crucially important, let me adopt the following notation:

$$e_i = r_i - \epsilon_i$$

Thus, for (4.1.1) it follows:

$$E\left(max((K-S_i)^+, \hat{Y}_i + e_i)|\mathcal{F}(\mathrm{t_k})\right) = E\left(max((K-S_i)^+ - \hat{Y}_i, e_i)|\mathcal{F}(\mathrm{t_k})\right) + \hat{Y}_i \quad (4.1.2)$$

Let me further assume that the distribution of e, the difference of two normally distributed variable, is:

$$e_i \sim N(0, \delta_i^2)$$

S_i and $\hat{Y}_i$ are known constants at t_k since they are $\mathcal{F}(\mathrm{t_k})$-measurable functions, thus (4.1.2) is the expectation of the maximum of a constant and a normally distributed random variable. Hence, after normalizing, the well known properties of the standard normal random variables are applicable:

$$\delta_i \cdot E\left(max(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}, e_0)\right) + \hat{Y}_i$$

where e_0 is now a standard normal random variable, hence using the definition of the expected value, the above further equals to:

$$\delta_i \cdot \int_{-\infty}^{\infty} max(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}, x)\varphi(x)dx + \hat{Y}_i =$$

$$\delta_i \cdot \int_{-\infty}^{\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}} \frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\varphi(x)dx + \delta_i \cdot \int_{\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}}^{\infty} x\varphi(x)dx + \hat{Y}_i$$

Since δ_i along with S_i and $\hat{Y}_i$ are all constants at t_k because once again they are $\mathcal{F}(\mathrm{t_k})$-measurable functions:

$$\left((K-S_i)^+ - \hat{Y}_i\right) \cdot \int_{-\infty}^{\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}} \varphi(x)dx + \delta_i \cdot \int_{\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}}^{\infty} x\varphi(x)dx + \hat{Y}_i$$

Using the well known property of the standard normal distribution, such that $-x\varphi(x) =$

$\varphi'(x)$:

$$\left((K-S_i)^+ - \hat{Y}_i\right) \cdot \int_{-\infty}^{\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}} \varphi(x)dx - \delta_i \cdot \int_{\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}}^{\infty} \varphi'(x)dx + \hat{Y}_i =$$

$$\left((K-S_i)^+ - \hat{Y}_i\right) \cdot \Phi\left(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\right) + \delta_i \cdot \varphi\left(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\right) + \hat{Y}_i$$

Consequently, the difference between the unbiased estimated value of the put option and the approximation of the regression at time t_k is summarized by:

$$b_i \doteq \left((K-S_i)^+ - \hat{Y}_i\right) \cdot \Phi\left(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\right) + \delta_i \cdot \varphi\left(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\right) \qquad (4.1.3)$$

Let me refer to expression (4.1.3) as the bias term b_i, which is the additional information or in other words the look ahead obtained by the algorithm. Furthermore, let b denote the bias vector. As all variables of the bias term are $\mathcal{F}(\mathrm{t_k})$-measurable, at any given t_k it is possible to determine the exact value of the foresight. Thus, the conditional expectation cleaned with the bias term is a computable and fair estimate of the theoretical value of the option at t_k.

4.2 Process of analytic bias correction

At every trading event, the value of continuation is estimated with the least squares method of LSM, thus the analytic foresight removal has to be applied before each decision making, so that the final option value is a fair, unbiased estimate of the real value of the instrument. Consequently, the bias term has to be calculated at each trading event. For a particular t_k time, on a given ω_i trajectory K, S_i and $\hat{Y}_i$ are known, but the distributional properties of e_i and particularly $var(e_i) = \delta_i^2$ are not trivial outputs of the model. However, it is possible to derive them once the distribution of r_i and ϵ_i are known. Using the results of the previous sections, now I introduce a new six-step approach that calculates the bias term and determines a new exercise policy. Since ϵ_i is unobservable, the intuition is to start with the residuals. The analytic bias removal is summarized by the following six-step process:

