DIPLOMARBEIT

On American Asian Option Duality and Monte Carlo Simulations

zur Erlangung des akademischen Grades

Diplom-Ingenieurin

im Rahmen des Studiums

Masterstudium Finanz- und Versicherungsmathematik

eingereicht von

Dora Kličko

Matrikelnummer 1463032

ausgeführt am Institut für Financial and Actuarial Mathematics
der Fakultät für Mathematik und Geoinformation der Technischen Universität Wien
(in Zusammenarbeit mit Politecnico di Milano)

Betreuer
Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Friedrich Hubalek

Wien, 13.6.2019.

(Unterschrift Verfasserin) (Unterschrift Betreuer)

Abstract

Asian options are fully path dependent options whose payoff depends on the history of the random walk of the asset price via some sort of average. The added feature of early exercise makes the American type particularly complex and hard to price. In such cases where no closed-form solution is available, duality relations become very useful as they transfer knowledge from one option to another. A recent symmetry result by Gounden and O'Hara relates two types of Asian American options that differ according to the role of the average mechanism. This symmetry between the floating and the fixed strike Asian option is of particular interest since much is known about the fixed strike, but comparatively little work has been done for the floating strike case. However, as the proof involves a change of numéraire and a time reversal of the Brownian motion, we suspect this might change the set of stopping times over which the claim holder optimizes. In this paper we develop the necessary tools for pricing Asian American options via simulation and review these duality relations on a set of examples.

Because of the path dependency and early exercise feature, the Asian American option is one of the most representative examples of options that are hard to be priced in terms of speed and/or accuracy. In the past years, Monte Carlo simulation has emerged as the most popular approach in computational finance for determining the prices of such high dimensional options. We employ the now well-known least squares method of Longstaff and Schwartz and use it in conjunction with a dual based method of Rogers and Haugh and Kogan. The result is a bound on the true value of the option.

Contents

Chapter 1

Introduction

1.1 Organization of the paper

In order to get the full benefit from the paper, some knowledge of mathematics and probability theory is implied. Readers familiar with the Black Scholes model and arbitrage theory can skip Chapter 2 without loss of continuity.

The first chapter contains a brief introduction to financial derivatives and Asian options followed by a few intuitive examples. We then take a look at the Black-Scholes-Merton framework, where the underlying asset pays a continuous dividend yield. An introduction to arbitrage pricing concepts and risk-neutral valuation leads to the expression of the Black Scholes price for both European and American option types. The chapter is concluded with a short overview of the duality theory and some interesting symmetry results. Aside from plain vanilla puts and calls, we consider the contributions of Henderson and Wojakovski [16] which relate different types of European arithmetic Asian options. In a recent article, Gounden and O'Hara [14] attempt to extend these results to options of American type, but due to a change of numéraire and time reversal of Brownian motion in conjunction with early exercise this symmetry seems counter intuitive. The purpose of the paper will therefore be to examine the accuracy of the noted duality relations. Since there is no analytical formula for the price of an arithmetic Asian option, we resort to a simulation approach.

Chapter 3 lays out the concepts of the Monte Carlos scheme for European options. We explain in detail how the simulation is carried out and consider improvements like variance reduction techniques. To address the added difficulty of optimal stopping pertaining all American options, in Chapter 4 we introduce a regression-based method of Longstaff and Schwartz [21]. The procedure approximates the optimal exercise policy and yields a lower bound on the price. An upper bound will be derived according to a dual based method developed independently by Rogers [27] and Haugh and Kogan [15].

Having described the necessary concepts, in the final chapter we apply the noted algorithms on some concrete examples. In particular, for a fixed set of parameters we evaluate both options for each duality relation and make a comparison of the results. Note that Monte Carlo only produces an estimate to the price, so the findings are merely indications of the validity of the symmetries.

1.2 Option basics

A vanilla option is a financial instrument that gives the holder the right, but not the obligation, to buy (call option) or sell (put option) an underlying asset at a predetermined price within a given timeframe. The agreed upon price is called the strike price. If the strike price is better than the price in the underlying market at maturity, the option is deemed "in the money" and can be exercised by its owner. An option gains intrinsic value, or moves into the money, as the underlying surpasses the strike price – above the strike for a call and below the strike for a put. Both calls and puts have an expiry date, which puts a time limit on how long the underlying asset has to move. Standard options define the strike price and expiration date at the onset of the contract. However, some options do not define these parameters, leaving them up to the holder and/or to the path the underlying asset takes to decide. The premium is the price paid to own the option. It is based on how close the strike is to the price of the underlying (in the money, out of the money, or at the money), the volatility of the underlying asset, and the time until expiration. Higher volatility and a longer maturity increase the premium. The seller of the option is referred to as its writer. Shorting or writing an option creates an obligation to buy or sell the instrument if the option is exercised by its owner. However, option traders don't need to wait until expiry to close out an options trade, nor do they need to exercise the option. They can take an offsetting position at any time to close the options trade and realize their profit or loss on the option.

Vanilla options have no special or unusual features and are used to hedge the exposure in a particular asset or to speculate on the price movement of a financial instrument. If such an option is not the right fit, exotic options like barrier options, Asian options and digital options are more customizable. They can be combined into complex structures to reduce the net cost or increase leverage. We now turn our attention to Asian options as part of a larger class of path dependent exotic options.

1.3 Asian options

While the payoff of a vanilla option depends only on the relative magnitude of the spot and strike price at maturity, regardless of how it is reached (from above, below or in a zigzag pattern), a path dependent option is designed to take into account the full price path the underlying takes. There are two varieties of path dependent options. One type, called a soft path dependent option, bases its value on a single price event that occurred during the life of the option. It could be the highest or lowest traded price of the underlying asset or it could be a triggering event such as the underlying touching a specific price.

The other, called a hard path dependent option, takes into account the entire trading history of the underlying asset. Some options take the average price, sampled at specific intervals. Option types in this group include Asian options, which allow the buyer to purchase (or sell) the underlying asset at the average price instead of the spot price and are therefore also known as average options.

An Asian option’s payoff is based on an average of underlying asset prices, interest rates, indices, or the like on some specific dates. There are two basic forms

1. Fixed strike Asian option (Average price/rate option)

   An option whose payoff depends on the average spot rate of the underlying asset during at least some part of the life of the option. It is determined as the amount by which the average price of the underlying exceeds a fixed strike price. In other words, if the average rate is higher than the strike price, the option's holder will get the difference. Otherwise, no payoff comes about. The fixed strike is preset and pre agreed by the two parties to the contract.

2. Floating strike Asian option (Average strike option)

   Option whose payoff depends on the difference between the final asset price and the average strike price. More specifically, the strike of an average strike option is not determined at the contract date. Rather it is calculated as an average of the spot rate as observed on a specific series of dates (or fixings) throughout the option's life. If the average price is less favorable than the strike price, the option is in the money and it is cash settled, i.e. without having to physically deliver the underlying asset. The issuer will reimburse the buyer for the difference. However, if the average rate is more favorable, the option will expire worthless.

The average price is usually a geometric or an arithmetic average of the price of the underlying asset at discreet intervals, which need to be specified in the options contract. Due to the averaging mechanism, Asian options have relatively low volatility which makes them less expensive than their standard counterparts. They are often used by traders to solve business problems that ordinary options cannot.

Typical uses include

1. When a business is concerned about the average exchange rate over time.
2. When a single price at a point in time might be subject to manipulation.
3. When the market for the underlying asset is highly volatile.
4. When pricing becomes inefficient due to thinly traded markets (low liquidity markets).

As seen in [17], we offer two simple examples to clarify how the averages are computed and how the option can be used to gain leverage.

Ex. 1 – Average price option

We consider an average price call option using arithmetic averaging and a 30-day period for sampling the data. On November 1st, a trader purchased a 90-day arithmetic call option on stock XYZ with an exercise price of $22, where the averaging is based on the value of the stock after each 30-day period. The stock price after 30, 60, and 90 days was $21.00, $22.00, and $24.00.

The arithmetic average (mean) is (21.00 + 22.00 + 24.00) / 3 = 22.33.

The profit is the average minus the strike price 22.33 - 22 = 0.33 or $33.00 per 100 share contract. As with standard options, if the average price is below the strike price, the loss is limited to the premium paid for the call option.

Ex. 2 – Average rate option

Average rate options (ARO) are often used by companies that receive payments over time that are denominated in a foreign currency. For example, a U.S. manufacturer agrees to import materials from a Chinese company for 12 months and pays the supplier in yuan. The monthly payment is 50,000 yuan. The manufacturer budgets for a particular exchange rate and purchases an ARO that matures in 12 months to hedge against the exchange rate falling below the budgeted level. At the end of each month, the manufacturer purchases 50,000 yuan on the spot market to pay the supplier. Upon maturity of the ARO, the strike price of the ARO is compared to the average rate that the manufacturer has paid for the purchase of 50,000 yuan. If the average is lower than the strike, the manufacturer will exercise the option and the issuer will pay the manufacturer the difference between the strike price and average price.

1.3.1 Option payoff

The Asian options under consideration are that of continuously sampled arithmetic averages. The contracts are initiated at time $t = 0$ and the payoff for exercising a fixed strike Asian option (with strike price $K \in \mathbb{R}^+$) is defined at time $t \in [0, T]$ as

$$C_{fix} = \max\{\varrho(A_t - K), 0\}$$

while the floating-strike option is defined to be the security with the following payoff upon exercise

$$C_{float} = \max\{\varrho(\lambda S_t - A_t), 0\}$$

where $\varrho = 1$ corresponds to a call, and $\varrho = -1$ to a put option. The stock average process A_t is calculated at $t \in [0, T]$, given the price history from the initial time $t = 0$. Formally, we have

$$A_t = \frac{1}{t}\int_0^t S_u du$$

1.4 Dividend paying stock

As we will assume the underlying distributes dividends, we conclude the section with some specific examples where such formulation is commonly used (see [11]).

- Commodities

  In pricing currency options, the relevant underlying variable is an exchange rate. We may think of an exchange rate S (quoted as the number of units of domestic currency per unit of foreign currency) as the price of the foreign currency. A unit of foreign currency earns interest at some risk-free rate, and this interest may be viewed as a dividend stream.

- Exchange rates

  A physical commodity like gold or oil may in some cases behave like an asset that pays negative dividends because of the cost of storing the commodity.

- Futures contracts

  The owner of a long futures contract does not receive dividends; hence this is a disadvantage compared to owning the underlying stock. If the dividend is increased, and the future price would not change, there is an arbitrage possibility. For the sake of simplicity, assume that the stock suddenly starts paying a dividend.

Chapter 2

Principles of Derivatives Pricing

The mathematical theory of derivatives pricing is both elegant and remarkably practical. A proper development of the tools needed even to state precisely its main results requires a book-length treatment. We therefore assume familiarity with at least the basic ideas of mathematical finance and highlight the main concepts, especially those that bear on the applicability of Monte Carlo to the calculation of prices. Three ideas are particularly important.

1. If a derivative security can be perfectly replicated (equivalently, hedged) through trading in other assets, then the price of the derivative security is the cost of the replicating trading strategy.

2. Discounted (or deflated) asset prices are martingales under a probability measure associated with the choice of discount factor (or numéraire). Prices are expectations of discounted payoffs under such a martingale measure.

3. In a complete market, any payoff (satisfying modest regularity conditions) can be synthesized through a trading strategy, and the martingale measure associated with a numéraire is unique.

2.2 The Black Scholes Merton model

In this chapter we discuss the Black Scholes model from the martingale point of view, as introduced in [5].We start with a basic setup, extend it to include dividends and use the risk neutral approach to derive the expression for the option price.

Consider the standard Black Scholes economy with a single risky asset and a money market account or bond. We take as given a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where the information structure is represented by the augmented filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \le t \le T}$ satisfying the "usual hypothesis", i.e. the filtration $(\mathcal{F}_t)_{0 \le t \le T}$ is right continuous and $\mathcal{F}_0$ contains all the $\mathbb{P}$-null sets of $\mathcal{F}$. Following are the assumptions of the original model.

- Market is efficient – people cannot predict the direction of the market or an individual stock
- The market is populated by equally informed traders who do not incur transaction costs (i.e. frictionless market).
- There are no arbitrage opportunities – no way to make a riskless profit
- Interest rates remain constant and known – the rate of return on the riskless asset is the risk-free interest rate
- The stock pays no dividends
- European exercise times are used
- Returns of a stock are log-normally distributed – modelled as a Brownian motion (random walk) with constant drift and volatility

Let S_t^0 denote the unit price of the riskless asset and S_t that of the risky asset at time $t \in [0, T]$. The random vector $\bar{S}_t := (S_t^0, S_t)$ is assumed to be measurable with respect to the σ-algebra $\mathcal{F}_t \subset \mathcal{F}$ and one should think of $\mathcal{F}_t$ as the class of all events observable up to time t.

Definition 2.1 The Black Scholes model consists of two assets with $\mathbb{P}$-dynamics given by

$$dS_t^0 = rS_t^0 dt$$

$$dS_t = \mu S_t dt + \sigma S_t dB_t$$

$r \in \mathbb{R}_+$ is the riskless rate of return

$\sigma \in \mathbb{R}_+$ is the instantaneous volatility of the stock price

μ is the expected rate of return under the actual probability measure $\mathbb{P}$

$B = (B_t)_{t \ge 0}$ is the standard $\mathbb{P}$-Brownian motion

The dynamics of S^0 suggest that an infinitesimal change of the bond value must be equal to the risk-free rate earned over the change in time dt. The percentage changes dS/S in the stock price are stochastic, modelled as the increments of the Brownian motion B. In particular, the filtration $\mathbb{F}^S$ generated by the price process S coincides with the Brownian filtration $\mathbb{F}^B$.

The discount factor β_t plays a prominent role in risk neutral pricing, due to the fact that we will be pricing the derivative asset in terms of the price of the underlying asset. Strictly speaking, determining asset price dynamics under the risk neutral measure will entail specifying dynamics that make the ratios $\beta_t S_t$ martingales. In the given framework, the prices are discounted at the risk-free rate r. The discount factor, determined by the riskless asset, is then defined as

$$\beta_t = \frac{1}{S_t^0} = \exp\left(-\int_0^t r_s ds\right) = e^{-rt}$$

The above model can be extended to include the impact of a dividend yield. When an asset pays a dividend at a constant rate it will be shown that the martingale property also holds, but with S replaced by the sum of S, any dividends paid by S, and any interest earned from investing the dividends at the risk free rate r. This was first done by Merton [23], and the result is also known as the Black-Scholes-Merton formula.

To that end, let D_t denote the cumulative dividend payment of a security over the time interval $[0, t]$. Suppose that the stock price S continuously distributes dividends at some fixed rate $q \in \mathbb{R}$ and that the effective dividend payment is proportional to the level of the stock price. This means that the holder of one unit of asset receives the wealth $dD_t = D(t + dt) - D_t = qS_t dt$ during the infinitesimal time interval $[t, t + dt]$. D grows at a rate

$$\frac{dD_t}{dt} = qS_t + rD_t$$

The first term on the right reflects the influx of new dividends, while the second reflects the interest earned on dividends already accumulated.

By absence of arbitrage, the payment of a dividend must entail a drop in the stock price by the same amount, reducing its value from S_t to $S_t e^{-q(T-t)}$. In order to apply the risk neutral valuation, we proceed by defining an auxiliary process $\hat{S}_t = e^{qt} S_t$ that reinvests all dividends back into the stock. The new process pays no dividends which allows us to price the derivatives for a dividend paying stock in terms of pricing functions for a non-dividend case.

The dynamics of the auxiliary process $\hat{S}_t$ under the real-world measure $\mathbb{P}$ are given as

$$d\hat{S}_t = d(e^{qt}S_t) = qe^{qt}S_t dt + e^{qt}dS_t$$

$$= e^{qt}S_t(qdt + \mu dt + \sigma dB_t)$$

$$= (\mu + q)\hat{S}_t dt + \sigma \hat{S}_t dB_t$$

Our next goal is to study the concept of a martingale measure within the Black-Scholes setup. For purposes of derivative pricing, the most important point is that the risk-neutral measure $\mathbb{Q}$ which makes $\beta_t \hat{S}_t$ a martingale, does not make $\beta_t S_t$ a martingale. This of course affects how we model the dynamics of S under $\mathbb{Q}$.

2.2.1 Risk neutral measure

The price process will on average have a trend and is not necessarily a martingale. However, it can be transformed into one using a change of measure. We assume the existence of a risk-neutral probability measure $\mathbb{Q}$ (equivalent to $\mathbb{P}$) under which the discounted asset prices are martingales, which implies no arbitrage. If $S \sim GBM(\mu, \sigma^2)$, then the drift in $(S_t + D_t)$ equals

$$\mu S_t + qS_t + rD_t$$

For the combined process $(S_t + D_t)$ to be a martingale the above drift coefficient has to be $r(S_t + D_t)$ instead. This is achieved when $\mu + q = r$ i.e. $\mu = r - q$. Indeed, by plugging in the above equation we get

$$\mu S_t + qS_t + rD_t = (r - q + q)S_t + rD_t = rS_t + rD_t = r(S_t + D_t)$$

To find a measure $\mathbb{Q}$ under which these assumptions are met, we first have to calculate the shape ratio (risk premium) θ

$$\frac{d\hat{S}_t}{\hat{S}_t} = (\mu + q)dt + \sigma dB_t$$

$$= rdt + \sigma\left(\frac{\mu + q - r}{\sigma}\right)dt + \sigma dB_t$$

$$= rdt + \sigma d\left(B_t + \int_0^t \frac{\mu + q - r}{\sigma} ds\right)$$

$$= rdt + \sigma dW_t$$

where $W_t = B_t + \theta t$, and $\theta = \frac{\mu - (r-q)}{\sigma}$.

This indicates that the objective and the risk-neutral measure are related through a change of drift in the driving Brownian motion. In particular, the diffusion terms are equal to the real-world ones, which ensures that the coefficients required to describe the dynamics of asset prices under the risk neutral measure can be estimated from data observed under the real-world measure.

To employ the change of measure approach we require the stochastic exponential $Z_t \coloneqq \mathcal{E}(-\theta B_t)$, given as a unique solution to the following stochastic differential equation (SDE)

$$Z_t = \mathcal{E}(-\theta B_0) - \theta \int_0^t \mathcal{E}(-\theta B_s) dB_s = 1 - \theta \int_0^t Z_s dB_s \quad \Rightarrow \quad dZ_t = \theta Z_t dB_t$$

Since Z is a stochastic integral it is a local martingale, and by the Novikov criterion also a real martingale, meaning it satisfies the assumptions of the Radon Nikodym theorem and can act as a density process of a measure change.

The corresponding measure $\mathbb{Q}$ is defined as $\mathbb{Q}(F) = \mathbb{E}_{\mathbb{P}}[Z_T 1_F]$ for some $F \in \mathcal{F}$. More precisely, we have

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_T} = Z_T = exp\left(-\theta B_T - \frac{1}{2}[-\theta B]_T\right) = exp\left(-\theta B_T - \frac{1}{2}\theta^2 T\right), \qquad \mathbb{P} - a.s.$$

Now a consequence of the Girsanov's theorem (Removal of Drift) tells us that for such a measure, the process

$$W_t = B_t + \theta t = B_t + \frac{(\mu + q - r)t}{\sigma} \qquad \Rightarrow \qquad dW_t = dB_t + \frac{(\mu + q - r)}{\sigma} dt$$

is a martingale under $\mathbb{Q}$. Since the quadratic variation process of W is deterministic i.e. $[W]_t = [B]_t = t$, by the Lévy's characterization, W is also a $(\mathbb{F}, \mathbb{Q})$-Brownian motion with values in $\mathbb{R}$.

Using a measure change, we can now revise the initial Black Scholes equation for the stock in terms of the $\mathbb{Q}$-Brownian motion W as follows

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu dt + \sigma dB_t \\ &= \mu dt + \sigma\left(\frac{(\mu + q - r)}{\sigma} dt + dW_t\right) \\ &= (r - q)dt + \sigma dW_t \end{aligned}$$

The coefficient on dt is the mean rate of return and by taking it to be equal the interest rate r minus the dividend yield q we are implicitly describing the risk-neutral dynamics of the stock price. Note that the volatility is the same in a risk-neutral world as it is in the real world.