1. At any given t_k, a realization of $\hat{Y}$ the expected value of continuation is calculated by regressing y, the known, discounted cash flow value at t_{k+1} on a given set of basis functions. Hence, a realization of the r residual random vector is computable such as:

$$r = y - \hat{Y}$$

2. Once the residual vector is known, using the relationship between errors and residuals makes it possible to compute the realization of the error vector. The hat matrix V is calculated as a sub-process of the regression; consequently, it is known, hence based on (3.5.6), the error vector is:

$$\epsilon = (I - V)^{-1} \cdot r$$

3. Since ϵ_i and ϵ_j are independent and identically distributed if $i \neq j$, the sample variance serves as an adequate approximation of σ^2 the variance of ϵ_i for $\forall i \in 1, \ldots, n$. This is calculated as:

$$\sigma^2 \approx \frac{\sum_{i=1}^{n} (\epsilon_i - E(\epsilon_i))^2}{n} = \frac{\sum_{i=1}^{n} (\epsilon_i)^2}{n}$$

4. The bias term b_i is a function of δ_i which is the variance of $e_i = r_i - \epsilon_i$, this is calculated very similarly to the variance of $\hat{Y}_i$:

$$var(e_i) = var(r_i - \epsilon_i) = var(r_i) + var(\epsilon_i) - 2cov(r_i, \epsilon_i)$$

Now the variance of ϵ_i and r_i are known and using the same arguments as for $\hat{Y}_i$, the above is:

$$\sigma^2 + \sigma^2(1 - v_{ii}) - 2E\left(r_i\epsilon_i - E(r_i)E(\epsilon_i)\right) = \sigma^2 + \sigma^2(1 - v_{ii}) - 2E(r_i\epsilon_i) =$$

$$\sigma^2 + \sigma^2(1 - v_{ii}) - 2E\left(\epsilon_i(\epsilon_i - \sum_{j=1}^{n} v_{ij}\epsilon_j)\right) = \sigma^2 + \sigma^2(1 - v_{ii}) + 2E(\epsilon_i^2 - v_{ii}\epsilon_i^2) =$$

$$\sigma^2 + \sigma^2(1 - v_{ii}) - 2\sigma^2(1 - v_{ii}) = v_{ii}\sigma^2$$

Consequently, the variance of e_i is computed as:

$$\delta_i^2 = var(e_i) = v_{ii}\sigma^2$$

5. Now all the building blocks of the bias term are calculated, hence b_i is computable as follows:

$$\left((K - S_i)^+ - \hat{Y}_i\right) \cdot \Phi\left(\frac{(K - S_i)^+ - \hat{Y}_i}{\delta_i}\right) + \delta_i \cdot \varphi\left(\frac{(K - S_i)^+ - \hat{Y}_i}{\delta_i}\right)$$

6. The investor's new decision is based on the newly calculated bias free estimate of continuation. The strategy is to exercise the option if it is in the money and $(K - S_i)^+$ the current exercise value is greater than or equal to $\hat{Y}_i + b_i$ the bias free estimate of continuation and not to exercise every other case. As a result for

in-the-money paths at t_k, the exercise strategy is summarized by:

$$\begin{cases} exercise & if\ (K-S_i)^+ \geq \hat{Y}_i + b_i \\ no\, exercise & else \end{cases}$$

4.3 The bias term

In order to better understand the concept of analytic bias removal, in this chapter I am going to examine the bias term b_i. Since S_i and $\hat{Y}_i$ are known at a given time t_k, it makes sense to look at b_i as a function of the standard deviation δ_i. Using the result for $var(e_i)$ of the previous section, for δ_i it follows that:

$$\delta_i = \sqrt{v_{ii}\sigma^2} = \sqrt{v_{ii}}\sigma$$

This further justifies the approach to view b_i as a function of δ_i, since σ is constant and v_{ii} describes the ith trajectory's properties relative to the rest of the paths, meaning that $b_i(\delta_i)$ does carry information about the look ahead on the particular path. Let me note that if $(K-S_i)^+ < \hat{Y}_i$ at a given time t_k then it is not optimal to exercise regardless the size of the corresponding bias b_i. Consequently, I assume that $\left((K-S_i)^+ - \hat{Y}_i\right)$ is a non-negative constant. The bias term consists of two parts:

1. The first part is $\left((K-S_i)^+ - \hat{Y}_i\right) \cdot \Phi\left(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\right)$. It converges to its supremum of $\left((K-S_i)^+ - \hat{Y}_i\right)$ as δ_i goes to zero and it converges to its infimum of $\frac{\left((K-S_i)^+ - \hat{Y}_i\right)}{2}$ as δ_i goes to infinity.