Definition 2.3 (The Black Scholes Merton model)
The model consists of two assets with $\mathbb{Q}$-dynamics given by

$$dS_t^0 = rS_t^0 dt$$

$$dS_t = (r-q)S_t dt + \sigma S_t dW_t$$

where $W_t = B_t + \theta t$ with $\theta = \frac{\mu-(r-q)}{\sigma}$ is a Brownian motion under the martingale measure $\mathbb{Q}$.

By discounting the equation for the stock price at the risk-free rate, we obtain a process of the form $S_t^* = \beta_t S_t = e^{-rt} S_t$, with the following price dynamics

$$dS_t^* = -qS_t^* dt + \sigma S_t^* dW_t$$

Note that the obtained process is not a martingale under the derived measure $\mathbb{Q}$. Instead, we turn back to the introduced auxiliary process $\hat{S}_t$. Its dynamics under the risk neutral measure are given by

$$d\hat{S}_t = r\hat{S}_t dt + \sigma \hat{S}_t dW_t$$

And in terms of discounted quantities $\hat{S}_t^* = \beta_t \hat{S}_t = e^{-rt}\hat{S}_t$ we have

$$d\hat{S}_t^* = \sigma \hat{S}_t^* dW_t$$

With the solution given by the equation

$$\hat{S}_t^* = \mathcal{E}(\sigma W_t) = \hat{S}_0^* \exp\left(\sigma W_t - \frac{1}{2}\sigma^2 t\right)$$

At this point it is fairly easy to check that the driftless geometric Brownian motion is a martingale. Simply by invoking the properties of the Brownian motion W_t we obtain

$$\mathbb{E}_{\mathbb{Q}}\left(\hat{S}_t^* \middle| \mathcal{F}_s\right) = \mathbb{E}_{\mathbb{Q}}\left(\hat{S}_0^* e^{\sigma W_t - \frac{1}{2}\sigma^2 t} \middle| \mathcal{F}_s\right)$$

$$= \hat{S}_0^* e^{-\frac{1}{2}\sigma^2 t} \mathbb{E}_{\mathbb{Q}}\left(e^{\sigma(W_t - W_s) + \sigma W_s} \middle| \mathcal{F}_s\right)$$

$$= \hat{S}_0^* e^{-\frac{1}{2}\sigma^2 t + \sigma W_s} \mathbb{E}_{\mathbb{Q}}\left(e^{\sigma(W_t - W_s)}\right)$$

$$= \hat{S}_0^* e^{-\frac{1}{2}\sigma^2 t + \sigma W_s} e^{\frac{\sigma^2 (t-s)}{2}}$$

$$= \hat{S}_0^* e^{\sigma W_s - \frac{1}{2}\sigma^2 s}$$

As intended, under the risk neutral measure $\mathbb{Q}$, the discounted and reinvested process $\left(\hat{S}_t^*\right)_{0 \le t \le T}$ is a martingale. We conclude this section by outlining some results with respect to the risk neutral distribution in the Black-Scholes model. The return equation of the stock price

$$\hat{R}_t = \frac{d\hat{S}_t}{\hat{S}_t} = rdt + \sigma dW_t$$

has the following properties under the risk neutral measure $\mathbb{Q}$

- It is also lognormally distributed with mean r and variance σ
- It has the same volatility parameter σ
- The mean parameter μ under the actual probability measure $\mathbb{P}$ is changed so that the expected return on the stock is the riskless interest rate r

Under the actual probability measure the expected stock price discounted at the risk-free rate is a submartingale, meaning that the discounted value is greater than the stock price at time t. This is so because the investors are risk-averse and demand a premium above the risk-free rate for taking on equity risk. By undertaking a change of measure, we shift the actual probability distribution to the left, thus putting more weight on negative outcomes and less weight on positive outcomes while keeping the shape of the distribution unchanged. Because we have shifted the actual probability distribution, the investor's risk premium is zero under the risk-neutral measure. It is instead embedded in the risk-adjusted probabilities.

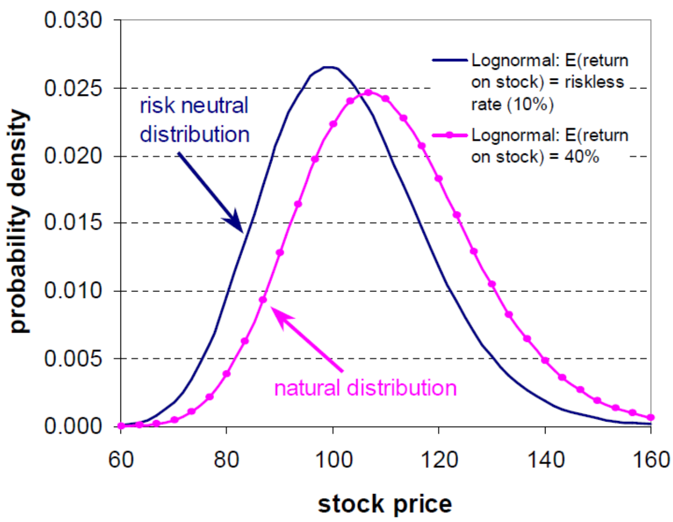

2.3 European option setting

In this chapter we present the general theory of European option pricing in a complete market, as developed by [1], [8], [26]. We start with the standard model and extend it to include dividends.

2.3.1 Strategies and arbitrage

A strategy is defined as a pair $\xi = (\psi, \vartheta)$ where the predictable, S-integrable process ϑ denotes the number of units of the risky asset held in the portfolio and the adapted S^0-integrable process ψ describes the holdings in the savings account. If we fix the portfolio holdings ξ_t over the time interval $[t, t+h]$, then the change in the value of the portfolio is given by

$$\xi_t(\bar{S}_{t+h} - \bar{S}_t) = (\psi, \vartheta)_t \cdot (\bar{S}_{t+h} - \bar{S}_t) = \vartheta_t(S_{t+h} - S_t) + \psi_t(S^0_{t+h} - S^0_t)$$

This suggests that, in a continuous-time case, we may describe the gains from trading over $[0, t]$ through a stochastic integral

$$\int_0^t \xi_u d\bar{S}_u = \int_0^t (\psi, \vartheta)_u \cdot d\bar{S}_u = \int_0^t \vartheta_u dS_u + \int_0^t \psi_u dS^0_u$$

A trading strategy is then self-financing if it satisfies

$$\xi_t\bar{S}_t - \xi_0\bar{S}_0 = (\vartheta_t S_t + \psi_t S^0_t) + (\vartheta_0 S_0 + \psi_0 S^0_0) = \int_0^t \xi_u d\bar{S}_u$$

The left-hand side of the equation is the change in the portfolio value from time 0 to time t, and the right-hand side gives the gains from trading over the interval $[0, t]$. Intuitively, it states that the portfolio is always rearranged in such a way that its present value is preserved. There is no external source of wealth – no gains withdrawn and no funds added, and changes in the portfolio value result solely from asset price fluctuations.

Let V_t denote the wealth process. By rewriting the above equation in the following way

$$\xi_t\bar{S}_t = \xi_0\bar{S}_0 + \int_0^t \xi_u d\bar{S}_u \quad \Rightarrow \quad V_t = V_0 + \int_0^t \xi_u d\bar{S}_u$$

we can interpret it as starting from an initial investment of $V_0 = \xi_0\bar{S}_0$ and following a strategy ξ over $[0, t]$ to ultimately achieve a portfolio gain of $\int_0^t \xi_u d\bar{S}_u$.

In the study of the no-arbitrage property of the model and the derivation of the risk-neutral valuation formula, it is convenient to focus directly on the martingale property of the discounted wealth process.

Definition 2.4 A probability measure $\mathbb{P}^*$ equivalent to $\mathbb{P}$ is called a spot martingale measure if the discounted wealth process $V_t^*(\xi) = \beta_t V_t^*(\xi)$ of any self-financing trading strategy ξ is a local martingale under $\mathbb{P}^*$.

If the stock S does not pay dividends, the notion of a spot martingale measure for the wealth process and the risk neutral measure for the stock price process coincide. This, however, won't be the case if the underlying stock distributes dividends.

2.3.1.1 Dividend case

We have so far assumed that the underlying asset S_t does not pay dividends. This is implied in the definition of a self-financing strategy, where the only gains reflected are those resulting from the change in the price of the underlying asset $\xi_t d\bar{S}_t$. But if each share pays dividends, then the portfolio gains would also include terms of the form $\xi_t dD_t = \xi_t q S_t dt$. Therefore, in the presence of dividends, a simple strategy of holding a single risky asset is no longer self-financing as it entails withdrawal of the dividends from the portfolio. In contrast, a strategy that continuously reinvests all dividends from an asset back into that asset is self-financing in the sense that it involves neither the withdrawal nor addition of funds from the portfolio. These observations further support the introduction of the auxiliary process $\hat{S}$ from the beginning of the chapter. We have

$$\frac{d\hat{S}_t}{\hat{S}_t} = \frac{dS_t + dD_t}{S_t} = (\mu + q)dt + \sigma dB_t$$

The expression on the right is the instantaneous return on the risky asset, including both capital gains and dividends, while the expression on the left is the instantaneous return on the new asset in which all dividends are reinvested. The new process $\hat{S}$ pays no dividends so we may apply the same ideas as those developed in the absence of dividends to the asset.

A strategy (ψ, ϑ) in terms of (S^0, S) may be written as a corresponding strategy $(\psi, \hat{\vartheta})$ in terms of $(S^0, \hat{S})$ where $\hat{\vartheta}_t = e^{-qt}\vartheta_t$, $\hat{S}_t = e^{qt}S_t$. As a result, the trading strategy $\xi = (\psi, \vartheta)$ is said to be self-financing when the wealth process $V(\xi) = \vartheta S + \psi S^0 = \hat{\vartheta}\hat{S} + \psi S^0$ satisfies the following condition

$$\begin{aligned} dV_t &= \hat{\vartheta}_t d\hat{S}_t + \psi_t dS_t^0 \\ &= e^{-qt}\vartheta_t d(e^{qt}S_t) + \psi_t dS_t^0 \end{aligned}$$

$$= e^{-qt}\vartheta_t(e^{qt}dS_t + qe^{qt}S_t dt) + \psi_t dS_t^0$$

$$= \vartheta_t(dS_t + qS_t dt) + \psi_t dS_t^0$$

The essential point is that, while the replicating portfolio is in terms of $\hat{S}_t$, the derivative is in terms of S_t. Note that the fluctuations in the portfolio dynamics are now also impacted by the effective dividend rate.

In order to simplify computations, we want to reduce this situation to one where the savings account is trivial ($S_t^0 = 1$). This corresponds to a "change of currency". By dividing all financial quantities with S_t^0 our new currency is a "riskless bond which pays one Euro at the end of the considered time period". We will represent prices using the correct number of these zero bonds. Only at the very end of the discussion will we translate the results back to the usual currency.

We start by dividing all financial quantities with S_t^0 and apply the Itô integration by parts formula which shows that the strategy (ψ_t, ϑ_t) for (S_t^0, S_t) is self-financing if and only if (ψ_t, ϑ_t) is self-financing for $(1, S_t/S_t^0)$. As $dS_t^0 = 0$ for $S_t^0 = 1$, the self-financing constraint then determines the wealth process V and the bank account holdings uniquely.

More precisely, since the stock price process S is modelled as a $GBM(r-q, \sigma^2)$ under the risk neutral measure $\mathbb{Q}$, we can calculate the dynamics of the discounted wealth process $V_t^* = \beta_t V_t$ in the following way

$$dV_t^* = d(e^{-rt}V_t) = -re^{-rt}V_t dt + e^{-rt}dV_t$$

$$= e^{-rt}(-rV_t dt + \vartheta_t(dS_t + qS_t dt) + \psi_t dS_t^0)$$

$$= e^{-rt}\big(-rV_t dt + \vartheta_t\big((r-q)S_t dt + \sigma S_t dW_t + qS_t dt\big) + \psi_t r S_t^0 dt\big)$$

$$= e^{-rt}(-rV_t dt + \vartheta_t(rS_t dt + \sigma S_t dW_t) + \psi_t r S_t^0 dt)$$

$$= e^{-rt}(-rV_t dt + r(\vartheta_t S_t + \psi_t S_t^0)dt + \vartheta_t \sigma S_t dW_t)$$

$$= e^{-rt}\vartheta_t \sigma S_t dW_t$$

$$= \vartheta_t \sigma S_t^* dW_t$$

$$= \hat{\vartheta}_t \sigma \hat{S}_t^* dW_t$$

$$= \hat{\vartheta}_t d\hat{S}_t^*$$

It is now clear that the spot martingale measure $\mathbb{P}^*$ for V^* is also the risk neutral measure for the reinvested process $\hat{S}^*$, and not for the discounted stock price process S^* (unless $q = 0$).

Lemma 2.5 For any self-financing trading strategy $\xi = (\psi, \vartheta)$, the dynamics of the discounted wealth process $V^*(\xi)$ under $\mathbb{P}^* = \mathbb{Q}$ are given by

$$dV_t^* = \vartheta_t \sigma S_t^* dW_t = \hat{\vartheta}_t \sigma \hat{S}_t^* dW_t = \hat{\vartheta}_t d\hat{S}_t^*$$

In integral form we have

$$V_t^* = V_0^* + \int_0^t \hat{\vartheta}_t d\hat{S}_t^*$$

Hence for a self-financing strategy we do not need to specify ψ or the savings account. The relation is completely determined with initial capital V_0^* and strategy ϑ. The S-integrable process ϑ can be interpreted as the number of shares held at time t, while the stochastic integral process $\int \hat{\vartheta}_t d\hat{S}_t^*$ presents the outcome of trading in $\hat{S}$, often called the gains process.

In many cases, the associated value process V is a $\mathbb{Q}$-local martingale for every martingale measure $\mathbb{Q}$ such that V is $\mathbb{Q}$-integrable. This holds for instance when the strategy ϑ is locally bounded or when V is a special semi-martingale. However, it is not true in full generality and merely ensuring that V is a $\mathbb{Q}$-local martingale does not necessarily exclude arbitrage phenomena. Take for example the following situation

Example 2.6 Suicide strategy [26]

Let the price process of a risky asset evolve according to a stochastic exponential $Z = \exp\left(B - \frac{1}{2}t\right)$ defined on the whole $\mathbb{R}^+$ where B denotes the standard Brownian motion. It can be shown that investing in an underlying with such price process, one starts with initial capital of one and ends up with no money at time infinity. This happens due to the fact that Z is not a martingale on the closed interval $[0, \infty]$. The same setting can be modified to generate an arbitrage opportunity, where we invest in a stock whose price process follows a Brownian motion via the self-financing strategy and generates a sure profit of one (albeit after an infinite amount of time), and yet requires no initial capital.

In order to exclude pathological behavior like the one described, it is necessary to restrict ourselves to considering only admissible strategies. We find it convenient to introduce the concept of admissibility of a trading strategy in terms of a martingale measure.

Definition 2.7 A trading strategy ϑ is called $\mathbb{Q}$-admissible if the associated discounted wealth process $V_t^* = e^{-rt} V_t$ follows a martingale under measure $\mathbb{Q} \in \mathcal{M}^e$.

Note that this applies if and only if the (normalized) gains process $\int \vartheta dS_t$ is a $\mathbb{Q}$-martingale.

2.3.2 Hedging in a complete market

In the Black Scholes setting, the martingale measure for the discounted stock price is unique and it is known explicitly. By the second fundamental theorem of asset pricing this is equivalent to a market being complete. We begin this section with a brief outline of the logic behind this result then turn our attention to the pricing problem for European contingent claims in a complete market.

The existence of a self-financing trading strategy that replicates a derivative security determines the price of the security. Under the assumptions of risk neutral dynamics, this argument leads to a representation of the derivative price as an expectation in terms of random objects we can simulate. As noted before, it is possible to construct an arbitrage-free market in which a risk-free bond and a dividend-paying stock are primary securities. Assuming that this is done, the valuation of stock-dependent contingent claims is now standard.

A European claim is an $\mathcal{F}_T$-measurable random variable $h(S_T)$ for some Borel measurable payoff function $h: [0, \infty) \to \mathbb{R}$. It can be replicated by trading in the underlying assets precisely if its payoff at expiration $h(S_T)$ coincides with the value V_T of some attainable (replicable) price process at time T. Clearly, a contingent claim H is attainable if the corresponding discounted claim $H^* = h(S_T)/S_T^0$ is attainable. To arrive at an expression for the price of a claim, it will be convenient to work with discounted quantities together with the trivial savings account (for which $dS^0 = 0$). With a slight abuse of notation, we denote $\hat{S}/S^0$ by S, $\hat{\vartheta}$ by ϑ, V^* by V and $H^*$by H until further notice.

Definition 2.8 A discounted European claim H is attainable if there exists a constant c and an admissible replicating strategy ϑ such that

$$H = c + \int_0^T \vartheta dS_t$$

The standing assumption of no arbitrage allows us to fix some equivalent martingale measure $\mathbb{Q}$ and for each bounded claim consider the associated $\mathbb{Q}$-martingale M given by

$$M_t = \mathbb{E}_{\mathbb{Q}}(H|F_t)\,, \qquad t \le T$$

To identify a complete market, we need to be able to tell whether it is possible to represent M as a stochastic integral with respect to the underlying process S. By the Kunita-Watanabe theorem we already know that M allows for a unique decomposition into a stochastic integral with respect to S and a strongly orthogonal local martingale L (corresponding to an attainable and non-attainable part). Nevertheless, to specify attainable claims we will need the following, stronger property to hold.

Let $\mathcal{M}^2$ denote the space of L^2-bounded martingales B where $\mathbb{E}([B]_T) < \infty$ and $\mathbb{F}^B$ the filtration generated by a process B in $\mathcal{M}^2$. The space $L^2(B)$ consists of all predictable processes ϑ such that $\mathbb{E}\left(\int_0^T \vartheta_t^2 d[B]_t\right) < \infty$.

Definition 2.9 For any $M \in \mathcal{M}^2(\Omega, \mathcal{F}^B, \mathbb{Q})$ there exists a unique $\mathbb{F}$-predictable process ϑ in $L^2(B)$ such that the predictable representation property (PRP) holds

$$M = M_0 + \int_0^T \vartheta dB$$

The condition $\vartheta \in L^2(B)$ implies that the Itô integral $\int_0^T \vartheta dB$ is a square-integrable $\mathbb{F}$-martingale and thus we also have for every $t \in [0, T]$ that

$$\mathbb{E}_{\mathbb{Q}}(M|\mathcal{F}_t^B) = M_0 + \int_0^T \vartheta dB$$

By the second fundamental theorem of asset pricing the PRP is equivalent to the existence of a unique equivalent martingale measure (EMM) i.e. to the market being complete and every claim attainable.

Example 2.10 Brownian motion

Brownian motion B satisfies the PRP with respect to its natural augmented filtration $\mathbb{F}^B$, implying thereby that all martingales with respect to the Brownian filtration are continuous. As a result, markets containing a single price process S driven by a one-dimensional Brownian motion B with $\mathbb{F}^S = \mathbb{F}^B$, are complete. This in particular applies to the considered Black Scholes economy.

Let us now consider a bounded claim H in a complete market. As the market is complete, H is attainable and there exist accordingly a constant c and an admissible strategy ϑ such that

$$H = c + \int_0^T \vartheta dS_t$$

The constant $c \in \mathbb{R}$ is the initial capital required to drive the self-financing replicating strategy ϑ (by financing the investment in the underlying asset). As the sum of initial capital and the outcome from trading matches the payoff of the claim H, it is reasonable to call c the fair price of the claim. Following arbitrage arguments, we can now show how this price can be calculated.

By the second fundamental theorem of asset pricing, there exists a unique EMM $\mathbb{Q}$. Given $\mathbb{Q}$ we can associate with claim H a value process V, obtained by taking the $\mathbb{Q}$-conditional expectation (thus converting H into a process!)

$$V_t = \mathbb{E}_{\mathbb{Q}}(H|\mathcal{F}_t)$$

The constructed process is a $\mathbb{Q}$-martingale which can easily be verified using the tower property for conditional expectations

$$\mathbb{E}_{\mathbb{Q}}(V_t|\mathcal{F}_s) = \mathbb{E}_{\mathbb{Q}}\big(\mathbb{E}_{\mathbb{Q}}(H|\mathcal{F}_t)\big|\mathcal{F}_s\big) = \mathbb{E}_{\mathbb{Q}}(H|\mathcal{F}_s) = V_s\,, \qquad a.s.$$

At the beginning of the chapter we show that the process $W_t = B_t + \theta t$ is a $\mathbb{Q}$-Brownian motion with Brownian filtration $\mathbb{F}^{\mathrm{W}}$. Under $\mathbb{Q}$, the price process has the PRP, meaning every $(\mathbb{Q}, \mathbb{F}^W)$-martingale can be written as a stochastic integral with respect to W. In particular, the $\mathbb{Q}$-dynamics of S are given by $dS_t = \sigma S_t dW_t$, justifying thereby the representation

$$V_t = V_0 + \int_0^t \vartheta_t^H dS_t$$

for some predictable strategy ϑ^H. Note that in order for V to truly represent the value of the derivative, at maturity $t = T$ we must have

$$H = V_T = \mathbb{E}_{\mathbb{Q}}(H|\mathcal{F}_T) = V_0 + \int_0^T \vartheta^H dS_t$$

By taking the $\mathbb{Q}$-expectation of both sides of the equation and using the fact that $\int \vartheta^H dS_t$ is a $\mathbb{Q}$-martingale (by the admissibility of ϑ^H) with $\mathbb{Q}$-expectation equal to zero, we obtain the fair price (Black Scholes price) of the claim H

$$\mathbb{E}_{\mathbb{Q}}(H) = \mathbb{E}_{\mathbb{Q}}(V_0) + \mathbb{E}_{\mathbb{Q}}\left(\int_0^T \vartheta_t^H (dS + qS_t dt)\middle|\,\mathcal{F}_0\right) = V_0 + \int_0^0 \vartheta^H (dS + qS_t dt) = V_0$$

In a complete market all claims are redundant, in the sense that they can be perfectly replicated by the outcome from trading with a self-financing admissible strategy in the underlying asset. The integrand ϑ^H used to replicate the payoff of a claim in question is called the hedging strategy and the initial capital needed to drive such a replicating strategy then corresponds to the fair price of the claim. It can be calculated by taking the expectation of the payoff H under the unique martingale measure $\mathbb{Q}$, commonly called the risk neutral measure

$$V_0 = \mathbb{E}_{\mathbb{Q}}[H]$$

It would be a severe mistake to take the expectation under the statistical (sometimes referred to as the "real-world") measure $\mathbb{P}$ unless the drift μ is zero. Observe however, that we do not assume that the agents in our model are risk neutral. The method proposes that even though most investors are risk-averse, we can price derivatives as if they all were risk-neutral. In particular, the agents are allowed to have any attitude to risk whatsoever, as long as they all prefer a larger amount of (certain) money to a lesser amount. This is reflected in the fact that the pricing equation does not contain the local mean rate of return, indicating that we can price derivatives without knowing anything about the investor's risk preference.

Remark 2.11 In what sense is it justified to call this expression the arbitrage price of a claim? We may sell the derivative security for V_0 at time 0, use the proceeds to implement a self-financing trading strategy ϑ, and deliver the promised payoff of $V_T = H$ at time T with no risk. If anyone were willing to pay more than V_0, we could sell the derivative and be guaranteed a riskless profit from a net investment of zero. If anyone were willing to sell the derivative for less than V_0, we could buy it, implement the strategy $-\vartheta$ and again be ensured a riskless profit without investment. Thus, V_0 is the only price that rules out riskless profits from zero net investment.

2.3.3 Summary

The Black Scholes model with a continuous dividend yield consists of one risky asset and a riskless bond whose dynamics under the unique risk neutral measure $\mathbb{Q}$ are given by

$$dS_t = (r - q)dt + \sigma dW_t$$

$$dS_t^0 = rS_t^0 dt$$

where $W_t = B_t + \theta t$ is a standard $\mathbb{Q}$-Brownian motion with shape ratio $\theta = \frac{\mu-(r-q)}{\sigma}$. The risk neutral measure $\mathbb{Q}$ is given by the Radon-Nikodym derivative

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_T} = e^{-\theta B_T - \frac{1}{2}\theta^2 T}$$

We accommodate dividends by introducing an auxiliary process $\hat{S}$ with dividends reinvested, defined through the requirement

$$\frac{d\hat{S}_t}{\hat{S}_t} = \frac{dS_t}{S_t} + \frac{dD_t}{dS_t}$$

Under $\mathbb{Q}$ the dynamics of $\hat{S}$ are given by

$$\frac{d\hat{S}_t}{\hat{S}_t} = rdt + \sigma dW_t$$

Implying that the discounted process $\hat{S}_t^* = \beta_t \hat{S}_t$ is a $\mathbb{Q}$-martingale.

The reinvested process $\hat{S}$ pays no dividends, so we can calculate the value of the option by applying the same ideas as in the absence of dividends to these assets.

With every trading strategy one can associate a value process. For a self-financing trading strategy $\xi = (\psi, \vartheta)$ that process satisfies the following condition under the martingale measure $\mathbb{Q}$

$$V_t^* = e^{-rt}V_t = V_0^* + \int_0^t \hat{\vartheta}_u \, d\hat{S}_u^*$$

In fact, $\mathbb{Q}$ is by definition the measure that makes the discounted wealth process of any self-financing strategy a martingale.

If the market model is arbitrage free, it is justified to call the associated wealth process of an attainable claim $H = h(S_T)$ the arbitrage price process of H. How do we associate a process to a claim (payoff)?

First note that the considered Black Scholes market model is complete. In a complete model every claim H is attainable and can thus be replicated by an admissible trading strategy. Associated with that strategy is a (discounted) wealth process that is by definition a $\mathbb{Q}$-martingale and as such, satisfies the self-financing condition. In particular, if the replicating portfolio really hedges the claim, the associated wealth process must coincide with the payoff of the claim at maturity i.e. we have $V_T = H$. Finally, the arbitrage price process of H is given by that same wealth process associated to its hedging strategy. We have

$$V_t^* = \mathbb{E}_{\mathbb{Q}}(H^*|\mathcal{F}_t) = V_0^* + \int_0^t \hat{\vartheta}_u \, d\hat{S}_u^*$$

Ultimately, we want to convert the above pricing formula back to the case where the currency is euros and not bonds. When the currency consists of euros, the time of the pay-off becomes relevant. The above expression translates to

$$V_t = e^{rt} V_t^* = e^{rt} \mathbb{E}_{\mathbb{Q}}[H^*|\mathcal{F}_t] = e^{rt} \mathbb{E}_{\mathbb{Q}}[e^{-rT} H|\mathcal{F}_t] = e^{-(T-t)r} \mathbb{E}_{\mathbb{Q}}[H|\mathcal{F}_t]$$

Since the price of the derivative equals its value at inception, i.e. the initial investment needed to finance the replication of the price process (purchase of the hedging strategy), we have

$$V_0 = e^{-rT} \mathbb{E}_{\mathbb{Q}}[H]$$

In such a way, the present value of a derivative (some future stochastic payout) in a risk neutral world equals the expected value of all future (net) payments discounted to present value using the short rate of interest. This is exactly the same result that we would obtain if we assumed that the world was risk neutral.

2.4 American option setting

In this chapter we lay out the theory on American options as it was developed in [20]. Basic results on continuous time martingales can be found in [28] and for applications to American option pricing we refer to [4], [10] and [22].

According to [3], the risk-neutral approach described in the previous chapter is also valid for American options given that the underlying market is complete. However, since the option holder has to decide on an optimal exercise strategy, American contracts are more complicated to analyze than their European counterparts and require the use of more sophisticated tools. To that end, we begin with a recap of some essential results from stochastic analysis and continue to extend the risk neutral pricing method to the American case.

2.4.1 Optimal stopping vs. hedging

An American feature of the option allows the holder of the option to exercise it at a pre-determined set of dates. To put it more formally, let us fix a final exercise date T and a contract function g. The European version of this contract will, as usual, pay the amount $g(S_T)$ at maturity. If the contract on the other hand is of the American type, the holder will obtain the amount $g(S_t)$ if he/she chooses to exercise the contract at time t. The situation is further complicated by the fact that the exercise time does not have to be chosen a priori (i.e. at $t = 0$). Instead, it is chosen on the basis of the information generated by the stock price process ($\mathbb{F}^S$), and thus the holder will in fact choose a random exercise time τ. The decision on an exercise strategy τ still has to depend solely on the information available up to time τ. The mathematical formulation of this property is in terms of the so-called "stopping times".

The holder of the American contract thus has to decide on an optimal exercise strategy contingent on the information available up to that time. Mathematically this means that we have to solve the following "optimal stopping problem"

$$\sup_{\tau} \mathbb{E}_{\mathbb{Q}}\left(e^{-r\tau} g(S_\tau)\right)$$

Problems of this kind are hard to solve, and analytically they lead to the so called "free boundary value problems". At each time t we have two options: hold on to the option or exercise. With only two options given, there must exist a value of S that marks the boundary between these two regions. This price, which we can denote by S_τ, is called the optimal exercise price. We do not know S_τ a priori and thus we do not know where to apply the boundary conditions, hence the name "free-boundary problem". It will be shown that, in a complete market model, any exercise strategy that maximizes the expected payoff under the unique EMM is in fact optimal.

We now take the point of view of the seller, whose aim is to hedge against all possible claims of the buyer. As we do not know which stopping time τ will be used, we need to prepare for the worst possible case, and charge the maximum value (maximized over all possible stopping strategies)

$$V_0 = \sup_{\tau} \mathbb{E}_{\mathbb{Q}}(e^{-rT} g(S_\tau))$$

If the holder of the option does not exercise optimally, the issuer's hedge will produce a surplus by date T. Even in the simplest Black-Scholes model, the hedging and pricing of American options needs very refined mathematical tools. Following [22] we will solve this problem under the assumption of market completeness, using the notion of Snell envelope in continuous time and the Doob-Meyer decomposition.

2.4.2 Snell envelope and Doob-Meyer decomposition

We assume a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a Brownian filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$. The model is studied in a bounded time interval from 0 to T. A $[0, T]$-valued random variable τ is called an $\mathbb{F}$-stopping time if it is adapted to the filtration $\mathbb{F}$, that is $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \in [0, T]$. Since in the Black Scholes model we have $\mathbb{F} = \mathbb{F}^{\mathrm{B}} = \mathbb{F}^{\mathrm{W}} = \mathbb{F}^{\mathrm{S}} = \mathbb{F}^{\hat{\mathrm{S}}}$, any stopping time of the filtration $\mathbb{F}$ is also a stopping time on the filtration generated by the stock price process. This reflects on the assumption that the decision to exercise an American claim at time t is based on the observations of the stock price fluctuations up to that time, but not after this date.

We denote by $\mathcal{T}$ the set of all stopping times with respect to filtration $\mathbb{F}$ and introduce the following subsets of $\mathcal{T}$

$$\mathcal{T}_{[\mathrm{t,T}]} = \{\tau \in \mathcal{T} : \mathbb{P}(\tau \in [t, T]) = 1\}, \qquad 0 \leq t \leq T < \infty$$

$$\mathcal{T}_{[\mathrm{t},\infty]} = \{\tau \in \mathcal{T} : \mathbb{P}(\tau \in [t, +\infty]) = 1\}, \qquad t \geq 0$$

Theorem 2.12 (Doob's optional stopping theorem)
Let $(M_t)_{t\geq 0}$ be a right continuous process and σ, τ bounded stopping times such that $\sigma \leq \tau$. If

1. M is a martingale, then the random variables M_σ and M_τ are integrable and we have

$$\mathbb{E}(M_\tau | \mathcal{F}_\sigma) = M_\sigma, \qquad a.s.$$

2. M is a nonnegative supermartingale, then the limit $M_\infty = \lim_{t\to\infty} M_t$ exists a.s. and

$$\mathbb{E}(M_\tau | \mathcal{F}_\sigma) \leq M_\sigma, \qquad a.s.$$

The above theorem generalizes the defining (in)equality for (super)martingales from deterministic times to stopping times that take finitely many values. We will need the following terminology

Definition 2.13 An adapted right continuous process $(X_t)_{t\geq 0}$ is said to be

- Of class D if the family $(X_\tau)_{\tau\in\mathcal{T}_{0,\infty}}$ is uniformly integrable
- Regular if, for every $\tau \in \mathcal{T}_{0,\infty}$, X_τ is integrable and, for every nondecreasing sequence $(\tau_n)_{n\in\mathbb{N}}$ of stopping times with $\tau = lim_{n\to\infty}\tau_n$, we have $lim_{n\to\infty}\mathbb{E}\big(X_{\tau_n}\big) = \mathbb{E}(X_\tau)$

Throughout the chapter we consider an adapted, right continuous process $Z = (Z_t)_{t\geq 0}$ satisfying

$$Z_t \geq 0\ , \qquad \mathbb{E}\left(\sup_{t\geq 0} Z_t\right) < \infty\ , \qquad \forall t \geq 0$$

The process Z is obviously of class D and in the interest of defining the Snell envelope of Z we rely on the following theorem

Theorem 2.14 For $t \geq 0$, set

$$U_t = \operatorname*{ess\,sup}_{\tau\in\mathcal{T}_{t,\infty}} \mathbb{E}(Z_\tau|\mathcal{F}_t)$$

1. The process $(U_t)_{t\geq 0}$ is a supermartingale
2. For every $t \geq 0, \mathbb{E}(U_t) = \sup_{\tau\in\mathcal{T}_{t,\infty}} \mathbb{E}(Z_\tau)$
3. U admits a right continuous modification

The right continuous modification of U is called the Snell Envelope of Z and will still be denoted by U. The Snell Envelope is the smallest right continuous supermartingale dominating Z, and we can easily verify the following result

$$\lim_{t\to\infty} U_t = \limsup_{t\to\infty} Z_t\ , \qquad a.s.$$

Note that U is of class D, and as such satisfies the following theorem

Theorem 2.15 (The Doob-Meyer decomposition of a supermartingale)
Let $U = (U_t)_{t\geq 0}$ be a right continuous supermartingale of class D. There exists a martingale $(M_t)_{t\geq 0}$ and a nondecreasing right continuous predictable process $A = (A_t)_{t\geq 0}$ with $A_0 = 0$, which are uniformly integrable, unique up to indistinguishability and such that

$$U_t = U_0 + M_t - A_t\ , \qquad t \geq 0$$

Moreover, if U is a regular process, the process A has continuous paths with probability one.

We say that a stopping time $\tau^* \in \mathcal{T}$ is optimal if it maximizes $\mathbb{E}(Z_\tau)$ i.e. if

$$\mathbb{E}(Z_{\tau^*}) = \sup_{\tau \in \mathcal{T}} \mathbb{E}(Z_\tau)$$

and propose the following characterization

Theorem 2.16 A stopping time $\tau^* \in \mathcal{T}_{0,\infty}$ is optimal if and only if the following conditions hold

1. $U_{\tau^*} = Z_{\tau^*}\,, \qquad a.s.$
2. The stopped process U^{τ^*}, defined by $U_{\tau^*}^t = U_{\tau^* \wedge t}\,,\ 0 \leq t \leq T$, is a martingale.

Assume now that the process Z is regular. If Z is regular, so is its Snell envelope. In this case, the existence of an optimal stopping time is equivalent to $\mathbb{P}(\tau_0 < \infty) = 1$ where τ_0 denotes the smallest optimal stopping time defined as

$$\tau_0 = \inf\{t \geq 0 | U_t = Z_t\}$$

To conclude, as we are dealing with an optimal stopping problem of finite horizon, we introduce the Snell envelope with horizon T, defined as

$$U_t^T = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}(Z_\tau | \mathcal{F}_t)$$

We then have $U_t^T = Z_T\ a.s.$ and, if the process $(Z_t)_{0 \leq t \leq T}$ is nonnegative, right continuous, regular and satisfies $\mathbb{E}\left(\sup_{0 \leq t \leq T} Z_t\right) < \infty$, the stopping time τ_0 is in $\mathcal{T}_{0,T}$ since $Z_T = U_t^T$ and it is minimal among all optimal stopping times.

2.4.3 Portfolio and consumption

In the context of arbitrage pricing of American claims, we assume that an individual may withdraw funds to finance his consumption needs. The consumption process $C = (C_t)_{0 \leq t \leq T}$ is nondecreasing, adapted and continuous with $C_0 = 0$. By C_t we denote the cumulative amount of funds that are withdrawn and consumed by an investor up to time t (in the sense that the wealth is dynamically diminished according to the process C). We can think of the pair (ξ, C) as the trading and consumption strategy in (S^0, S).

We jump straight away to the extended Black Scholes economy with included dividends. Just like in the European case, we can write a strategy (ψ, ϑ) in terms of (S^0, S) as a corresponding strategy $(\psi, \hat{\vartheta})$ in terms of $(S^0, \hat{S})$ where $\hat{\vartheta}_t = e^{-qt}\vartheta_t$, $\hat{S}_t = e^{qt}S_t$. The self-financing restriction reads as follows

Definition 2.17 A trading and consumption strategy (ξ, C) in (S^0, S) is self-financing on $[0, T]$ if the wealth process $V(\xi, C)$ given by the formula

$$V_t(\xi, C) = \hat{\vartheta}_t \hat{S}_t + \psi_t S_t^0 = \vartheta_t S_t + \psi_t S_t^0, \qquad t \in [0, T]$$

satisfies for every $t \in [0, T]$ the following condition

$$dV_t(\xi, C) = \vartheta_t(dS_t + qS_t dt) + \psi_t dS_t^0 - dC_t$$

In integral form we have

$$V_t(\xi, C) = V_0(\xi, C) + \int_0^t \vartheta_t(dS_t + qS_t dt) + \int_0^t \psi_t dS_t^0 - C_t$$

The above equality must be interpreted as the indistinguishability of the two processes, which implies that V is continuous. It is now clear that C_t models the flow of funds that are not reinvested in the primary assets, but rather put aside forever. By convention, we say that the amount of funds represented by C_t is consumed by the holder of the portfolio up to time t. Note that, the role of the consumption process is not essential in the present context as one can alternatively assume that these funds are reinvested in the risk-free bonds.

In view of the postulated dynamics of the auxiliary process $\hat{S}$ under the real-world measure

$$d\hat{S}_t = (\mu + q)\hat{S}_t dt + \sigma \hat{S}_t dB_t$$

we can use the equality $V_t(\xi, C) = \vartheta_t S_t + \psi_t S_t^0$ to eliminate the component ψ from the self-financing condition. This yields the following equivalent formulation

$$\begin{aligned}
dV_t &= \hat{\vartheta}_t d\hat{S}_t + \psi dS_t^0 - dC_t \\
&= \hat{\vartheta}_t \left((\mu + q)\hat{S}_t dt + \sigma \hat{S}_t dB_t \right) + \psi(rS_t^0 dt) - dC_t \\
&= r\big(\hat{\vartheta}_t \hat{S}_t + \psi S_t^0\big)dt + \hat{\vartheta}_t \left((\mu + q - r)\hat{S}_t dt + \sigma \hat{S}_t dB_t \right) - dC_t \\
&= rV_t dt + \hat{\vartheta}_t \hat{S}_t \big((\mu + q - r)dt + \sigma dB_t\big) - dC_t \\
&= rV_t dt + \hat{\vartheta}_t \hat{S}_t \big((\mu + q - r)dt + \sigma dB_t\big) - dC_t
\end{aligned}$$

We conclude that the wealth process of any self-financing trading and consumption strategy is uniquely determined by the following quantities: the initial endowment V_0, the consumption process C_t, and the process $\hat{\vartheta}_t \hat{S}_t = \vartheta_t S_t$ representing the amount of cash invested in the risky asset. In other words, there is a one to one correspondence between (ξ, C) and

$(\vartheta S, C)$. For the purpose of this section we find it convenient to identify a self-financing trading and consumption strategy (ξ, C) with the corresponding pair $(\vartheta S, C)$ and use this convention without further mention.