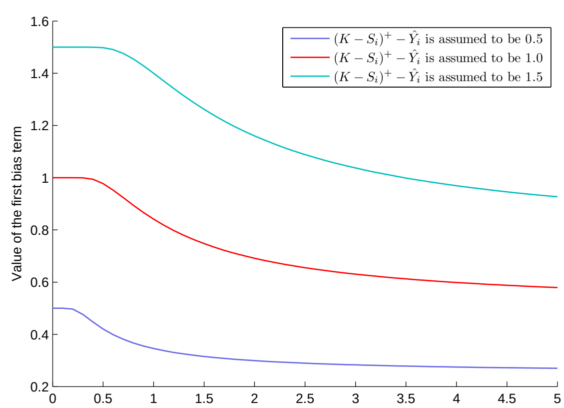

Figure 4.3.1: The first part of the bias term

2. The second part is $\delta_i \cdot \varphi\left(\frac{(K-S_i)^+ - \hat{Y}_i}{\delta_i}\right)$. It converges to its supremum of $+\infty$ as δ_i goes to infinity and it converges to its infimum of 0 as δ_i goes to zero.

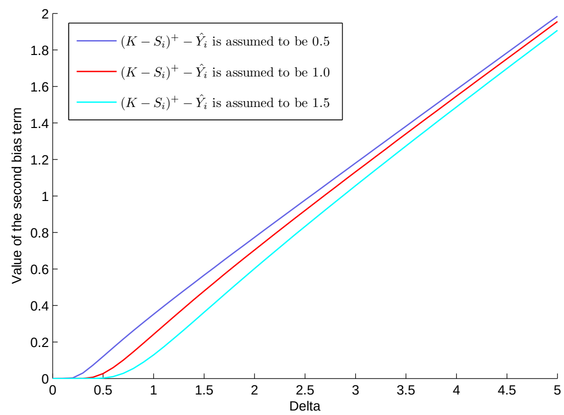

Figure 4.3.2: The second part of the bias term

The bias term itself converges to its supremum of $+\infty$ as δ_i goes to infinity and it converges to its infimum of $\left((K-S_i)^+ - \hat{Y}_i\right)$ as δ_i goes to zero.

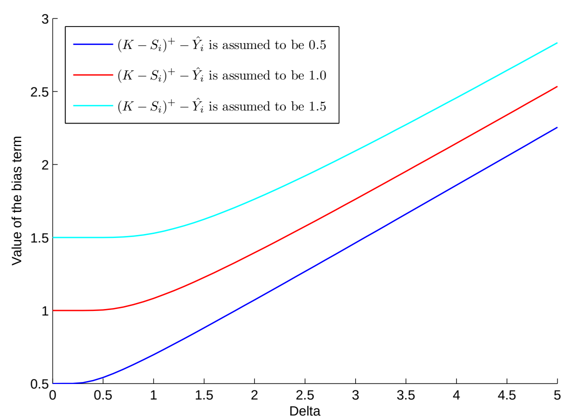

Figure 4.3.3: The bias term as a function of δ

This has two very important implications. First, as δ_i goes to 0, $\hat{Y}_i + b_i$ the true value of the option is $\hat{Y}_i + (K-S_i)^+ - \hat{Y}_i = (K-S_i)^+$. The intuition here is that as δ_i goes to zero the volatility of the stock needs to go to zero as well, meaning that the stochastic nature of the underlying is eliminated, thus it equals to its intrinsic value. Second, the bias term b_i is an increasing function of δ_i, or equivalently the greater the value of v_{ii}, the greater the bias term. Once again, v_{ii} is large if the trajectory is far removed from the bulk of the cases, or in other words if it is biased. This implies that b_i is large if the trajectory is biased and small otherwise. Using bias free estimation of the stopping policy, the condition for exercise is:

$$(K - S_i)^+ \geq \hat{Y}_i + b_i$$

this will less likely to hold on biased, extreme paths because the bias b_i is large since v_{ii} is big as well on these trajectories, hence implying no exercise. Given the backward nature of LSM, this will delay the exercise, meaning that the investor will only exercise at an earlier point in time, when the underlying's spot price is lower and thus the extra premium of the foresight bias is eliminated from the price of the American option.

4.4 Implementation of the new approach

I was using Matlab R2010a for implementing and testing LSM algorithm and the analytic bias correction method. The program starts with identifying the initial values for the simulation along with preallocation of the variables used in the code. For simplicity, I assumed 0% interest rate and chose the volatility parameter σ to be 0.2. The initial value of the stock is 100 and there are 50 trading events in a 1 year time frame:

```
tic %starting a timer
r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100; %strike
numberofpaths = 5000; %number of paths
N = 50; %number of exercise times

%preallocation of the below vectors and matrices
cashflow = zeros(numberofpaths,1);
residual = zeros(numberofpaths,1);
error = zeros(numberofpaths,1);
residualvar = zeros(numberofpaths,1);
errorvar = zeros(numberofpaths,1);
X = zeros(numberofpaths,3);
W = zeros(numberofpaths,N);
```

After initialization, for tractability purposes a seed is specified for the random number generation, this means that the program will generate the same random numbers, hence different algorithms are comparable since they can run on the same random paths. Monte Carlo simulation is done by evaluating the formula of the Geometric Brownian motion s_t

using the randomly generated normal variables stored in the vector *increments*. Each row of the matrix W is one trajectory generated by the simulation:

```
randn('seed',0) %to use seed 0
t=1:N; %N exercise events
time=t/N; %normalising the time frame to 1 year
s_t(:,1)=s_0*ones(numberofpaths,1);
dt=1/N;

for i=1:N;
   increments = randn(numberofpaths,1);
   s_t(:,i+1) = s_t(:,i) .* exp((r-.5 * volatility^2) * dt + volatility ...
       * sqrt(dt) * increments);
   W(:,i+1) = s_t(:,i+1);
end
W(:,1)=[];
```

The algorithm starts with identifying the intrinsic value at maturity on each path, particularly for a put this is done by the following. Let me note that since European options are only exercisable at maturity, this information alone is sufficient to approximate the European option value, this is stored in the variable *european_value*.

```
for i=1:numberofpaths
    if W(i,N)<strike
        cashflow(i) = (strike - W(i,N)) * exp(-r/N);
    else
        cashflow(i) = 0;
    end
end

european_value = sum(cashflow(1:numberofpaths))/numberofpaths;
```

Now that the exercise policy at maturity is known, the implementation of the recursion part of LSM is next. Working backwards for each $j = N-1, \ldots, 1$, in other words at each trading event, variable *index* stores the indices of all the in-the-money paths. The spot price on these trajectories are used to evaluate the regressor function and these figures are stored in the rows of the matrix X. This is then used to calculate the variable *conditionalexp* which is the regression value, the approximation of continuation.

```
for j = N-1:-1:1
    index = find(strike - W(:,j) > 0); %
    X = [ones(size(index)) W(index,j) W(index,j).^2]; %
    B = (X'*X)\X'*cashflow(index); %B = inv(X'*X)*X'*cashflow(index);
    conditionalexp = X * B;
```

Up until this point the algorithm is same for LSM and the new method. The following section will calculate the bias term based on the six steps introduced in Section 4.2. The only difficulty might arise is the calculation of the error term $\epsilon = (I - V)^{-1} \cdot r$, the problem is that $(I - V)^{-1}$ might be badly scaled or close to be singular. Therefore, instead of using the built in inverse function of Matlab, I used the Taylor series to calculate the inverse, specifically it is known that:

$$\frac{1}{1-x} = 1 + x + x^2 + \dots \qquad for\,|x| < 1$$

This also holds true for matrices, numerical tests suggested that even $I + V$ gives accurate results for the inverse of $I - V$.

```
    %bias correction
    V = X / (X'*X) * X'; %V = X*inv(X'*X)*X';
    residual = cashflow(index) - conditionalexp;
    error = (eye(size(index,1)) + V) * residual; %inv(I-V)*r
    errorvar = sum(error.^2) / size(index,1);
    delta = sqrt(diag(V) .* errorvar);
    b1 = (max(strike - W(index,j),0) - conditionalexp) .* ...
        normcdf((max(strike - W(index,j),0) - conditionalexp) ./ delta)
    b2 = delta .* normpdf((max(strike - W(index,j),0) - conditionalexp) ...
        ./ delta);
    bias = b1 + b2;
```