Recall now that the unique (spot) martingale measure $\mathbb{Q}$ for the Black Scholes market with dividends satisfies

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_T} = e^{-\frac{\mu+q-r}{\sigma}B_T - \frac{1}{2}\left(\frac{\mu+q-r}{\sigma}\right)^2 T}, \qquad \mathbb{P}-a.s.$$

Substituting $B_t = W_t - \frac{\mu+q-r}{\sigma}t$ in the self-financing equation we get the following dynamics of the wealth process V under the EMM $\mathbb{Q}$

$$dV_t = rV_t dt + \hat{\vartheta}_t \sigma \hat{S}_t dW_t - dC_t$$

It is known that the unique solution of the above SDE is given by the formula

$$V_t = \beta_t^{-1}\left(V_0 + \int_0^t \beta_u \hat{\vartheta}_u \hat{S}_u \sigma dW_u - \int_0^t \beta_u dC_u\right)$$

and with the $\mathbb{Q}$-dynamics of the discounted process $\hat{S}^* = \beta_t \hat{S}_t$ given by

$$d\hat{S}_t^* = d(\beta_t \hat{S}_t) = \sigma \hat{S}_t^* dW_t$$

we can reframe the above expression in terms of $d\hat{S}_t^*$

$$V_t = \beta_t^{-1}\left(V_0 + \int_0^t \hat{\vartheta}_u d\hat{S}_t^* - \int_0^t \beta_u dC_u\right)$$

We proceed by defining an auxiliary process M_t

$$M_t := V_t^* + \int_0^t \beta_u dC_u = V_0 + \int_0^t \hat{\vartheta}_u \hat{S}_u^* \sigma dW_u$$

where $V_t^* = \beta_t V_t$ denotes the discounted wealth process. Because $V_t^* \geq 0$ and C is nondecreasing with $C_0 = 0$, M follows a local martingale under the risk neutral measure $\mathbb{Q}$. The same is true for V_t^*, since the process $\int_0^t \beta_u dC_u$ is nondecreasing.

Similar to the previous section, we will consider only admissible trading strategies in order to exclude pathological examples of arbitrage opportunities from the market model.

Definition 2.18 A self-financing trading and consumption strategy (ξ, C) is admissible if $\vartheta S \in L^2(W)$ for a $\mathbb{Q}$-Brownian motion W. In other words, we restrict ourselves to predictable processes ϑS that satisfy

$$\mathbb{E}_{\mathbb{Q}}\left(\int_0^T (\vartheta_t S_t)^2 d[W]_t\right) = \mathbb{E}_{\mathbb{Q}}\left(\int_0^T \left(\hat{\vartheta}_u \hat{S}_u\right)^2 dt\right) < \infty$$

As a result, a process M driven by an admissible strategy is a $\mathbb{Q}$-martingale. M in turn determines the discounted wealth process V_t^* which is then a supermartingale under $\mathbb{Q}$. We are now in a position to formally introduce the concept of an American style contingent claim.

2.4.4 Arbitrage price of an American claim

For American claims, the payoff function has certain additional properties. We say that a continuous function $g: \mathbb{R}_+ \times [0, T] \longrightarrow \mathbb{R}$ is a reward function if it satisfies the linear growth condition $|g(s,t)| \leq k_1 + k_2 s$ for some constants k_1, k_2. In this way, an American claim with the reward function g and expiry date T is a financial instrument that pays to its holder the amount $g(S_t, t)$ when exercised at time t. The amount obviously depends on the exercise date, but for simplicity, we will denote the payoff at time t by $g(S_t)$. For more exotic options, the reward function will additionally depend on other state variables. We include this notion by viewing S as a vector $\boldsymbol{S}$ taking all the relevant variables into account.

Definition 2.19 An American contingent claim is characterized by a nonnegative adapted right-continuous process $Z = (Z_t)_{0 \leq t \leq T}$ with the reward function $g: \mathbb{R}_+ \longrightarrow \mathbb{R}$. It is a financial instrument consisting of

a. an expiry date T
b. the selection of a stopping time $\tau \in \mathcal{T}_{0,T}$
c. a payoff $Z_\tau = g(S_\tau)$ upon exercise

We also impose the following integrability condition

$$\mathbb{E}_{\mathbb{Q}}\left(\sup_{0 \leq t \leq T} Z_t\right) < \infty$$

Since the interest rate is deterministic, the assumption still holds if we replace Z_t by $Z_t^* = \beta_t Z_t$.

Our aim is to derive the "rational" price and determine the "rational" exercise time of an American contingent claim using purely no-arbitrage line of reasoning. For simplicity, we concentrate on the value at time 0. The general case can be treated similarly but is more cumbersome in terms of notation. It will then be sufficient to consider a very specific class of trading strategies, namely the buy-and-hold strategies.

Definition 2.20 A buy-and-hold strategy associated with a claim Z is a pair (a, τ) where $\tau \in \mathcal{T}_{0,T}$ and a is a real number. It proposes that $a > 0$ units of the American security Z are acquired (or shorted for $a < 0$) at time 0, and then held in the portfolio up to the exercise time τ.

Observe that such strategy excludes trading in the American claim after the initial date. In other words, dynamical trading of American claims will not be considered.

Let us assume there exists a "market" price U_0 at which the American claim Z trades in the market at time 0. Our first task is to find the right value of U_0 by means of arbitrage arguments.

Definition 2.21 A trading strategy in (S^0, S, Z) is a collection $(\xi, C, a, \tau) \coloneqq \Psi$, where (ξ, C) is the trading and consumption strategy in (S^0, S) and (a, τ) is a buy-and-hold strategy associated with claim Z. It is self-financing if on the random interval $(\tau, T]$ we have

$$\vartheta_t = 0\,, \qquad \psi_t = \vartheta_\tau S_\tau \beta_\tau + \psi_\tau + a g(S_\tau)\beta_\tau$$

The above condition implicitly assumes that the American claim is exercised at a random time τ with the existing positions in shares closed and all proceeds invested in the risk-free bond.

In what follows, we restrict ourselves to the class of admissible trading strategies defined as

Definition 2.22 A self-financing trading strategy $\Psi = (\xi, C, a, \tau)$ in (S, B, Z) is said to be admissible if a trading and consumption strategy (ξ, C) is admissible and $C_T = C_\tau$. The class of all admissible strategies (ξ, C, a, τ) is denoted by $\widetilde{\Psi}$.

Furthermore, we introduce the class $\widetilde{\Psi}_0$ of those admissible trading strategies $\widetilde{\Psi}$ such that

1. $V_0(\widetilde{\Psi}) < 0$
2. $V_T(\widetilde{\Psi}) = \psi_T B_t \geq 0$

In order to precisely define an arbitrage opportunity, we have to take into account the early exercise feature of American claims. It is intuitively clear that it is enough to consider two cases – a long and a short position in one unit of an American claim (i.e. for $a = 1$ and $a = -1$). This is due to the fact that we need to exclude the existence of arbitrage opportunities for both the seller and the buyer of an American claim. Indeed, the position of both parties involved in a contract of American style is no longer symmetric as it was in the case of European claims. The holder of an American claim can actively choose his exercise policy. The seller of a claim, on the contrary, should be ready to meet his obligations at any (random) time. We therefore impose the following definition

Definition 2.23 There is absence of arbitrage in the market if the following conditions hold

a. For any stopping time τ and any trading and consumption strategy (ξ, C), the strategy $(\xi, C, 1, \tau)$ is not in $\widetilde{\Psi}_0$

b. For any trading and consumption strategy (ξ, C), there exists a stopping time τ such that the strategy $(\xi, C, -1, \tau)$ is not in $\widetilde{\Psi}_0$

Intuitively, under the absence of arbitrage in the market, the holder of an American claim is unable to find an exercise policy τ and a trading and consumption strategy (ξ, C) that would yield a risk-free profit. On the other hand, under the absence of arbitrage, it is not possible to make a risk-free profit by selling the American claim at time 0, provided that the buyer makes a clever choice of the exercise date. More precisely, there exists an exercise policy for the long party that prevents the short party from locking in a risk-free profit.

By definition, the arbitrage price at time 0 of the American claim Z, denoted by $\pi_0(Z)$, is that level of the price U_0 that makes the model arbitrage-free. Our aim now is to show that the assumed absence of arbitrage in the sense of the above definition leads to a unique value for the arbitrage price $\pi_0(Z)$ of the American claim Z. We shall also find the rational exercise policy of the holder – that is, the stopping time that excludes the possibility of short arbitrage.

The following result relates the value process associated with the specific optimal stopping problem to the wealth process of a certain admissible trading strategy. For any reward function g, we define an adapted process V by setting

$$V_t = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_{\mathbb{Q}}\left(e^{-r(\tau - t)} g(S_\tau) \middle| \mathcal{F}_t\right)$$

for every $t \in [0, T]$, provided that the right-hand side is well-defined.

Proposition 2.24 Let V be an adapted process defined by the above formula for some reward function g. Then there exists an admissible trading and consumption strategy (ξ, C) such that $V_t = V_t(\xi, C)$ for every $t \in [0, T]$.

Proof. Let U denote the Snell envelope (under $\mathbb{Q}$) of the discounted reward process $Z_t^* = e^{-rt} g(S_t)$. By definition, the Snell envelope is the smallest supermartingale dominating Z^* and is given by

$$U_t = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_{\mathbb{Q}}(e^{-r\tau} g(S_\tau) | \mathcal{F}_t) = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_{\mathbb{Q}}(Z_\tau^* | \mathcal{F}_t)$$

for every $t \in [0,T]$, and thus $V_t = e^{rt}U_t$. Since U is a regular supermartingale of class D, it admits the unique Doob-Meyer decomposition $U = M - A$, where M is a (square-integrable) $\mathbb{Q}$- martingale and A is a continuous, nondecreasing process with $A_0 = 0$. Consequently,

$$d(e^{rt}U_t) = re^{rt}U_t dt + e^{rt}dM_t - e^{rt}dA_t$$

The predictable representation property in terms of the $\mathbb{Q}$-Brownian W motion yields for some $\phi \in L^2(W)$ the following decomposition of M

$$M_t = M_0 + \int_0^t \phi_u dW_u \ , \qquad \forall t \in [0,T]$$

Now upon setting

$$\vartheta_t = e^{rt}\phi_t \sigma^{-1} S_t^{-1} \ , \qquad C_t = \int_0^t e^{ru} dA_u$$

we can rewrite M in terms of ϑ

$$M_t = M_0 + \int_0^t \vartheta_u \sigma S_u^* dW_u = M_0 + \int_0^t \vartheta_u (dS_u^* + qS_u^* du)$$

and derive the following expression for V

$$\begin{aligned} dV_t = d(e^{rt}U_t) &= re^{rt}U_t dt + e^{rt}dM_t - e^{rt}dA_t \\ &= rV_t dt + e^{rt}(\vartheta_t \sigma S_t^* dW_t) - dC_t \\ &= rV_t dt + \vartheta_t \sigma S_t dW_t - dC_t \end{aligned}$$

Process ψ enters the above equation by setting

$$\psi_t = U_t - \phi_t \sigma^{-1} = e^{-rt}(V_t - \vartheta_t S_t) \quad \Longrightarrow \quad V_t = \psi_t S_t^0 + \vartheta_t S_t$$

It can easily be verified that the defined trading and consumption strategy is admissible, concluding thereby that process V indeed represents the wealth process of some (admissible) strategy.

■

By the general theory of optimal stopping, we also know that the random time τ_t that maximizes the expected discounted reward after the date t is the first instant at which the Snell envelope process U drops to the level of the discounted reward, i.e.

$$\tau_t = \inf \{u \in [t,T] \mid U_t = Z_u^*\}$$

In other words, the optimal (under $\mathbb{Q}$) exercise policy of the American claim with reward function g is given by the equality

$$\tau_0 = \inf \{t \in [0,T] \mid U_t = e^{-rt} g(S_t)\}$$

Observe that the stopping time τ_0 is well defined (i.e. the set on the right-hand side is non-empty with probability 1), and necessarily

$$V_{\tau_0} = g(S_{\tau_0})$$

In addition, the stopped process $U_{t \wedge \tau_0}$ is a martingale, so that the process A is constant on the interval $[0, \tau_0]$. This means also that $C_t = 0$ on $[0, \tau_0]$, implying there is no consumption present before time τ_0. We find it convenient to introduce the following definition

Definition 2.25 An admissible trading and consumption strategy (ξ, C) is a perfect hedging strategy against the American contingent claim $Z = g(S)$ if its wealth process V satisfies

$$V_t(\phi) \geq g(S_t), \qquad \forall t \in [0,T]$$

with probability 1. We denote by $\Phi(Z)$ the class of all perfect hedging strategies against the American claim Z.

From the majorizing property of the Snell envelope, we can infer that the trading and consumption strategy (ξ, C) introduced in the proof of Proposition 2.24 is a perfect hedging strategy against the American claim Z. Moreover, this strategy has the special property of minimal initial endowment amongst all admissible perfect hedging strategies for Z.

Finally, the following theorem determines $\pi_0(Z)$ explicitly, by assuming that trading in the American claim would not destroy the arbitrage-free features of the Black-Scholes model.

Theorem 2.26 There is no arbitrage in the market model with trading in an American claim if and only if the price $\pi_0(Z)$ is given by the formula

$$\pi_0(Z) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}_{\mathbb{Q}}\left(e^{-r\tau} g(S_\tau)\right)$$

More general, the arbitrage price at time t of an American claim with reward function g equals

$$\pi_t(Z) = V_t = \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_{\mathbb{Q}}\left(e^{-r(\tau - t)} g(S_\tau) \middle| \mathcal{F}_t\right)$$

The rational price formula for American-style options is sometimes also referred to as the Karatzas formula.

2.4.5 Markovian setting

Note that so far we have implicitly restricted our attention to path-independent contingent claims – that is, to claims whose payoff at exercise depends on the value of the underlying asset at the exercise date only. To incorporate the notion of path dependent options in our setting we need the following results regarding the Markovian environment of our model.

Recall that the asset price dynamics under the risk neutral measure $\mathbb{Q}$ are given by

$$dS_t = (r - q)S_t dt + \sigma S_t dW_t$$

and observe that the parameters r, q, σ are deterministic and that the filtration is one generated by the Brownian motion. From general theory on stochastic analysis we know that the Brownian motion W has the Markov property with respect to $\mathbb{F}^W$. Since S is a one-to-one transformation of the Brownian motions, this leads to the following conclusion

Proposition 2.27 The stock price S and the discounted stock price S^* are Markov processes under the EMM $\mathbb{Q}$ with respect to the filtration $\mathbb{F}^S = \mathbb{F}^{S^*}$.

Consider now a simple European claim $H = f(S_T)$ for some bounded Borel measurable payoff function $f: [0, \infty) \to \mathbb{R}$. Since S is Markovian we have

$$e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[f(S_T)|\mathcal{F}_t] = e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[f(S_T)|S_t = s]$$

Strictly speaking, the payoff of a European claim is a function $f(S_T)$ of the stock price at maturity. By the Markov property, the arbitrage price of this option is obtained as the conditional expectation of the discounted payoff given the current stock price. Furthermore, for some arbitrary t, h the following equality also holds

$$e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[f(S_{t+h})|\mathcal{F}_t] = e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[f(S_{t+h})|S_t = s] = e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[f(S_h)|S_0 = S_t]$$

In this sense, we can interpret S_{t+h} as a new process S_h starting at the current state S_t. Which means that, once the current state of the process is known, all future changes will be determined as if the process had just started, independent of the its history up to the given time. This property can be extended to bounded random times.

For a reward function $g: [0, \infty) \to \mathbb{R}$ and any stopping time τ such that $\mathbb{Q}(\tau < \infty) = 1$, we define the strong Markov property of a Brownian motion W with respect to $\mathbb{F}^W$

$$\mathbb{E}_{\mathbb{Q}}[f(W_{\tau+t})|\mathcal{F}_\tau] = \mathbb{E}_{\mathbb{Q}}[f(W_{\tau+h})|W_\tau]$$

From there it can be deduced that a Brownian motion W starts fresh after any finite stopping time τ. That is, the process $\widetilde{W} = \left(\widetilde{W}_t\right)_{0\leq t\leq T}$ given by the equation $\widetilde{W}_t = W_{\tau+t} - W_\tau$ is a Brownian motion with respect to its natural filtration $\mathbb{F}^{\widetilde{W}}$.

Now recall that for an American contingent claim Z with payoff function g we have

$$V_t(Z) = \operatorname*{ess\,sup}_{\tau\in\mathcal{T}_{[t,T]}} \mathbb{E}_{\mathbb{Q}}\left(e^{-r(\tau-t)}g(S_\tau)\middle|\mathcal{F}_t\right)$$

Since S is strongly Markovian, we can derive the following result for the price process of an American claim (see [24])

$$V_t = V(S_t, t)$$

where

$$V(S,t) = \sup_{\tau\in\mathcal{T}_{[t,T]}} \mathbb{E}_{\mathbb{Q}}\left[e^{-r(\tau-t)}g\left(S_\tau^{t,S}\right)\right]$$

$$\equiv \sup_{\tau\in\mathcal{T}_{[t,T]}} \mathbb{E}_{\mathbb{Q}}\left[e^{-r(\tau-t)}g(S_\tau)\middle|S_t = S\right]$$

In particular at time $t = 0$ we have

$$V(S,0) = \sup_{\tau\in\mathcal{T}_{[0,T]}} \mathbb{E}_{\mathbb{Q}}\left[e^{-r\tau}g(S_\tau)|S_0 = S\right]$$

It can be shown directly by arguing that the supremum and the essential supremum are the same when $\mathcal{T}_{[t,T]}$ is replaced by the set of stopping times with respect to the filtration generated by the Brownian motion increments. This equality is essentially the expression of the Snell envelope. Mathematically we have moved from pricing a family of random variables contingent on the whole market history $\mathcal{F}_t$ to a family contingent only on the variable S_t, making the formula more analytically tractable. Financially we are pricing a claim whose drift (in the real world) is stochastic, and possibly path-dependent. However, this drift is linked (through the market price of risk) to the drift of the asset price, all path dependency collapses to dependence only on the current price of the underlying asset.

2.4.6 The curse of dimensionality

In this section we turn our attention to American Asian options of arithmetic averaging. We begin by deriving the expression for the arbitrage price of the option and show how the optimal stopping problem of a floating strike Asian option can be reduced from a three- to a two-dimensional one. This can be particularly useful when applying a Monte Carlo simulation, as it makes the pricing more efficient in terms of computing time.

An application of the Karatzas formula yields the arbitrage-free price for the floating-strike Asian option

$$V_t = \underset{\tau \in \mathcal{T}_{[t,T]}}{\text{ess sup}}\, \mathbb{E}_{\mathbb{Q}} \left[e^{-r(\tau - t)} max(\varrho(S_\tau - A_\tau), 0) \middle| \mathcal{F}_t \right]$$

and the fixed type

$$V_t = \underset{\tau \in \mathcal{T}_{[t,T]}}{\text{ess sup}}\, \mathbb{E}_{\mathbb{Q}} \left[e^{-r(\tau - t)} max(\varrho(A_\tau - K), 0) \middle| \mathcal{F}_t \right]$$

The above optimal stopping problems are three-dimensional but changing the numéraire which amounts to a change of measure approach, the problem reduces to a two-dimensional. An appropriate transformation can be found for any floating strike Asian option (arithmetic or geometric, European or American). As for the fixed strike Asian, this is true only for the European type.