Now that the bias term is known, it is possible to derive the stopping policy by:

```
    for i = 1:size(index,1)
        if conditionalexp(i) + bias(i) <= strike - W(index(i),j)
            cashflow(index(i)) = (strike - W(index(i),j));
        end
    end
end

cashflow = cashflow * exp(-r/N);
toc
american_value = sum(cashflow(1:numberofpaths))/numberofpaths;
```

At last, the value of the American option is calculated with analytic foresight bias correction method and it is stored in the variable *american_value*. In case the conditional expectation is not cleaned from the look ahead bias, the exercise policy is reduced to:

```
    for i = 1:size(index,1)
        if conditionalexp(i) <= strike - W(index(i),j)
            cashflow(index(i)) = (strike - W(index(i),j));
        end
    end
end
```

Hence, the algorithm simplifies to the original method introduced by Longstaff and Schwartz.

4.5 Testing of the new approach

In this section I am going to compare the three introduced methods, namely the original LSM, the independent path algorithm and the analytic bias correction method. The main objective is to see if the price calculated by the analytic bias correction method is in the optimal range. Since LSM algorithm always results in a value greater than the theoretical price and on the other hand the independent path method is known to be suboptimal, it is a natural requirement for the new method to compute option prices in this optimal range. In the chart below, the $y - axis$ represents the value of the option whereas the $x - axis$ shows the number of paths.

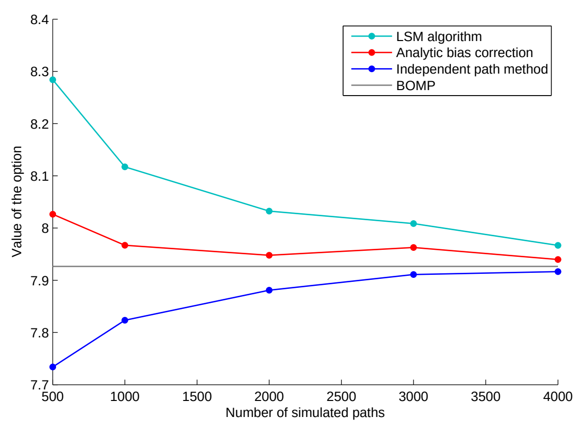

Figure 4.5.1: Analytic bias correction test 1

The value determined by the new method is clearly in the optimal range and it seems to fluctuate a lot less then the other two methods. The next chart shows the same test; however, this time four basis functions are used instead of three.

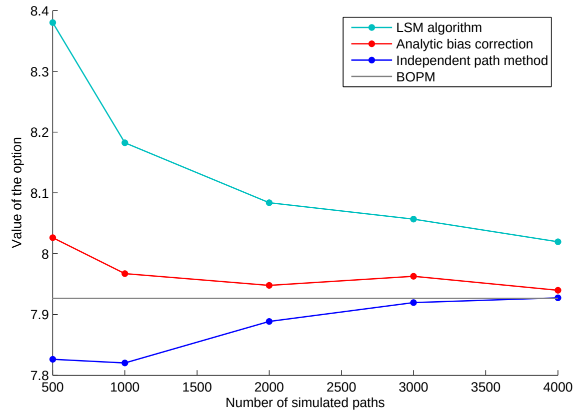

Figure 4.5.2: Analytic bias correction test 2

The last test for the optimal range uses three basis functions; however, this time variance of the underlying is 0.4 instead of 0.2 as in the previous two examples.

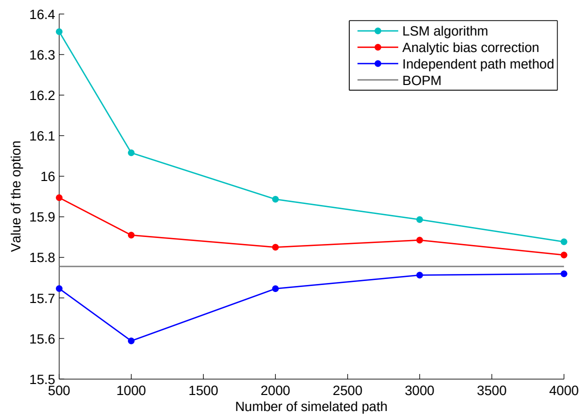

Figure 4.5.3: Analytic bias correction test 3

The option prices generated by the analytic bias correction method are in the optimal range for all the test cases.

4.6 Areas of further research

The new method possesses the most important qualities required from any American style contingent claim pricing model. Since it is fundamentally based on LSM algorithm, it is flexible and easy to implement because it only consists of six simple additional steps. Moreover, it does address the look ahead bias of the original algorithm, hence it derives a more reliable estimate of the real American option value. Even though, the analytic bias removal method uses large matrices, for instance the hat matrix V, it is still computationally tractable; however, simulations with large number of paths might be very time and memory consuming. One possible solution to avoid dealing with large matrices is to divide the in-the-money paths into buckets and then calculate the bias in these various buckets. Since the number of elements of V is the square of the number of in-the-money paths, using the bucket method will substantially reduce the size of the matrices used to calculate the bias. Broadie and Glasserman [2004] propose a new idea called policy fixing, meaning that the investor only considers exercising if the immediate payoff is greater than a threshold. A natural choice for this threshold is the current European option price of the underlying security. Using this approach reduces the range over which bias has to be calculated. In addition, I conjecture that buckets close to strike or alternatively close to the threshold will tend to be very low biased since they are not removed from the bulk of the cases; consequently, calculating the bias term might be avoided for these particular bins altogether. These methods further trim the size of matrix V, resulting the algorithm to be time-wise competitive as well, even for simulations with higher number of paths.

The analytic bias correction method is now proved to be a reliable algorithm for pricing of American call and put options. Further areas of research include testings on derivatives with different and possibly more complex payoffs.

Chapter 5

Conclusion

One of the most popular areas of quantitative finance is the ongoing struggle to determine the optimal exercise strategy used for the pricing of American style contingent claims. Deriving the best exercise policy is the common goal of investors, hedge funds and investment banks to ultimately maximize their profit. Monte Carlo simulation methods are more and more popular in derivative pricing as a result of rapid development of computational efficiency and stochastic calculus. LSM algorithm introduced by Longstaff and Schwartz is a simple yet powerful method for valuing American options. The goal of my thesis was to gain a deep understanding of the algorithm itself along with its strengths and weaknesses and then to address the issue of the embedded foresight bias.

The fundamental incentive was to examine the underlying least squares regression model and furthermore to derive the distributional properties of residuals and theoretical errors. A sound understanding of these principles enabled me to reveal an unbiased approximation of the theoretical conditional expectation value of continuation. Consequently, I introduced a six-step methodology for path-wise calculation of the bias term, a new approach for eliminating the foresight of the original algorithm. The analytic bias removal method fulfills all the natural requirements one might have towards any American option pricing algorithm and upon further testings it proved to be a reliable and accurate new model.

Appendix A

Implementations

Implementation 1:

LSM algorithm:

```
clear all
tic r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100;
numberofpaths = 1000000;
testnum=100;
N = 50; %number of the standard normals ~ number of trading events

european_value = zeros(1,testnum);
american_value = zeros(2,testnum);
cashflow = zeros(numberofpaths,1);

randn('seed',100) %to use seed 100

for g = 1:testnum
   t = 1:N;
   time = t/N;
   W(:,1) = s_0*ones(numberofpaths,1);
   dt = 1/N;
   for i=1:N;
      disp = randn(numberofpaths,1); %increments
      W(:,i+1) = W(:,i) .* exp((r - 0.5 * volatility^2) * dt + ...
          volatility * sqrt(dt) * disp);
   end
   W(:,1) = [];

   for i = 1:numberofpaths
      if W(i,N) < strike
```

```
            cashflow(i) = (strike - W(i,N)) * exp(-r * dt);
        else
            cashflow(i) = 0;
        end
    end

    european_value(g) = sum(cashflow(1:numberofpaths)) / numberofpaths;

    for j = N-1:-1:1
        index = find(strike-W(:,j) > 0);
        X = [ones(size(index)) W(index,j) W(index,j).^2];
        B = inv(X'*X) * X' * cashflow(index);
        conditionalexp = X * B;

        for i=1:size(index,1)
            if conditionalexp(i) <= strike - W(index(i),j)
                cashflow(index(i)) = (strike - W(index(i),j));
            end
        end
        cashflow = cashflow * exp(-r * dt);
    end
american_value(:,g) = [mean(cashflow) std(cashflow)/sqrt(numberofpaths)];
toc
end
american_value' european_value(:);
mean(american_value')
```