We therefore turn our attention to the floating strike case; take the underlying asset as the numéraire and define the new equivalent martingale measure $\mathbb{Q}^*$ by

$$\left. \frac{d\widetilde{\mathbb{Q}}}{d\mathbb{Q}} \right|_{\mathcal{F}_T} = e^{-(r-q)T} \frac{S_T}{S_0} = e^{-\frac{1}{2}\sigma^2 T + \sigma W_T}$$

where via Girsanov's theorem $\widetilde{W}_t \equiv -W_t + \sigma t$ is a standard Brownian motion under $\widetilde{\mathbb{Q}}$. The expectation of $\int_t^T \frac{S_u}{S_t} dt$ depends on the drift and volatility parameters, with no dependence on the state variables S_t and tA_t. In the following, it will be convenient to use tA_t instead of A_t as the averaging state variable. We therefore proceed by defining the new state variable $Y_t \coloneqq tA_t / S_t$ and derive the rational pricing formula for the American Asian floating-strike as

$$V_t = \underset{\tau \in \mathcal{T}_{[t,T]}}{\text{ess sup}}\, \mathbb{E}_{\mathbb{Q}} \left[e^{-(q+r-q)(\tau - t)} S_t \frac{S_\tau}{S_t} max\left(\varrho \left(1 - \frac{1}{\tau} Y_\tau \right), 0 \right) \middle| \mathcal{F}_t \right]$$

$$= \underset{\tau \in \mathcal{T}_{[t,T]}}{\text{ess sup}}\, S_t \mathbb{E}_{\widetilde{\mathbb{Q}}} \left[e^{-q(\tau - t)} max\left(\varrho \left(1 - \frac{1}{\tau} Y_\tau \right), 0 \right) \middle| \mathcal{F}_t \right]$$

where the change of measure performed does not affect the set of stopping times over which the claim holder optimizes. Using Itô's Lemma we have that Y evolves according to the stochastic differential equation

$$dY_t = Y_t \left[\left(q - r + \frac{1}{Y_t} \right) dt + \sigma d\tilde{S}_t \right]$$

Therefore, we can then view Y as the underlying asset, with dividend yield $r - \frac{1}{Y_t}$ in the auxiliary economy with interest rate q. In this way, the optimal stopping problem of an American Asian call with floating strike no longer depends on three state variables, but two. With a slight abuse of notation, we have the following result

$$C_{float}(S_t, tA_t, t) = S_t C_{float}(Y_t, t)$$

2.4.7 Concluding remark

In the preceding chapters we have successfully derived the Black Scholes pricing formula in terms of a discounted expectation of some future payoff. Unfortunately, this does not mean that we already have a way to compute these prices. When it comes to arithmetic Asian options, it turns out that the expectation at hand is quite difficult to evaluate, and even though they have attracted much attention in recent literature, there is still no closed-form solution available.

To elaborate, note that for an Asian option the role of the averaging mechanism is the same as that of S in the payoff function. On one side, the geometric average of stock prices G_{ave} is lognormally distributed (as a product of lognormal random variables) and based on the Black Scholes formula, the price for geometric average option can be derived straightforward.

On the other hand, the arithmetic averaging mechanism A_{ave} does not have a lognormal distribution. This feature makes the task of finding an explicit formula for the price of an Asian option surprisingly involved. Early studies were based either on approximations or on the direct application of the Monte Carlo method. The numerical approach proposed by [18] is based on the approximation of the arithmetic average using the geometric average. It appears however that such an approach significantly underprices Asian call options. To overcome this deficiency, one may directly approximate the true distribution of the arithmetic average by a lognormal distribution with appropriate parameters (first and second moments of A_{ave}) and use the Black Scholes pricing formula.

Nevertheless, in this paper we adapt a direct simulation approach. Monte Carlo simulation is convenient and flexible and is very useful to Asian options which are highly path dependent. It is applicable as long as the underlying follows a Markovian diffusion.

2.5 The duality principle in option pricing

Monte Carlo simulation is a useful technique used to evaluate options for which no closed form solution is available. However, another tool was proven to be of extreme importance when dealing with Asian options in particular. To motivate the results of the chapter we begin with an example.

Consider an investor who buys a call option in the foreign exchange markets, for example the Euro/Dollar market. Owning the call option, he has the right to buy euros for a strike rate K. In case he exercises the option, he has to pay in dollars. Therefore, the right to buy euros is at the same time a right to sell dollars at the inverse rate $\frac{1}{K}$. Thus, the call option on the Euro/Dollar rate is equivalent to a put option on the Dollar/Euro rate. The prices of these options determine each other.

Behind this simple observation lies a much deeper result, which we call the duality principle. Duality relations arise when the price process is taken as a numéraire which leads to the notion of the dual measure and process. The simplest example is the call-put duality described above. It allows to express the price of a call in terms of a put written on the dual process with inverse strike. These tricks become very useful when perhaps no closed form solution exists, but an equivalence relation holds. We can then transfer knowledge of one option to another and simplify the whole valuation procedure. In view of Monte Carlo simulations, it would save a great deal of computational effort if we could relate one option to another that is already coded. This is of particular interest for the Asian option since much is known about the fixed strike case, but comparatively little work has been done for the floating strike option.

We will lay out the theory on duality in a bit more general setup where the underlying asset follows an exponential Lévy process X. Later we argue how this can be applied to the Black Scholes environment. We will consider both a European and American setting. In order to better illustrate the principle at hand, we first look at plain vanilla options, and conclude with a symmetry result for Asian options.

Definition 2.28 An adapted process $X = (X_t)_{t \geq 0}$ is a standard Lévy process on $(\Omega, \mathbb{F}, \mathbb{Q})$ if the following properties hold

1. $\mathbb{Q}(X_0 = 0) = 1$
2. The paths $t \longrightarrow \mathrm{X}_{\mathrm{t}}(\omega)$ are a.s. càdlàg (right continuous with left-hand side limits)
3. X has stationary increments $\mathcal{L}(X_t - X_s) = \mathcal{L}(X_{t-s})$, $\forall s \leq t$
4. The increments of X are independent of the past: $X_t - X_s$ independent of $\sigma(X_u : u \leq s)$

We assume that the initial measure $\mathbb{Q}$ is the unique martingale measure. All our calculations will be done with respect to this measure. As an introduction, we begin with a simplified setting. The following results are easy to prove, and offer a good illustration of the duality principle.

2.5.1 Simplified setting

To reduce notation we assume that the current interest rate is zero and the price process $S = (S_t)_{0 \le t \le T}$ is a non-dividend paying stock. We begin by defining the dual measure and process. A change of probability measure or numéraire and a time reversal argument are then used to prove the results for models where the underlying asset follows an exponential Lévy process.

Let $S = \exp(X)$ be a martingale under the measure $\mathbb{Q}$. Thus $\mathbb{E}_{\mathbb{Q}}(S_T) = 1$ which allows us to define the dual probability measure $\mathbb{Q}'$ by

$$\left.\frac{d\mathbb{Q}'}{d\mathbb{Q}}\right|_{\mathcal{F}_T} = S_T$$

The dual process S' is $$S' = \frac{1}{S}$$

Then with $X' = -X$ we have $$S' = \exp(X')$$

A known result from martingale theory states that if $Z = \frac{d\mathbb{Q}'}{d\mathbb{Q}}$ then S' is a $\mathbb{Q}'$ martingale if and only if $S'Z$ is a $\mathbb{Q}$-martingale. In our case $Z = S$ and $S'S \equiv 1$, thus proving the following lemma, crucial to duality theory.

Lemma 2.29 Suppose that $S = \exp(X)$ is a $\mathbb{Q}$-martingale. Then the process S' is a martingale with respect to the dual measure $\mathbb{Q}'$.

A. Standard call and put options of European type

A European option H is characterized by a Borel measurable payoff function $h: [0, \infty) \to \mathbb{R}$ which yields the payoff $H = h(S_T)$ at maturity T. As is known from previous chapters, in a complete market where the martingale measure $\mathbb{Q}$ is unique, the rational (arbitrage free) price of a claim is given by

$$\mathbb{E}_{\mathbb{Q}}\big(h(S_T)\big)$$

In the case of a standard call respectively put with underlying process S, strike price $K > 0$ and maturity T the payoff functions are

$$H_{call} = (S_T - K)^+$$

$$H_{put} = (K - S_T)^+$$

The corresponding option prices are given by the formulae

$$C_T(S,K) = \mathbb{E}(S_T - K)^+$$

$$P_T(K,S) = \mathbb{E}(K - S_T)^+$$

These expectations are taken with respect to the initial martingale measure $\mathbb{Q}$. If the expectation is to be taken under the dual measure $\mathbb{Q}'$, we denote the expectation operator by $\mathbb{E}'$, and the corresponding prices $C'_T(S,K)$, $P'_T(K,S)$. We will also use the notation $K' = \frac{1}{K}$.

Proposition 2.30 (Call-Put Duality)

For a standard call and put, option prices satisfy the following duality relations

$$\frac{1}{K} C_T(S,K) = P'_T(K',S')$$

$$\frac{1}{K} P_T(K,S) = C'_T(S',K')$$

Proof. We offer a proof for the first symmetry result. The second result can be obtained in the same way. Since the density $\mathbb{Q}'$ with respect to $\mathbb{Q}$ is S_T, by changing to the dual measure we get

$$\begin{aligned} C_T(S,K) &= \mathbb{E}(S_T - K)^+ \\ &= \mathbb{E}'\left(\frac{(S_T - K)^+}{S_T}\right) \\ &= \mathbb{E}'(1 - KS_T)^+ \\ &= K\mathbb{E}'\left(\frac{1}{K} - S_T\right)^+ \\ &= KP'_T(S',K') \end{aligned}$$

■

From the identity $(S_T - K)^+ = (K - S_T)^+ + S_T - K$ we can obtain the well-known "put-call parity" by simply taking the expectation with respect to measure $\mathbb{Q}$

$$C_T(S,K) = P_T(K,S) + 1 - K$$

Corollary 2.31 Calls and put prices in markets $(S,\mathbb{Q})$ and $(S',\mathbb{Q}')$ which satisfy the duality relation, are connected by the following "call-call parity"

$$C_T(S,K) = KC'_T(S',K') + 1 - K$$

and the following "put-put parity"

$$P_T(K,S) = K'P'_T(K',S') + K - 1$$

B. Fixed and Floating strike Asian options

Turning our attention to floating strike Asian options, the duality principle allows to replace them with simpler fixed strike Asians, given the following property holds.

Definition 2.32 X' satisfies the time inversion principle if we have the following equality (in law)

$$X'_T - X'_{(T-t)-} \sim X'_t$$

Note that the time inversion principle holds for the processes under consideration.

Proposition 2.33 (Fixed-Floating Asian option Duality)
Let X' satisfy the time inversion principle. Then the price of an Asian option with floating strike equals the price of an Asian option with fixed strike in the following way

$$C_T\left(S,\frac{1}{T}\int S\right) = P'_T\left(1,\frac{1}{T}\int S'\right)$$

$$P_T\left(\frac{1}{T}\int S\,,S\right) = C'_T\left(\frac{1}{T}\int S'\,,1\right)$$

Proof. Again, by employing the change of measure we have

$$C_T\left(S,\frac{1}{T}\int Sdt\right) = \mathbb{E}\left(S_T - \frac{1}{T}\int_0^T S_t dt\right)^+$$

$$= \mathbb{E}'\left(1 - \frac{1}{T}\int_0^T \frac{S_t}{S_T}dt\right)^+$$

$$= \mathbb{E}'\left(1 - \frac{1}{T}\int_0^T \frac{S'_T}{S'_t}dt\right)^+$$

$$= \mathbb{E}'\left(1 - \frac{1}{T}\int_0^T e^{X'_T - X'_t} dt\right)^+$$

$$= \mathbb{E}'\left(1 - \frac{1}{T}\int_0^T e^{X'_T - X'_{(T-u)-}} du\right)^+$$

The last equality follows from a change of variables and the properties of the Lebesgue integral for càdlàg functions. By the time inversion principle, the final term equals

$$\mathbb{E}'\left(1 - \frac{1}{T}\int_0^T e^{X'_u} du\right)^+ = \mathbb{E}'\left(1 - \frac{1}{T}\int_0^T S'_u du\right)^+ = P'_T\left(1, \frac{1}{T}\int S'\right)$$

■

C. Standard call and put options of American type

Consider now an American claim $Z = (Z_t)_{0 \le t \le T}$ characterized by a reward function $g: \mathbb{R}_+ \times [0, T] \longrightarrow \mathbb{R}$. If the claim is exercised at a random time $\tau \in \mathcal{T}_{[0,T]}$ we would receive a payoff $g(S_\tau, \tau)$. In the case of a standard call and put options this translates to

$$Z^\tau_{call} = (S_\tau - K)^+$$

$$Z^\tau_{put} = (K - S_\tau)^+$$

The price of an American put and call option are given by

$$\hat{C}_T(S, K) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}(S_\tau - K)^+$$

$$\hat{P}_T(S, K) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}(K - S_\tau)^+$$

Using the change of measure and tower property of conditional expectations we have

$$\hat{C}_T(S, K) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}(S_\tau - K)^+ = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'\left(\frac{(S_\tau - K)^+}{S_T}\right)$$

$$= \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'((S_\tau - K)^+ S'_T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'\big((S_\tau - K)^+ \mathbb{E}'(S'_T | \mathcal{F}_\tau)\big)$$

$$= \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'((S_\tau - K)^+ S'_\tau) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'((1 - KS'_\tau)^+)$$

$$= K \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'((K' - S'_\tau)^+) = \sup_{\tau \in \mathcal{T}_{[0,T]}} \mathbb{E}'((S_\tau - K)^+ S'_T)$$

$$= K\hat{P}'_T(K', S')$$

Thus, similar to the results of the European case, we have

Proposition 2.34 (American Call-Put Duality)

$$\frac{1}{K}\hat{C}_T(S,K) = \hat{P}'_T(S',K')$$

$$\frac{1}{K}\hat{P}_T(S,K) = \hat{C}'_T(S',K')$$

2.5.2 Black Scholes Merton setting

In this section we specify the assumptions of a Lévy market and show how they extend to our model.

Definition 2.35 A Lévy market consists of two assets. A deterministic savings account $B = (B_t)_{0 \le t \le T}$ with

$$B_t = e^{rt}, \qquad r \ge 0$$

where we take $B_0 = 1$ for simplicity, and a stock $S = (S_t)_{0 \le t \le T}$ with random evolution modelled as

$$S_t = S_0 \exp(X_t), \qquad S_0 = e^x > 0$$

where $X = (X_t)_{0 \le t \le T}$ is a Lévy process.

First, note that the standard Brownian motion B (starting a.s. at 0) is indeed a Lévy process satisfying an additional distributional requirement and right continuity of paths. To apply these results to our model we will assume that the stock pays dividends at a constant rate $q \ge 0$ and that the given probability measure $\mathbb{Q}$ is the (unique) EMM derived at the very beginning of the chapter. In other words, prices are computed as expectations with respect to $\mathbb{Q}$, and the discounted and reinvested process $\hat{S}^*_t = e^{-(r-q)} S_t$ is a $\mathbb{Q}$-martingale.

The Black Scholes Merton model consists of two assets with $\mathbb{Q}$-dynamics given by

$$dB_t = rB_t dt$$

$$dS_t = (r-q)S_t dt + \sigma S_t dW_t$$

where W is a standard $\mathbb{Q}$-Brownian motion. Since the parameters r, q, σ are deterministic, the above stochastic differential equations have unique solutions given by

$$B_t = \mathcal{E}(rt) = \exp\left(rt - \frac{1}{2}[r]_t\right) = e^{rt}$$

$$S_t = \mathcal{E}\big((r-q)t + \sigma W_t\big) = S_0 \exp\left((r-q)t + \sigma W_t - \frac{1}{2}[\sigma W]_t\right)$$

$$= S_0 \exp\left(\left(r - q - \frac{1}{2}\sigma^2\right)t + \sigma W_t\right)$$

This is exactly the case we would get if we were to model X in the definition of the Lévy market as an exponential Brownian motion $e^{\left(r-q-\frac{1}{2}\sigma^2\right)t+\sigma W_t}$.

Now consider the martingale $Z = (Z_t)_{t\geq 0}$ defined by

$$Z_t = \frac{e^{-rt}S_t}{e^{-qt}S_0} = \exp\left(-\frac{1}{2}\sigma^2 t + \sigma W_t\right)$$

Z can act as the Radon Nikodym density process of measure change and will be used to introduce the notion of a dual measure as follows.

Definition 2.36 The dual martingale measure $\mathbb{Q}'$ for the Black Scholes Merton model is defined by changing the numéraire to S via

$$\left.\frac{d\mathbb{Q}'}{d\mathbb{Q}}\right|_{\mathcal{F}_T} = Z_T = \frac{e^{-rt}S_T}{e^{-qt}S_0} = \exp\left(-\frac{1}{2}\sigma^2 T + \sigma W_T\right)$$

The dual process Z' is defined as

$$Z' = \frac{1}{Z} = \frac{e^{-qT}S_0}{e^{-rT}S_T} = \exp\left(\frac{1}{2}\sigma^2 T - \sigma W_T'\right)$$

where (via Girsanov's theorem) $W_t' = W_t - \sigma t$ is a standard $\mathbb{Q}'$-Brownian motion.

For simplicity of notation, we define the process $X = (X_t)_{t\geq 0}$ by $X_t = -\frac{1}{2}\sigma^2 t + \sigma W_t$, and rewrite Z in terms of X as

$$Z_t = \exp X_t$$

Now as in the previous section, by defining $X' = -X'$, we can express the dual process Z'_T as

$$Z'_T = \exp X'_T$$

The dual process Z'_t is a martingale with respect to the dual measure $\mathbb{Q}'$.

A. Call and Put of European type

With payoff functions defined just as in the previous section, we denote the prices of a plan vanilla call and put option respectively by

$$c(S_0, K, r, q, 0, T) = e^{-rT}\mathbb{E}(S_T - K)^+$$

$$p(S_0, K, r, q, 0, T) = e^{-rT}\mathbb{E}(K - S_T)^+$$

Theorem 2.37 (European Call-Put Duality)
With the prices of a European call and put defined as above, we have the following symmetry

$$c(S_0, K, r, q, 0, T) = p(K, S_0, q, r, 0, T)$$

Note that the roles of the interest rate r and dividend q have been reversed in the symmetry results.

Proof. We start from a call option and apply a change of measure in order to obtain

$$c(S_0, K, r, q, 0, T) = e^{-rT}\mathbb{E}(S_T - K)^+ = e^{-rT}\mathbb{E}(S_0 e^{X_T} - K)^+$$

$$= \mathbb{E}\big(e^{-rT}e^{X_T}(S_0 - Ke^{-X_T})\big)^+$$

$$= \mathbb{E}\left(e^{-qT}Z_T\big(S_0 - Ke^{X'_T}\big)\right)^+$$

$$= e^{-qT}\mathbb{E}'\big((S_0 - KZ'_T)\big)^+$$

To conclude the proof, i.e. to show that

$$e^{-qT}\mathbb{E}'\big((S_0 - KZ'_T)\big)^+ = p(K, S_0, q, r, 0, T)$$

we need to verify that the dual process Z'_T follows a log normal distribution under the dual measure $\mathbb{Q}'$. We have

$$Z'_T = \exp X'_T = \exp(-X_T) = \exp\left(\frac{1}{2}\sigma^2 T - \sigma W'_T\right)$$

where W'_t is a standard $\mathbb{Q}'$-Brownian motion. Obviously $-W'_t$ is the reflected $\mathbb{Q}'$-Brownian motion starting at zero where the distributional requirement $\mathcal{L}\big(-(W'_t - W'_s)\big) = N(0, t - s)$ follows from the symmetry of the mean zero Gaussian distribution.