Implementation 2:

Independent path method:

```
clear all
tic r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100;
numberofpaths = 5000;
N = 50; %number of the standard normals ~ number of trading events
testnum = 100;

american_value = zeros(2,testnum);
randn('seed',100) %to use seed 100

for g = 1:testnum
    t = 1:N;
    time = t/N;
    cashflow = zeros(numberofpaths,1);
    cashflowS = zeros(numberofpaths,1);
    W = zeros(numberofpaths,N);
    WS = zeros(numberofpaths,N);
```

```
W(:,1) = s_0 * ones(numberofpaths,1);
dt = 1 / N;
for i = 1:N;
   disp = randn(numberofpaths,1); %increments
   W(:,i+1) = W(:,i) .* exp((r - 0.5 * volatility^2) * dt + ...
       volatility * sqrt(dt) * disp);
end
W(:,1) = [];

for i = 1:numberofpaths
   if W(i,N) < strike
      cashflow(i) = (strike - W(i,N)) * exp(-r * dt);
   else
      cashflow(i) = 0;
   end
end

european_value = sum(cashflow(1:numberofpaths))/numberofpaths;

WS(:,1) = s_0 * ones(numberofpaths,1);
dt = 1/N;
for i = 1:N;
   disp = randn(numberofpaths,1); %increments
   WS(:,i+1) = WS(:,i) .* exp((r - 0.5 * volatility^2) * dt + ...
       volatility * sqrt(dt) * disp);
end
WS(:,1) = [];

for i = 1:numberofpaths
   if WS(i,N) < strike
      cashflowS(i) = (strike - WS(i,N)) * exp(-r * dt);
   else
      cashflowS(i) = 0;
   end
end

european_valueS = sum(cashflowS(1:numberofpaths))/numberofpaths;

for j = N-1:-1:1
   index = find(strike - W(:,j) > 0);
   X = [ones(size(index)) W(index,j) W(index,j).^2];
   B = inv(X'*X) * X' * cashflow(index);
   conditionalexp = X * B;

   for i = 1:size(index,1)
      if conditionalexp(i) <= strike - W(index(i),j)
         cashflow(index(i)) = (strike - W(index(i),j));
      end
   end
```

```
        indexS = find(strike - WS(:,j) > 0);
        XS = [ones(size(indexS)) WS(indexS,j) WS(indexS,j).^2];
        conditionalexpS = XS * B;

        for i = 1:size(indexS,1)
           if conditionalexpS(i) <= strike - WS(indexS(i),j)
              cashflowS(indexS(i)) = (strike - WS(indexS(i),j));
           end
        end
        cashflow = cashflow * exp(-r * dt);
        cashflowS = cashflowS * exp(-r * dt);
      end
     american_value(:,g) = [mean(cashflowS) ...
         std(cashflowS)/sqrt(numberofpaths)];
     toc
end

american_value'
european_valueS(:); mean(american_value')
```

Implementation 3:

BOPM model:

```
clear all
r = 0; %interest rate
volatility = 0.2; %volatility
T = 1; %length of the period
s_0 = 100; %price of the underlying at time 0
strike = 100;
N = 50; %number of trading events

dt = T / N;
nudt = (r - 0.5 * volatility^2) * dt;
dx = sqrt(volatility^2 * dt + nudt^2);
pu = 0.5 + 0.5 * nudt / dx;
pd = 0.5 - 0.5 * nudt / dx;

S = s_0 * exp(([0:N] * 2 - N) * dx);
V = max(strike - S,0);
for tt = N:-1:1
   Vup = V([2:tt+1]);
   Vdown = V([1:tt]);
   V = max(max(strike - s_0 * exp(([0:tt-1] * 2 - (tt - 1)) * ...
       dx),0),exp(-r * dt) * (pu * Vup + pd * Vdown));
end

value=V
```