■

B. Call and Put of American type

To extend the theory to American options, recall that $\mathcal{T}$ denotes the set of all stopping times with respect to $\mathbb{F}$, and $\mathcal{T}_{[0,\mathrm{T}]}$ the class of stopping times up to a fixed time $T < \infty$ defined as

$$\mathcal{T}_{[0,\mathrm{T}]} = \{\tau \in \mathcal{T} : \mathbb{P}(\tau \in [0, T]) = 1\}, \qquad 0 \leq t \leq T < \infty$$

We proceed by replacing the fixed, deterministic time can be replaced by some bounded stopping time $\tau \in \mathcal{T}_{[0,T]}$ to obtain

$$\left.\frac{d\mathbb{Q}'_\tau}{d\mathbb{Q}_\tau}\right|_{\mathcal{F}_\tau} = Z_\tau = \exp X_\tau$$

where $\mathbb{Q}_\tau$ (resp. $\mathbb{Q}'_\tau$) is the restriction of $\mathbb{Q}$ (resp. $\mathbb{Q}'$) to the σ-algebra

$$\mathcal{F}_\tau = \{F \in \mathcal{F} | F \cap \{\tau \leq t\} \in \mathcal{F}_t\}$$

The expectations in this section will be taken with respect to the restricted martingale measure $\mathbb{Q}_\tau$ and restricted dual measure $\mathbb{Q}'_\tau$. The prices for an American call and put option are then respectively defined by

$$C(S_0, K, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau} \mathbb{E}(S_\tau - K)^+$$

$$P(S_0, Kr, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau} \mathbb{E}(K - S_\tau)^+$$

The following result is obtained simply by taking the supremum over all $\tau \in \mathcal{T}_{[0,T]}$ in the duality relation of European put and call options.

Theorem 2.38 (American Call-Put Duality)
With the value functions of an American call and put defined as above, the following symmetry holds

$$C(S_0, K, r, q, 0, T) = P(K, S_0, q, r, 0, T)$$

C. Asian options of European type

In this section we present the results form Henderson and Wojakovski [16], which unlike other symmetry results, relate two different types of arithmetic Asian options. Consistent with the previous sections, the expectations are taken under the martingale measure $\mathbb{Q}$ and the dual measure $\mathbb{Q}'$, where $Z = \exp X$ denotes the process of measure change and $Z' = \exp X'$ the corresponding dual process.

The arithmetic average A is calculated at time $t = T$ given the price history from time 0. Formally we have

$$A_T = \frac{1}{T}\int_0^T S_t dt$$

The fixed strike Asian option is written on the average of asset prices A_T and assumes a fixed strike K. By arbitrage arguments, the time 0 price of a fixed strike Asian call resp. put is defined as

$$c_x(K, S_0, r, q, 0, T) = e^{-rT}\mathbb{E}(A_T - K)^+$$

$$p_x(K, S_0, r, q, 0, T) = e^{-rT}\mathbb{E}(K - A_T)^+$$

The floating strike Asian is typically interpreted as written on S with floating strike A_T. When exercising, the holder receives or buys λ units of stock and pays the average of past prices A_T. The price of a floating strike Asian call resp. put is given by

$$c_f(S_0, \lambda, r, q, 0, T) = e^{-rT}\mathbb{E}(\lambda S_T - A_T)^+$$

$$p_f(S_0, \lambda, r, q, 0, T) = e^{-rT}\mathbb{E}(A_T - \lambda S_T)^+$$

Theorem 2.39 (Fixed-Floating European Asian Duality)
With the notation as above, the following duality relations hold

a. $c_f(S_0, \lambda, r, q, 0, T) = p_x(\lambda S_0, S_0, q, r, 0, T)$
b. $c_x(K, S_0, r, q, 0, T) = p_f\left(S_0, \frac{K}{S_0}, q, r, 0, T\right)$

Proof. We begin by proving the symmetry in a). Consider the floating strike Asian call with price (at time $t = 0$) expressed in units of stock as numéraire

$$c_f' = \frac{c_f(S_0, \lambda, r, q, 0, T)}{S_0} = \frac{e^{-rT}}{S_0}\mathbb{E}(\lambda S_T - A_T)^+$$

$$= \mathbb{E}\left(\frac{S_T e^{-rT}}{S_0}\frac{(\lambda S_T - A_T)^+}{S_T}\right) = \mathbb{E}\left(e^{-qT}Z_T\frac{(\lambda S_T - A_T)^+}{S_T}\right)$$

Note that

$$\frac{(\lambda S_T - A_T)^+}{S_T} = (\lambda - A_T')^+$$

where A_T' denotes the average of stock prices denominated in units of stock S

$$A_T' \equiv \frac{A_T}{S_T} = \int_0^T \frac{S_u}{S_T}du$$

By passing to the dual measure $\mathbb{Q}'$ defined via $d\mathbb{Q}' = Z_T d\mathbb{Q}$ where Z is the process of measure change defined at the beginning of the section, we have

$$c_f' = \mathbb{E}\left(e^{-qT}Z_T\frac{(\lambda S_T - A_T)^+}{S_T}\right) = e^{-qT}\mathbb{E}'(\lambda - A_T')^+$$

This is essentially a fixed strike put option written on an average of asset prices A_T' and a fixed strike $K = \lambda$. To conclude the proof, we consider the distribution of A_T'.

Remember that for the dual measure, the process $W_t' = W_t - \sigma t$ is a standard $\mathbb{Q}'$-Brownian motion. In terms of Brownian motion W' we have

$$\frac{S_u}{S_T} = \exp\left(-\left(r - q - \frac{1}{2}\sigma^2\right)(T - u) - \sigma(W_T - W_u)\right)$$

$$= \exp\left(\left(r - q + \frac{1}{2}\sigma^2\right)(u - T) - \sigma(W_u' - W_T')\right)$$

Now by defining the reflected $\mathbb{Q}'$-Brownian motion $\widetilde{W}_t \equiv -W_t'$ we have (in law)

$$W_u' - W_T' \equiv \widetilde{W}_T - \widetilde{W}_u \sim \widetilde{W}_{T-u}$$

Thus obtaining the equality (in law)

$$A'_T \sim \tilde{A}_T \equiv \frac{1}{T}\int_0^T \exp\left(\sigma \widetilde{W}_{T-u} + \left(r - q + \frac{1}{2}\sigma^2\right)(u - T)\right) du$$

And now by resorting to a variable change $s = T - u$, we can apply the time inversion principle of Brownian motions

$$\widetilde{W}'_T - \widetilde{W}'_{(T-s)-} \sim \widetilde{W}'_s$$

to obtain

$$\tilde{A}_T = \frac{1}{T}\int_0^T \exp\left(\sigma \widetilde{W}_s - \left(r - q + \frac{1}{2}\sigma^2\right) s\right) ds$$

From here it is obvious that $\frac{S_u}{S_T}$ are log-normally distributed and $A'_T \sim \tilde{A}_T$ is a sum of such log-normally distributed variates. We thus have

$$c'_f = e^{-qT}\mathbb{E}'(\lambda - A'_T)^+ = e^{-qT}\mathbb{E}'\left(\lambda - \tilde{A}_T\right)^+$$

which concludes the proof for a). The second relation can be obtained using an analogous procedure.

Alternatively, we can use the put-call parity results for Asian options. Starting this time with a fixed strike call

$$c_x(K, S_0, r, q, 0, T) = e^{-rT}\mathbb{E}(A_T - K)^+$$

and consider the put-call parity for Asian options of floating strike

$$p_f(S_0, \lambda, r, q, 0, T) - c_f(S_0, \lambda, r, q, 0, T) = \frac{1}{(r - q)T}(e^{-qT} - e^{-rT})S_0 - \lambda S_0$$

and the analogous for Asian options of fixed strike

$$c_x(K, S_0, r, q, 0, T) - p_x(K, S_0, r, q, 0, T) = \frac{1}{(r - q)T}(e^{-qT} - e^{-rT})S_0 - e^{-rT}K$$

Combining these and the symmetry in a), the results in b) follow.

∎

D. Asian options of American type

Gounden and O'Hara [14] derive an analogue result to the one in the previous section for Asian options of American type. Their proof involves once again a change of numéraire and time reversal of Brownian motion which they argue can be carried through to the American case as the change of measure performed will not affect the set of stopping times over which the claim holder optimizes.

The stock average process A_t is calculated at time $t \in [0, T]$ given the price history from the initial time $t = 0$. The continuous (arithmetic) average process is defined as

$$A_t = \frac{1}{t} \int_0^t S_u du$$

For Asian options of American type, we denote the price of a fixed call resp. put by

$$C_x(K, S_0, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau} \mathbb{E}(A_\tau - K)^+$$

$$P_x(K, S_0, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau} \mathbb{E}(K - A_\tau)^+$$

and the price of a floating call resp. put by

$$C_f(S_0, \lambda, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau} \mathbb{E}(\lambda S_\tau - A_\tau)^+$$

$$P_f(S_0, \lambda, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau} \mathbb{E}(A_\tau - \lambda S_\tau)^+$$

With the notation as above, Gounden and O'Hara argue the following duality relation

a. $C_f(S, \lambda, r, q, 0, T) = P_x(\lambda S, S, q, r, 0, T)$
b. $C_x(K, S, r, q, 0, T) = P_f\left(S, \frac{K}{S}, q, r, 0, T\right)$

In the final chapter we use explicit examples in an attempt to verify the noted symmetries within the Black Scholes Merton framework.

Chapter 3

Monte Carlo Simulation

"Consider a game of heads and tails. We suspect that the coin is not symmetric and want to investigate how this influences the game. A simple solution would be to toss the coin many times, and to compute the average of the outcomes. This way we can estimate the expected outcome of the game. Such computation of expectation is the underlying idea behind Monte Carlo methods."

Behind that simple idea lies a very powerful tool for estimating integrals and expected values. Monte Carlo is a random algorithm, so a different run of simulation will yield a different estimate, making it unlike any other deterministic numerical schemes usually designed for problems of lower dimensions. On the contrary, it is very flexible when it comes to accommodating models that are otherwise inaccessible. For example, it is widely used in financial industry, where we often need to estimate integrals of larger or infinite dimensions. The reason why it works so well compared to other deterministic methods is that the time taken to carry out a Monte Carlo simulation increases approximately linearly with the number of variables, while for most others that increase is exponential. Moreover, the error of the estimate decays in the order of $\mathcal{O}(\frac{1}{\sqrt{m}})$ with respect to the sample size m, regardless of the dimension.

However, the method is not without shortcomings. In its basic form, it is computationally inefficient in the sense that a very large number of simulations are generally required to achieve a high degree of accuracy. Even though its efficiency can be improved using some form of variance reduction techniques, little can be done to accelerate the convergence above the rate of $\mathcal{O}(\frac{1}{\sqrt{m}})$. On top of that, as we will see in the following chapter, it cannot easily handle situations where there are early exercise opportunities. This chapter aims to give a quick introduction to Monte Carlo simulation applied to pricing European options, as well as its pros and cons.

3.1 Basics

In practice, the method chosen to value options is likely to depend on the characteristics of the derivative being evaluated and the accuracy required. Monte Carlo simulation can cope with a great deal of complexity but is of particular use when the payoff is dependent on the history of the underlying variable. This makes Monte Carlo just the right fit for estimating Asian option prices. We begin with a demonstration of the method for plain vanilla options of European exercise type and show how it can be extended to estimate exotic Asian puts and calls.

In the previous chapter we have discussed the theory behind risk neutral evaluation of a derivative security in a complete market within the Black Scholes Merton setting. For a European claim H with payoff at maturity specified by a function h of the underlying asset price process, i.e. $H = h(S_T)$ the arbitrage price is given by

$$\mathbb{E}_{\mathbb{Q}}[e^{-rT}h(S_T)]$$

The name implies that only by trading at the specified price do we not introduce arbitrage in our (otherwise arbitrage free) model. Note that the expectation is taken with respect to the (unique) risk neutral measure $\mathbb{Q}$, the only such measure under which both the value process and the underlying price process are martingales. With the expression at hand, we still need a way to compute the price. The simplest way to calculate the expected discounted payoff is to simulate a (large) number of the entire price path of the underlying asset, calculate the payoff and discount at the risk-free rate. The approximate value of the derivative is an arithmetic average of the discounted payoffs for each path.

Note that we will be simulating paths for a process whose dynamics are given under the risk neutral measure $\mathbb{Q}$, not the real-world measure $\mathbb{P}$. This means that we can price options by Monte Carlo simulations without estimating the drift of underlying assets or modeling stochastic discount factors and still get the same price as with no-arbitrage valuation. The approach becomes more efficient as we increase the number of simulated paths.

Intuitively, this procedure easily extends to Asian options. The only difference is the calculation of the payoff, which now depends on the average of the simulated prices of the entire path, not only its spot price at expiration. So we would again simulate the a price path of the underlying, and now compute the average along it, calculate the payoff at expiration and discount it. The estimate of the derivative price is then obtained by taking the average over all paths. This sounds reasonable since the only added difficulty is the computation of the average, which can be done whilst simulating paths.

The Monte Carlo scheme can be summarized in four distinct actions

1. Simulating stock price paths
2. Calculating the payoffs at maturity for each given path
3. Discounting the payoffs at the risk-free rate
4. Averaging discounted payoffs over all paths

Observe that the four-step procedure described above assumes that the risk neutral measure is derived and the dynamics of the process under this measure are given explicitly. In reality, we would first need to replace the drift in the initial $\mathbb{P}$-equations. For a dividend paying stock, by substituting μ with $(r-q)$ we implicitly describe the dynamics in a risk neutral world. Only with this modification can we resort to simulation. Assuming the dynamics of the stock process under the risk neutral measure we now formalize exactly how this is done.

Consider a derivative dependent on a single market variable that provides a payoff $h(S_T)$ at maturity. We want to estimate the price of the derivative at time $t=0$, i.e.

$$H = \mathbb{E}_{\mathbb{Q}}[e^{-rT}h(S_T)]$$

The dynamics of the asset price process in a risk neutral world are the following

$$dS = (r-q)Sdt + \sigma SdW$$

where the random variable W is a standard Brownian motion under the risk neutral measure $\mathbb{Q}$. Using Itô's formula, the solution to the above SDE determines the following random evolution of the stock

$$S_t = S_0 exp\left[\left(r - q - \frac{\sigma^2}{2}\right)t + \sigma W_t\right]$$

As S_0 is the current price of the stock (we are computing the price at time zero), we may assume it is known. Since W is a Brownian motion, the exponential of such a process is a geometric Brownian motion, and we will use the notation $S \sim GMB(r-q,\sigma^2)$ to indicate that S follows the above dynamics.

The payoff of a standard European call/put option is determined by the terminal stock price S_T and does not otherwise depend on the evolution of S_t between times 0 and T. However, in valuing derivative securities, it is often necessary to simulate the entire random stream over multiple intermediate dates and not just at the initial and terminal date. Two considerations make this necessary:

- the payoff of a derivative security may depend explicitly on the values of underlying assets at multiple dates (path dependent options)
- we may not know how to sample transitions of the underlying assets exactly and thus need to divide a time interval $[0, T]$ into smaller subintervals to obtain a more accurate approximation to sampling from the distribution at time T

So, to simulate the payoffs at maturity we will in fact need to simulate a whole path for S. We describe this procedure in the following section.

3.1.1 Random walk construction

The previous section suggests a simulation of an entire a stream of values for the option during its lifetime from 0 to T. To this end, we begin by dividing the interval $[0, T]$ into n equal time steps of length $\Delta t = \frac{T}{n} = t_j - t_{j-1}$, $j = 1,2, \dots, n$.

The evolution of the stock price process follows a geometric Brownian motion with mean $(r - q)$ and variance σ^2

$$S_t = S_0 exp\left[\left(r - q - \frac{\sigma^2}{2}\right)t + \sigma W_t\right]$$

The use of a geometric Brownian motion to model stock prices is convenient since it does not assume negative values. More importantly, unlike for an ordinary Brownian motion, the percentage changes (returns)

$$\frac{S_{t_2} - S_{t_1}}{S_{t_1}}, \frac{S_{t_3} - S_{t_2}}{S_{t_2}}, \dots, \frac{S_{t_n} - S_{t_{n-1}}}{S_{t_{n-1}}}$$

are independent. This independence of the returns allows for a more general representation of the stock equation. For $u < t$ we have

$$S_t = S_u exp\left[\left(r - q - \frac{\sigma^2}{2}\right)(t - u) + \sigma(W_t - W_u)\right]$$

It is evident that the randomness of the price process arises from the driving Brownian motion. So, as the above equation suggests, a simulation of a Markov chain $\{S_{t_1}, S_{t_2}, \dots, S_{t_n}\}$ will entail simulating increments $\{W_{t_1}, W_{t_2} - W_{t_1}, \dots, W_{t_n} - W_{t_{n-1}}\}$ of the standard Brownian motion W.

Recall that for an ordinary Brownian motion the increments are independent and normally distributed, i.e. $W_{t+\Delta t} - W_t \sim N(0, \Delta t)$. This is also the distribution of $\sqrt{\Delta t}Z$ if $Z \sim N(0,1)$.

Now let $Z_1, Z_2, \dots, Z_n$ be independent standard normal random variables. To simulate a standard Brownian motion (starting a.s. at zero) we set $t_0 = 0$ and $W_0 = 0$. Subsequent values can then be generated in the following way

$$W_{t_{j+1}} = W_{t_j} + \sqrt{t_{j+1} - t_j}Z\,, \qquad j = 1,2, \dots, n-1$$

This gives a simple recursive procedure for simulating S. Starting at the known state (initial price) S_0 and over each subinterval $[t, t + \Delta t]$ we approximate the governing equation for the stock price as

$$S_{t+\Delta t} = S_t exp\left[\left((r-q) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z\right]$$

This allows the value of S at time Δt to be calculated from the initial value S_0, the value at time $2\Delta t$ from the preceding value at Δt, and so on. This method is exact, in the sense that the produced estimates of the Markov chain $\{S_{t_1}, S_{t_2}, \dots, S_{t_n}\}$ have the same joint distribution as $S \sim GBM(r-q, \sigma^2)$ but only at the fixed time points $t_1, t_2, \dots, t_n$ of simulation. There is no way of saying what happens between these time points, so the distribution is subject to discretization error compared to a true Brownian motion. We can attempt to extend the simulated values to other time point through linear interpolation for example. However, it can be shown that this introduces further discretization error even at the simulated time points $t_1, t_2, \dots, t_n$ and will thus not be discussed in this paper.

One simulation trial involves constructing a complete path for S using n independent draws from a normal distribution. Repeating this procedure for *n* steps produces a value of $\hat{S}_T$ whose distribution approximates the exact distribution. We expect that as n becomes larger (so that Δt becomes smaller) the approximating distribution of $\hat{S}_T$ draws closer to the exact distribution.

In practice, it is usually more accurate to simulate lnS rather than S. From Itô's lemma the process followed by lnS is

$$dlnS_t = \left((r-q) - \frac{\sigma^2}{2}\right)dt + \sigma dW$$

And using the same arguments as above we have

$$lnS_{t+\Delta t} - lnS_t = \left((r-q) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z$$

At each time step we would then simply take the exponential of the simulated logarithmic prices.

3.1.2 Monte Carlo scheme for plain vanilla options

Finally, we formalize our initial four-step procedure for approximating option prices via Monte Carlo simulation. It will involve simulating m independent paths of n time steps each. For simplicity we will denote Z_{i,t_j} by Z_{ij}. To be precise, we use Z_{ij} to denote the j-th draw from the standard normal distribution along the i-th path. The $\{Z_{ij}\}$ are mutually independent and identically distributed as Z in a risk neutral world. We do the following

1. Simulate a Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$ of the stock price in a risk neutral world
   For $j = 1,2, \dots, n$
   a. Generate a sample Z_{ij} from the standard normal distribution $Z_{ij} \sim N(0,1)$
   b. Simulate the logarithmic spot price

$$lnS_{ij} = lnS_{i,j-1} + \left((r-d) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z, \qquad \Delta t = t_j - t_{j-1}$$

   c. The spot price at time step t_j then equals $S_{ij} = \exp\left(lnS_{ij}\right)$
2. Calculate the payoff $h(S_{in})$ for the computed path
3. Discount each payoff at the risk-free rate to get an estimate of the value of the derivative for the i-th path

$$H_i = e^{-rT}h(S_{in})$$

4. Repeat steps 1. to 3. for $i = 1,2, \dots, m$, and then calculate the mean of the sample payoffs to get an estimate of the expected payoff in a risk-neutral world

$$\widehat{H}_m = \frac{1}{m}[H_1 + H_2 + \cdots + H_m]$$

3.1.3 Monte Carlo scheme for Asian options

The previously described algorithm suffices for approximation of simple options like vanilla puts or calls. In order to map Asian options into our framework we would additionally have to compute spot averages along each path. Let A_{iT} denote the average along the i-th path calculated at maturity given the spot price history from time 0. The modified procedure reads

1. Simulate a Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$ of the stock price in a risk neutral world
   Set $A_{i1} = S_0$ and for $j = 1,2, \dots, n$ do the following
   a. Generate a sample Z_{ij} from the standard normal distribution, $Z_{ij} \sim N(0,1)$
   b. Simulate the logarithmic spot price

$$lnS_{ij} = lnS_{i,j-1} + \left((r-d) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z, \qquad \Delta t = t_j - t_{j-1}$$

   c. Calculate the spot price $S_{ij} = \exp(lnS_{ij})$
   d. Add the simulated spot price into the sum $A_{ij} = A_{i,j-1} + S_{ij}$
2. Calculate the payoff $h(A_{iT})$ for the computed path
   a. For a given path compute the arithmetic average of spot prices at maturity

$$A_{iT} = \frac{A_{in}}{n}$$

   b. Calculate $h(A_{iT})$
3. Discount each payoff at the risk-free rate to get an estimate of the value of the derivative for the i-th path

$$H_i = e^{-rT} h(A_{iT})$$

4. Repeat steps 1. to 3. for $i = 1,2, \dots, m$, and then calculate the mean of the sample payoffs to get an estimate of the expected payoff in a risk-neutral world

$$\widehat{H}_m = \frac{1}{m}[H_1 + H_2 + \cdots + H_m]$$

Observe that the implementation of the described algorithm would only entail storing a single price path at a time. When the payoffs for a given path are calculated and discounted, we can add them into a sum, after which we simply simulate the next path in place of the current one. The Monte Carlo estimate is then obtained by averaging the summed-up payoffs.

3.1.4 Generating random samples from the normal distribution

The former chapter suggest that, in order to carry out a Monte Carlo simulation, we have to be able to repeatedly sample from a given distribution. At the core of this procedure is a random number generator that produces iid sequences of uniformly distributed random numbers on $[0,1]$. In reality, it is not actually possible to generate a truly random number. Instead we come up with (deterministic) algorithms like "Box-Muller" that mimic randomness, when in fact the numbers produced are only pseudo random. However, for practical purposes, we treat the said numbers as if they were de facto random.

Now assume that we have a way to generate random numbers of the interval $[0,1]$, and can thus simulate uniformly distributed random variables. The following method is by far the simplest and most commonly used in, and it relies on one key observation. The inverse transform method asserts that any random variable can be represent as a function of a uniformly distributed random variable.

To that end, suppose we have a random variable X whose cumulative distribution function F is known. We have that $F(x) = \mathbb{P}(X \leq x)$ and define the inverse of F as

$$F^{-1}(u) = \min\{x : F(x) \geq u\}, \qquad u \in [0,1]$$

Consequently, for some uniformly distributed random variable U on $[0,1]$, the inverse transform method says that by setting

$$X = F^{-1}(U)$$

the cumulative distribution function of X is in fact F. Therefore, once we obtain the inverse of the underlying distribution, then $X = F^{-1}(U)$ will have the desired distribution if $U \sim Uniform([0,1])$. Indeed, we have

$$\mathbb{P}(X \leq x) = \mathbb{P}(F^{-1}(U) \leq x) = \mathbb{P}\big(U \leq F(x)\big) = F(x)$$

The method is applicable even for distribution with jumps, but for the time being we turn back to our setting where the underlying distribution is the standard normal $N(0,1)$. The observations in this section suggest the following reframe of step 1. a. of the Monte Carlo scheme

a. Generate a sample Z_{ij} from the standard normal distribution, $Z_{ij} \sim N(0,1)$

   i. Generate a sample u from the uniform distribution on $[0,1]$
   ii. Set $Z_{ij} = \Phi^{-1}(u)$

3.2 Quality of an estimate

This chapter aims to address further developments of the realized method. To that end, recall that the quantity H we wish to estimate is an unknown fixed number, whereas the Monte Carlo estimate $\widehat{H}_m$ is a random variable. To discuss improvements, we first need to define the criteria used to determine the quality of an estimate. Three considerations will be of particular importance: bias, variance and computational time.

1. *Bias*

In the absence of bias, estimator variability and computational effort are the most important considerations. However, reducing variability or computational time would be pointless if it merely accelerated convergence to an incorrect value. In our case of concern, the value of $\widehat{H}_m$ will of course vary depending on the drawn samples, but is, for any $m \geq 1$, unbiased (in the sense that its expectation is the target quantity):

$$\mathbb{E}\left[\widehat{H}_m\right] = H \equiv \mathbb{E}[e^{-rT}h(S_T)]$$

We continue the discussion of efficiency under the assumption that the compared estimators are unbiased.

2. *Variance*

In Monte Carlo simulation, in addition to the estimate, it is customary to report the standard error. We show how it can be calculated and demonstrate the effect it has on the quality of an estimate.

Let H be the unknown quantity of interest and $H_1, H_2, \ldots, H_n$ independent and identically distributed random variables such that $\mathbb{E}[H_i] = H$ and $Var[H_i] = \sigma_H^2 < \infty$. A Monte Carlo estimate for H is then:

$$\widehat{H}_n = \frac{1}{n}\sum_{i=1}^{n} H_i$$

Then, by the strong law of large numbers, the above empirical arithmetic averages are strongly consistent, meaning that

$$\lim_{n\to\infty} \widehat{H}_n = \mathbb{E}(H_1)\,, \qquad a.s.$$

Simply put, $\widehat{H}_n$ is close to H when n is large. But how close? Since $\widehat{H}_n$ is a random variable, so is the error $\widehat{H}_n - H$, and what we are really looking for, is the distribution of this error, not the error bound in general sense.

If $\widehat{H}_n$ has finite variance σ_H^2, it follows from the central limit theorem that

$$\frac{H_1 + H_2 + \cdots + H_n - nH}{\sigma_H \sqrt{n}} \xrightarrow{n \to \infty} \frac{\sqrt{n}\left(\widehat{H}_n - H\right)}{\sigma_H}$$

Meaning that, as the sample size n tends to infinity, the above distributions converge weakly to the standard normal distribution $N(0,1)$. More formally, for every $a \in \mathbb{R}$ we have

$$\mathbb{P}\left\{\frac{\sqrt{n}\left(\widehat{H}_n - H\right)}{\sigma_H} \leq a\right\} \to \Phi(a)$$

And therefore, the error $\widehat{H}_n - H$ is approximately normally distributed with mean 0 and variance σ_H^2/n.

This asymptotic analysis also produces confidence intervals for the Monte Carlo simulation, which can, for moderately large n, be used as a supplement to the point estimate $\widehat{H}_n$. From the above equation it follows that as n tends to infinity, the $1 - \delta$ confidence interval for H is approximately

$$\widehat{H}_n \pm z_{\delta/2} \frac{\sigma_H}{\sqrt{n}}$$

where z_δ is the $1 - \delta$ quantile of the standard normal distribution (i.e. $\Phi(z_\delta) = 1 - \delta$). Which is to say that the probability of the confidence interval actually containing the true value H is approaching $1 - \delta$ as $n \to \infty$. This shows that uncertainty about the value of the derivative is inversely proportional to the square root of the number of trials, meaning the width of a confidence interval decreases as the sample size increases. If one quadruples the sample size, the width is reduced by half. To increase the accuracy by a factor of 10, the number of trials must increase by a factor of 100, and so on.

The standard deviation σ_H in the above formulation is rarely known in practice. Instead, the sample standard deviation s_H is used

$$s_H = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(H_i - \widehat{H}_n\right)^2} = \sqrt{\frac{1}{n-1}\left(\sum_{i=1}^{n} H_i^2 - n\widehat{H}_n^2\right)}$$

Note that an unbiased estimator of the actual variance involves the term $\frac{1}{n}$, while the unbiased estimator of the sample variance has the term $\frac{1}{n-1}$ instead. This substitution is appropriate since s_H converges to σ_H with probability one when the sample size n tends to infinity, and hence the central limit theorem still holds if σ_H is replaced by s_H. The empirical1 $1-\delta$ confidence interval thus becomes

$$\widehat{H}_n \pm z_{\delta/2}\frac{s_H}{\sqrt{n}}$$

Therefore, the commonly used 95% confidence interval for the price H of the derivative is just the estimate $\widehat{H}_n$ plus/minus twice the standard error (note that for a 95% confidence interval, $\delta = 0.05$ and $z_{\delta/2=0.025} \approx 1.96$).

$$\widehat{H}_n - \frac{1.96\ s_H}{\sqrt{n}} < H < \widehat{H}_n + \frac{1.96\ s_H}{\sqrt{n}}$$

The quantity $s_H/\sqrt{n}$ is often called the standard error of $\widehat{H}_n$ and we denote it by

$$S.E. = \sqrt{\frac{1}{n(n-1)}\left(\sum_{i=1}^{n} H_i^2 - n\widehat{H}_n^2\right)}$$

Since the standard error of H determines the width of a confidence interval, in can be inferred that in some sense it also determines the size of the error $\widehat{H}_n - H$. Given a fixed sample size, the smaller the variance, the tighter the confidence interval, and hence the more accurate the estimate. This leads to the following criterion: other things being equal, in comparing two estimates of the same quantity, the one with the smaller variance is more efficient.

From here on, when using Monte Carlo simulation, we will report both the estimate and the standard error as it is just as important. The following step is added to our scheme

5. Calculate the standard error of the produced estimate $\widehat{H}_m$
   a. Estimate the sample variance by

$$Var = \frac{1}{m-1}\sum_{i=1}^{m}\left(H_i - \widehat{H}_m\right)^2$$

   b. And calculate the standard error as

$$S.E. = \sqrt{\frac{Var}{m}}$$

3. *Computational time*

Naturally, the accuracy of the results obtained via simulation must depends on the number of trials. Other things equal, the estimator that requires less computational effort is more efficient. But what happens if we are given a choice between two unbiased estimators, and that the one with smaller variance takes longer to compute? How does one balance variance reduction and computational effort? We give a brief discussion leading to a simple enough criterion and for more details refer the reader to (Glasserman, 2003).

Suppose that generating a replication H_i takes a fixed amount of computing time k. Let s denote our "computational budget", and $\widehat{H}_{\lfloor s/k \rfloor}$ an estimate for the number of replications we can complete given budget s. In this case we can use the normalized formulation of the central limit theorem

$$\sqrt{s}\left[\widehat{H}_{\lfloor s/k \rfloor} - H\right] \xrightarrow{n \to \infty} N(0, \sigma_H^2 k)$$

which tells us that given a budget s, the approximate distribution of the estimate error is a normal distribution with mean 0 and variance $\sigma_H^2 k/s$. This property suggests the following criteria when choosing between a faster, more variable estimator, and a slower, less variable one. Asymptotically (as the computational budget grows), we should prefer the estimator with the smaller value of:

$$\text{(variance per replication) x (computing time per replication)}$$

because this is the one that will produce a more precise estimate (and narrower confidence interval) from the restricted budget s. We would also like to point out that the requirement leading to this result (constant computing time per replication) is satisfied in the case of Asian options, i.e. all replications require simulating the same number of transitions, and the time per transition is nearly constant.

3.3 Variance reduction techniques

Now that we know which factors influence the quality of an estimate, we point out some fundamental techniques of improvement.

Estimating a value of a derivative via Monte Carlo calls upon simulation of the underlying price process. A large number of trials is usually necessary for this estimate to be of reasonable accuracy. This can be very expensive in terms of computational effort. However, variance reduction procedures described in this chapter can be used to dramatically reduce computation time. In the previous chapter we have determined the standard error of $\widehat{H}_n$ to be $\sigma/\sqrt{n}$, with σ^2 denoting the variance of $\widehat{H}_n$. From there we can identify two sources that contribute to the standard error of an estimate. First is the factor $1/\sqrt{n}$. Except for increasing the number of simulated paths over which we estimate, not much can be done about it. The second one is the variance σ^2 of the sample estimate $\widehat{H}_n$. As the name already suggest, it can be reduced using a number of techniques. Four standard methods to reduce σ are:

- Antithetic Variates
- Control Variate
- Stratification
- Importance Sampling

In this section, we closely examine the first two. A detailed exposition on variance reduction techniques can be found in [1].

3.3.1 Antithetic sampling

In plain Monte Carlo simulation, samples are independent and identically distributed. The idea behind antithetic sampling is to reduce the variance by introducing samples that are negatively correlated. One simulation trial involves calculating two values of the derivative. The first value X is calculated in the usual way, while the second value Y is calculated simply by changing the sign of all the random samples from standard normal distributions. In other words, if ϵ is a sample used to calculate X, then $-\epsilon$ is the corresponding sample used to calculate Y. The sample value of the derivative is calculated as the arithmetic average of the two values, X and Y. This works well because when one value is above the true value, the other tends to be below, and vice versa.

To be more precise, consider the following two schemes for estimating $\mathbb{E}[X]$. In both scenarios, the estimate is the sample average.

1. *Plain Monte Carlo:* $2n$ iid samples $X_1, \dots, X_n, X_{n+1}, \dots, X_{2n}$
   The estimate is
   $$\frac{1}{2n}\sum_{i=1}^{2n} X_i$$
2. *Antithetic Sampling:* $2n$ samples (or n pairs of samples)
   $$\begin{matrix} X_1 & X_2 & \dots & X_n \\ Y_1 & Y_2 & \dots & Y_n \end{matrix}$$
   Pairs of samples (X_i, Y_i) are iid. Y_i has the same distribution as X_i, and they are dependent.

The estimate is

$$\hat{v} = \frac{1}{n}\sum_{i=1}^{n}\frac{X_i + Y_i}{2}$$

For brevity, we denote $\sigma^2 = Var[X_i]$. The plain Monte Carlo estimate is unbiased, and the variance associated to it is

$$Var\left[\frac{1}{2n}\sum_{i=1}^{2n} X_i\right] = \frac{1}{2n}\sigma^2$$

In the antithetic approach, since (X_i, Y_i) are independent and Y_i has the same distribution as X_i, the estimate $\hat{v}$ is again unbiased and

$$Var[\hat{v}] = \frac{1}{4n^2}\sum_{i=1}^{n} Var[X_i + Y_i]$$

Now suppose that the correlation coefficient between X_i and Y_i is ρ. From general probability theory we know that

$$Var[X_i + Y_i] = Var[X_i] + Var[Y_i] + 2Cov(X_i, Y_i) = 2\sigma^2 + 2\rho\sigma^2$$

Which in turn implies for the variance

$$Var[\hat{v}] = \frac{1}{2n}\sigma^2 + \frac{\rho}{2n}\sigma^2$$

The above expression is minimized for negative values of ρ, which leads to the conclusion that antithetic sampling reduces the variance when X_i and Y_i are negatively correlated. The degree of improvement is characterized by the magnitude of ρ. So a stronger negative correlation of X_i and Y_i, implies a more significant variance reduction. Note that if the samples X_i and Y_i are made to be positively correlated, then antithetic sampling will actually increase the variance and make the estimate less accurate.

In theory, it is always possible to construct a negatively correlated antithetic sample. Indeed, when discussing the inverse transform method we have seen that one can represent any random variable X as $X = F^{-1}(U)$ where U is a random variable uniformly distributed on $[0,1]$ and F^{-1} is the inverse of the cumulative distribution function of X. An antithetic sample of X can then be determined as $Y = F^{-1}(1-U)$. Clearly Y has the same distribution as X since $1-U$ is also uniformly distributed on $[0,1]$. Furthermore, X and Y are negatively correlated because X is an increasing function of U and Y is a decreasing function of U.

The standard deviation associated with antithetic sampling is

$$\sqrt{\frac{1}{n} Var\left[\frac{(X_i + Y_i)}{2}\right]}$$

By replacing the variance with sample variance, we obtain the following standard error

$$S.E. = \sqrt{\frac{1}{n(n-1)}\left(\sum_{i=1}^{n}\left[\frac{X_i + Y_i}{2}\right]^2 - n\hat{v}^2\right)}$$

which is usually much smaller than the standard error obtained when using $2n$ random trials.

The noted observations are used accordingly to improve the Monte Carlo scheme. We point out only the affected steps.

1. Simulate the Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$ with corresponding antithetic values $\{\check{S}_{i1}, \check{S}_{i2}, \dots, \check{S}_{in}\}$ of the stock price in a risk neutral world

   For $j = 1,2, \dots, n$

   a. Generate a sample Z_{ij} from the standard normal distribution, $Z_{ij} \sim N(0,1)$
   b. Set $lnS_{ij} = lnS_{i,j-1} + \left((r-d) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z$

      and $ln\check{S}_{ij} = ln\check{S}_{i,j-1} + \left((r-d) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z$
   c. Set $S_{ij} = \exp\left(lnS_{ij}\right)$

      and $\check{S}_{ij} = \exp\left(ln\check{S}_{ij}\right)$
2. Calculate the payoffs $h(S_{in})$ and $h\left(\check{S}_{in}\right)$ for the computed path
3. The estimate of the value of the derivative for the i-th path is the (discounted) average of the two antithetic payoffs

$$H_i = e^{-rT}\frac{\left(h(S_{in}) + h\left(\check{S}_{in}\right)\right)}{2}$$

3.3.2 Control variate

The control variate technique is suited when we want to approximate an option X, for which there exists a similar option Y whose value can be explicitly calculated. In fact, Asian options are a textbook example for the application of this technique since there exists no tractable solution for an arithmetic Asia, but there is one for a geometric Asian option. The key point is that we use the same numerical procedure to approximate both options. We carry out the simulation in parallel, simulate the same number of paths with the same number of time steps, and most importantly, using the same random number streams. If we denote by $\widehat{H}_X^*$ the resulting estimate for X, and by $\widehat{H}_Y^*$ that of Y, we can improve the estimated value for X in the following way

$$H_X = \widehat{H}_X^* - \widehat{H}_Y^* + H_Y$$

To approximate the value of Y when an accurate value H_Y already exists may seem like a waste of time. However, it can be shown that if the estimation errors arising from the approximation of X and Y are unbiased (or equally biased) and highly correlate, the technique produces a better estimate for the value of X and reduces the error of computation without having to increase the number of simulations. What is more, the scheme can be designed in such a wat that the sampling of both X and Y requires little extra computational effort compared to sampling X alone.

Suppose now we are interested in estimating the expected value $\mathbb{E}[X]$. The difference between the plain Monte Carlo scheme and the one using a control variate is demonstrated below.

1. *Plain Monte Carlo:* n iid samples $X_1, \ldots, X_n$
   The estimate is
$$\frac{1}{n}\sum_{i=1}^{n} X_i$$

2. *Control Variates:* n samples $\{X_i\}$ and n control variate samples $\{Y_i\}$
$$\begin{matrix} X_1 & X_2 & \ldots & X_n \\ Y_1 & Y_2 & \ldots & Y_n \end{matrix}$$
   Pairs of samples (X_i, Y_i) are iid. Y_i has known expected value $\bar{\mu}$.
   The estimate is
$$\hat{v} = \frac{1}{n}\sum_{i=1}^{n} X_i - b\left(\frac{1}{n}\sum_{i=1}^{n} Y_i - \bar{\mu}\right)$$

   where b is a fixed constant.

Both of the above estimates are clearly unbiased. Now denote by σ_X^2 the variance of X_i, by σ_Y^2 the variance of Y_i, and with ρ the correlation coefficient between X_i and Y_i. The variance of the plain Monte Carlo estimate is

$$\frac{1}{n}\sigma_X^2$$

As for the control variate estimate $\hat{v}$, we write

$$\hat{v} = \frac{1}{n}\sum_{i=1}^{n} H_i \ , \qquad H_i = X_i - b(Y_i - \bar{\mu})$$

Note that $H_1, \dots, H_n$ are iid random variables, implying that the variance of the control variate estimate can be written as

$$\frac{1}{n^2}\sum_{i=1}^{n} Var[H_i] = \frac{1}{n}[\sigma_X^2 - 2b\rho\sigma_X\sigma_Y + b^2\sigma_Y^2]$$

Since all the above variables are otherwise determined, the size of the variance reduction depends on the choice of the coefficient b. Therefore, the optimal choice for b is the one that minimizes the variance of the control variate estimate

$$b^* = \rho\frac{\sigma_X}{\sigma_Y} = \frac{Cov(X,Y)}{Var(Y)}$$

If one uses b^*, the variance of the control variate estimate comes down to

$$\frac{1}{n}(1-\rho^2)\sigma_X^2$$

In other words, the control variate approach reduces the variance of the plain Monte Carlo estimate by a factor of ρ^2. More often than not, there are many ways to select an appropriate control for X. Ideally it should be in strong correlation with X, either positive or negative. For that reason, the chosen control Y is often of similar structure to X.