Appendix B

Figures

Figure 3.6.1:

Parameter of the simulation are:

```
r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 110; %the strike is 100
numberofpaths = 11;
N = 50; %number of the standard normals ~ number of trading events
randn('seed',0) %used seed 0 to generate these paths
```

The equation of the exercise boundary is: $f(x) = 2 \cdot \sqrt{N - x} + 110$

Figure 3.6.2:

Shows the average of the means and variance of 100 re-simulations of LSM algorithm with the below parameters:

```
r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100; %the strike is 100
N = 50; %number of the standard normals ~ number of trading events
randn('seed',100) %used seed 100 to generate these paths
```

Figure 3.6.3:

Shows the average of the means and variance of 100 re-simulations of independent path method with the below parameters:

```
r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100; %the strike is 100
N = 50; %number of the standard normals ~ number of trading events
randn('seed',100) %used seed 100 to generate these paths
```

Let me further note that half of the indicated number of paths were used to create the stopping rule and the remaining half were used to calculate the option price.

Figure 3.6.4:

Shows the average of the means and variance of 100 re-simulations of LSM algorithm and independent path method with the below parameters:

```
r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100; %the strike is 100
N = 50; %number of the standard normals ~ number of trading events
randn('seed',100) %used seed 100 to generate these paths
```

Let me further note that for the independent path method half of the indicated number of paths were used to create the stopping rule and the remaining half were used to calculate the option price.

Figure 4.5.1:

Shows the average of the means 100 re-simulations of LSM algorithm, analytic bias correction method and independent path method with the below parameters:

```
r = 0; %interest rate
volatility = .2; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100; %the strike is 100
N = 50; %number of the standard normals ~ number of trading events
randn('seed',100) %used seed 100 to generate these paths
```

Let me further note that for the independent path method half of the indicated number of paths were used to create the stopping rule and the remaining half were used to calculate the option price.

Figure 4.5.2:

Shows the average of the means 100 re-simulations of LSM algorithm, analytic bias correction method and independent path method with the same parameters as figure 9, but 4 basis functions were used instead of 3.

Figure 4.5.3:

Shows the average of the means 100 re-simulations of LSM algorithm, analytic bias correction method and independent path method with the below parameters:

```
r = 0; %interest rate
volatility = .4; %volatility
s_0 = 100; %price of the underlying at time 0
strike = 100; %the strike is 100
N = 50; %number of the standard normals ~ number of trading events
randn('seed',100) %used seed 100 to generate these paths
```

Let me further note that for the independent path method half of the indicated number of paths were used to create the stopping rule and the remaining half were used to calculate the option price.

Bibliography

Takeshi Amemiya. *Advanced Econometrics.* Harvard University Press, 1985.

Alain Bensoussan. On the theory of option pricing. *Acta Applicandae Mathematicae*, 2, 1984.

Fisher Black and Myron Scholes. The pricing of options and corporate liabilities. *The Jornual of Political Economy*, 81, 1973.

Mark Broadie and Paul Glasserman. A stochastic mesh method for pricing high-dimensional american options. *Journal of Computational Finance*, 7, 2004.

Dennis Cook and Sanford Weisberg. Characterizations of an empirical influence function for detecting influential cases in regression. *Technometrics*, 22, 1980.

Dennis Cook and Sanford Weisberg. *Monographs on Statistics and Applied Probability.* 1982.

Christian Fries. The foresight bias in monte-carlo pricing of options with early exercise: Classification, calculation & removal. 2005. URL `www.christian-fries.de`.

John Hull. *Options, Futures and Other Derivatives.* Prentice Hall, 2011.

Kiyoshi Ito. *On stochastic differential equations.* American Mathematical Society, 1951.

Kenneth Judd. Numerical methods in economics. *MIT Press: Cambridge*, 1998.

Ioannis Karatzas. On the pricing of american options. *Applied Mathematics and Optimization*, 17, 1988.

Francis Longstaff and Eduardo Schwartz. Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14, 2001.

Steven Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models.* Springer Finance, 2004.

Halbert White. *Asymptotic Theory for Econometricians.* Academic Press, 1984.