Again, by replacing the variance with the sample variance, we can calculate the standard error associated with the control variate estimate $\hat{v}$ as

$$S.E. = \sqrt{\frac{1}{n(n-1)}\left(\sum_{i=1}^{n} H_i^2 - n\hat{v}^2\right)}$$

3.3.2.2 The optimal coefficient b^*

The optimal coefficient b^* is an unknown quantity in general and must itself be estimated. One way of doing that is by means of the sample variance and covariance of samples (X_i, Y_i). More precisely

$$\hat{b}^* = \frac{\sum_{i=1}^{n}(X_i - \bar{X})\,(Y_i - \bar{Y})}{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}\ , \qquad \bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i\ , \qquad \bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i$$

Since the sample size n is large, the estimate $\hat{b}^*$ will be close to the actual optimal b^*. However, it should be pointed out that by using $\hat{b}^*$ as an estimate, we introduce bias. The reason is that in general

$$\mathbb{E}\left[\hat{b}^*\left(\frac{1}{n}\sum_{i=1}^{n} Y_i - \bar{\mu}\right)\right] \neq 0$$

since $\hat{b}^*$ is clearly no longer independent of $\{Y_i\}$. The introduced bias is in practice usually very small compared to the standard error of the estimate and can thus be safely ignored.

Nonetheless, if one wishes to eliminate this bias, he could alternatively estimate b^* from samples other than (X_i, Y_i) (often called pilot samples). In that case b^* will be independent of $\{Y_i\}$, and the unbiasedness of the control variate estimate preserved. Since we do not want to incur a lot of extra computational effort, the size of the pilot samples is generally much smaller than the size of the simulation n. Another way to eliminate the bias is to forgo optimality of coefficient b. For example, for two positively correlated random variables of similar structures, we can let $b = 1$ or choose a quantity based on prior experience. Although a bit naïve, it turns out that this simple approach can be quite efficient sometimes.

To conclude, we point out the following fact which proved to be the case for Asian options. In terms of Monte Carlo simulation, as the option values are just estimates, finding an appropriate control is an imperative to the improvement of the pricing performance of one's simulation. For instance, in Asian options with a larger strike the simulation becomes less efficient even when the standard errors decrease. This is attributed to the fact that the estimates grow smaller much faster than the corresponding error. On top of all, an increase of intermediate time steps at which we simulate barely affects the performance. Since an antithetic method doesn't turn out to be that efficient, this makes the choice of an appropriate control all the more necessary. For this reason, we present the following example containing a brief discussion on the control variate technique for arithmetic Asian options.

Example 2.40 Analytically tractable derivatives as control variates

We want to estimate the price of an arithmetic Asian call option with fixed strike and maturity T. In terms of improving the standard error and producing narrower confidence intervals, we implement a control variate technique. Since a suitable control should have a high degree of correlation with the variable we want to estimate, a natural guess, and probably the best one can find, is the geometric Asian call option. The variance reduction achieved is significant even with the naive choice of $b = 1$, which reflects the strong correlation between the arithmetic and geometric mean.

The underlying asset price process follows a log-normal distribution. The product of log-normally distributed random variables is again log normal; hence we can expect that the pricing of geometric average Asians will be fairly easy to deal with. In fact, the pricing formula can be derived within the Black Scholes framework. For a call option, we have the following closed form expression

$$H_G^{call} = e^{-rT}\left\{\exp\left(\mu_G + \frac{1}{2}\sigma_G^2\right)\Phi(d_1) - K\Phi(\mathrm{d}_2)\right\}$$

With

$$\mu_G = \ln S_0 + \left(r - q - \frac{1}{2}\sigma^2\right)\frac{T + \Delta t}{2}$$

$$\sigma_G^2 = \sigma^2 \Delta t \frac{(2n+1)(n+1)}{6n}$$

And

$$d_1 = \frac{\mu_G - \ln K + \sigma_G^2}{\sigma_G}$$

$$d_2 = d_1 - \sigma_G$$

Now let $\bar{S}_A$ denote the arithmetic, and $\bar{S}_G$ the geometric mean of stock prices (for a given path)

$$\bar{S}_A = \frac{1}{n}\sum_{j=1}^{n} S_j \ , \qquad \bar{S}_G = \left(\prod_{j=1}^{n} S_j\right)^{\frac{1}{n}}$$

The exact price p of an Asian call option in the Black Scholes framework is determined using the above noted formulae. A control variate estimate for the price of an Asian call with arithmetic mean is then a sample average of iid copies of

$$e^{-rT}(\bar{S}_A - K)^+ - b[e^{-rT}(\bar{S}_G - K)^+ - p]$$

For the purpose of introducing the control variate technique into our Monte Carlo scheme, we let $b = 1$ or $b = \hat{b}^*$. The algorithm extends as follows

1. Simulate a Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$ of the stock price in a risk neutral world

   For $j = 1,2,\dots,n$ do the following

   a. Generate a sample Z_{ij} from the standard normal distribution, $Z_{ij} \sim N(0,1)$

   b. Simulate the logarithmic spot price

$$lnS_{ij} = lnS_{i,j-1} + \left((r-d) - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}Z, \qquad \Delta t = t_j - t_{j-1}$$

2. Calculate the payoffs $h(\bar{S}_A^i)$ and $h(\bar{S}_G^i)$ for the computed path

   a. Compute the arithmetic mean

$$\bar{S}_A^i = \frac{1}{n}\sum_{j=1}^{n} S_j^i$$

   b. Compute the geometric mean

$$\bar{S}_G^i = \left(\prod_{j=1}^{n} S_j^i\right)^{\frac{1}{n}}$$

3. Set $X_i = e^{-rT} h(\bar{S}_A^i)$

   Set $Y_i = e^{-rT} h(\bar{S}_G^i) - p$

   Set $H_i = X_i - bY_i$ for $b = 1$ or $b = \hat{b}^*$

4. Repeat steps 1. to 3. for $i = 1,2,\dots,m$, and compute the estimate and the standard error

$$\hat{H}_m = \frac{1}{m}[H_1 + H_2 + \cdots + H_m]$$

$$S.E. = \sqrt{\frac{1}{m(m-1)}\left(\sum_{i=1}^{m} H_i^2 - m\hat{H}_m^2\right)}$$

Chapter 4

Monte Carlo for American Options

Unlike European options, which can only be exercised at maturity, American options include early exercise, early termination or callable features. These features can be seen as further rights for the holder of the option as they allow the holder to choose the exercise time. In that respect, the value of an American option is the one obtained by exercising optimally. To determine this optimal exercise policy, one has to solve an optimal stopping problem which makes pricing via simulation that more difficult. However, aside from the substantial computational effort, Monte Carlo is a well-suited tool for pricing (path dependent) American options. In this chapter we implement two methods that will make use of its principles and provide reasonably tight bounds on the option price.

The value process of an American option can in general be characterized through dynamic programming (see [2]), a recursive algorithm that breaks down the optimal stopping problem into smaller parts and using value iteration and backward induction provides the option price. However, problems of higher dimension (like the ones of pricing Asian options) make the use of this technique impractical, as the “curse of dimensionality” implies that even the partial iterations are hard to evaluate explicitly. In this case we cannot hope to determine the optimal stopping rule exactly, but maybe we can still make use of the iterative procedure. Longstaff and Schwartz [21] develop an approximate dynamic programming (ADP) algorithm that follows the same reasoning but employs least squares regression to avoid any direct computations. The recursions ultimately produce an estimate to the option price, and as the results correspond to a feasible exercise strategy, the eventual effect of the approach is a lower bound.

While ADP methods works surprisingly well on realistic high dimensional problems, they fail to provide any information on the optimality of its estimate. Hence, we consider an approach developed by Rogers [27] and Haugh and Kogan [15] and use it to evaluate ADP solutions. They implement a method based on the dual representation of the optimal stopping problem and instead of maximizing over a set of stopping times, minimize over a class of (super) martingales. Now given an arbitrary price approximation, we can use this dual representation to construct an upper bound on the true value function. Finally, with a computed upper and lower bound, one can determine how far the approximate solution is from optimality.

Our approach can be summarized in the following steps:

1. Using an ADP algorithm in connection with cross-path regression compute an approximation of the option price as a function of time and state.

2. Estimate the lower bound on the option price by replicating the exercise strategy implicitly generated in step 1.

3. Based on the option price approximation, define a martingale process and use it to estimate the upper bound by Monte Carlo simulation.

The following results are derived in a Markovian environment under the usual assumption of market completeness. All the expectations are taken under the risk neutral measure $\mathbb{Q}$, but for simplicity, we omit pointing out such technicalities in the notation. The reliance on a specific measure is in fact not essential to dynamic programming. As one will see, it will only be implicit in the choice of the discount factor.

4.1 Problem formulation

A continuous-time American option is specified by a process $(Z_t)_{0\leq t\leq T}$ describing the discounted exercise value, and a class of admissible stopping times $\mathcal{T}$ with values in $[0,T]$. The problem is that of finding the optimal time-zero value of the option

$$V_0 \coloneqq \sup_{\tau\in\mathcal{T}} \mathbb{E}(Z_\tau)$$

Under certain regularity conditions, it is justified calling this expression the arbitrage price of the American claim. However, if we want to characterize the option value through a dynamic programming relation, we first need to slightly reframe this assertion.

Start by noting that the reward process Z is usually derived from more primitive elements. Per assumption, all relevant financial variables (like the price of the underlying assets) are embedded in an $\mathbb{R}^d$-valued Markov process $\{S_t, 0 \leq t \leq T\}$. Since the model in question (Black-Scholes-Merton) consists of a single risky asset, we only pursue the case where $d = 1$. A generalization is fairly easy to obtain and can be found in [11]. Now for some nonnegative reward function g and payoff upon exercise at time t denoted by $g(S_t)$, the above pricing problem becomes the one of computing

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}(e^{-r\tau} g(S_\tau))$$

Note how the structure of the discount factor $e^{-r\tau}$ (where r is the riskless rate of lending and borrowing) supports the assumption of valuation under the risk-neutral measure $\mathbb{Q}$.

Since we cannot exactly replicate the continuity of time in a simulation, we have to impose some further restraints. Recall that American options can be exercised at any point in a continuum of dates up to maturity and that the decision to exercise can be made at the very moment based on the information available. However, when discussing pricing via simulation, another notion comes in handy. Bermudan options, much like American, can also be exercised prior to maturity though only at a fixed, predetermined set of dates $t_1 < t_2 < \cdots < t_n$. Such a restriction is often part of the options contract, since it is sometimes known in advance that exercise is suboptimal on all but specific dates, like dividend payment dates. Alternatively, the restriction can be interpreted as a discrete approximation of a contract allowing continuous exercise. In this case one would want to consider many exercise opportunities and the effect of letting n increase to infinity. To introduce the necessary concepts, we focus solely on the finite-exercise problem. In other words, when speaking of pricing American options via simulation, we are tacitly referring to Bermudan options.

Consistent with prior chapters, we denote the state of the underlying Markov process at time t_j by S_j. In the context of Bermudan options, the fixed set of time points $\{t_1, t_2, \dots, t_n\}$ at which we simulate assume the role of exercise opportunities. To avoid further complications, we assume that each value S_{j+1} can be simulated directly from the value S_j at the earlier exercise date. That is, without the need to use a smaller time step than the intervals $\Delta t = t_{j+1} - t_j$ between the exercise dates. The simulation of the Markov chain $\{S_1, S_2, \dots, S_n\}$ in this case is exact, and the simulated values aren't subject to discretization error. This finite-exercise Markovian formulation will ultimately allow us to characterize the option value through dynamic programming.

4.1.1 Dynamic programming

The trick with dynamic programming is to somehow break down the optimal stopping problem into a reasonable number of subproblems. Solving these partial problems and combining the results will build up larger solutions and ultimately provide the final answer to our question. Given a sample path $\{S_1, \dots, S_n\}$ let us analyze possible payoffs received by the option holder at each time step $\{t_1, \dots, t_n\}$. Clearly, if the option is not exercised before maturity, the holder receives the terminal payoff $g(S_n)$. At any other exercise opportunity t_j, the holder of an American claim can decide whether to exercise the option or wait for the next exercise opportunity. The realized cashflow based on his decision at time t_j is

$$Y_j = \begin{cases} g(S_j)\,, & exercise \\ D_{j,j+1}Y_{j+1}\,, & continue \end{cases}$$

Which is to say that the holder receives the payoff $g(S_j)$ if his decision is to exercise. Otherwise, the life of the option is continued, and he receives a cash flow of Y_{j+1} at the succeeding time step. The present value of this future payment is obtained by multiplying with a "one step" discount factor $D_{j,j+1}$ from t_j to t_{j+1}. In a risk-neutral world we assume it is of the form

$$D_{j,j+1} = \exp\left(-\int_{t_j}^{t_{j+1}} r ds\right) = e^{-r(t_{j+1}-t_j)} = e^{-r\Delta t}\,, \qquad j = 0, \dots, n-1$$

where r is the continuously compounded interest rate.

However, one should not make the mistake of basing the optimal exercise decision on a simple comparison of the realized cash flows $g(S_j)$ and $D_{j,j+1}Y_{j+1}$ at each time step. The key aspect is to consider only the information available up to that time, and since there is no way of knowing the payoff Y_{j+1} at the previous time step, this would be the equivalent of looking into the future. The correct approach instead is to compare the immediate exercise cash flow $g(S_j)$ with the expected discounted future cash flow, contingent on the available information. Iterating this argument, we get the following recursive scheme for the value of the option at each time step

$$V_n = g(S_n)$$

$$V_{j-1} = \max\left\{ g(S_{j-1})\,,\ \mathbb{E}\left[D_{j-1,j}V_j \,\middle|\, \mathcal{F}_{t_j}\right] \right\}, \qquad j = 1, \dots, n$$

Note that this is essentially the expression of the Snell Envelope in discrete time. However, the effectiveness of this approach is more transparent in a Markovian setting. Using the Markov property of the underlying process, we can rewrite the above formulae as functions of time and space. Assuming the option has not previously been exercised, let $V_j(s)$ be the value at time t_j given the state $S_j = s$. We allow the payoff function to depend on time and denote it by g_j for exercise at t_j. Determining the option price (the time zero-value $V_0(S_0)$), involves computing the value function iteratively by working backwards in time. An approach also known as dynamic programming

$$V_n(s) = g_n(s)$$

$$V_{j-1}(s) = \max\{\ g_{j-1}(s)\ ,\ \mathbb{E}\big[D_{j-1,j}V_j(S_j)\ \big|\ S_{j-1} = s\big]\ \}\ , \qquad j = 1, \dots, n$$

This representation nicely reflects the fact that the values V_j depends on the asset values at time t_j. Since there is usually no exercise right in the time period $[0, t_1)$, we exclude the current time 0 from the exercise opportunities by setting $g_0 \equiv 0$. At each time step the holder of an American option optimally compares the payoff from immediate exercise with the expected payoff from continuation, and then exercises if the immediate payoff is higher. Thus, the optimal exercise strategy is fundamentally determined by the conditional expectation of the payoff from continuing to keep the option alive. Determining this value will be the main difficulty of the approach.

Note that the iterated values $V_j(S_j)$ are presented in time-t_j euros. To make a comparison with payoffs at time t_{j-1}, every iteration requires discounting $V_j(S_j)$ one time-step back. While this is useful for practical implementation, in discussing pricing algorithms it would be more convenient to work with a simpler notation. We achieve this by denominating all the values at the same point in time, thereby suppressing the explicit discounting at each time step.

The value of the option reported in time-zero euros is defined by

$$\tilde{g}_j(s) = D_{0,j}g_j(s)\ , \qquad j = 1, \dots, n$$

and

$$\tilde{V}_j(s) = D_{0,j}V_j(s)\ , \qquad j = 0, \dots, n$$

It is clear that $\tilde{V}_0 = V_0$, and since the discount factor obviously satisfies $D_{0,j-1}D_{j-1,j} = D_{0,j}$, we can derive the following simplified expressions

$$\tilde{V}_n(s) = \tilde{g}_n(s)$$

$$\begin{aligned}
\tilde{V}_{j-1}(s) &= D_{0,j-1}V_{j-1}(s) \\
&= D_{0,j-1}\, max\{\ g_{j-1}(s)\ ,\ \mathbb{E}\big[D_{j-1,j}V_j(S_j)\ \big|\ S_{j-1} = s\big]\ \} \\
&= max\{\ \tilde{g}_{j-1}(s)\ ,\ \mathbb{E}\big[D_{0,j-1}D_{j-1,j}V_j(S_j)\ \big|\ S_{j-1} = s\big]\ \} \\
&= max\{\ \tilde{g}_{j-1}(s)\ ,\ \mathbb{E}\big[D_{0,j}V_j(S_j)\ \big|\ S_{j-1} = s\big]\ \} \\
&= max\{\ \tilde{g}_{j-1}(s)\ ,\ \mathbb{E}\big[\tilde{V}_j(S_j)\ \big|\ S_{j-1} = s\big]\ \}
\end{aligned}$$

The discounted values obviously support the same dynamic programming recursions. Continuing in this way, we absorb the discount factors into the definitions of S and g and assume all values are denominated in time 0-euros. This formulation will be used to establish the main results in the upcoming chapter. Further on, when describing the algorithms, we will include discounting explicitly.

4.1.2 Continuation values and Q-value iteration

Conditional on the option not being previously exercised, the expected value of continuing $C_j^E(s)$ in state s at time t_j (measured in time-0 dollars) is given by

$$C_j^E(s) = \mathbb{E}\left[\tilde{V}_{j+1}(S_{j+1}) \mid S_j = s\right], \qquad j = 0, \dots, n-1$$

Prevailing that the current state of the price process equals s, it is the value resulting from holding, rather than exercising the American option until the next exercise possibility.

Let us now verify that the continuation value satisfies the dynamic programming recursion, which further clarifies the backward nature of the problem. We know the continuation value at maturity $C_n^E(x)$ equals zero, since there is no continuing after that point. Furthermore, we can express the expected continuation value at time t_j in terms of the expected and realized continuation value in the following way

$$C_j^E(s) = \mathbb{E}\left[\max\left\{\tilde{g}_{j+1}(S_{j+1}),\ C_{j+1}^E(S_{j+1})\right\} \mid S_j = s\right], \qquad j = 0, \dots n-1$$

Such formulation suggests interpreting $C_j^E(s)$ as a newly issued option at time t_j starting from the state s. Iterating this argument backward in time infers that the option value equals $C_0^E(S_0)$, the time zero value of continuing.

The continuation values are obviously defined by means of value functions. The converse also holds, i.e. the value functions V_j are determined by C_j^E via the following relation

$$\tilde{V}_j(s) = \max\left\{\tilde{g}_j(s), C_j^E(s)\right\}, \qquad j = 0, \dots, n$$

The above discussion justifies using continuation values as an alternative to value iteration, in some literature also referred to as Q-value iteration.

4.1.3 Stopping rules

Dynamic programming recursions focus on option values, but it will also be convenient to view the pricing problem through stopping rules and exercise regions. The optimal strategy for the option holder is to exercise the option at an optimal stopping time τ^*. However, stopping according to any other (not necessarily optimal) exercise policy also determines a value through

$$V_0^{(\tau)}(S_0) = \mathbb{E}[\tilde{g}_\tau(S_\tau)]$$

The generated output is usually suboptimal and can be improved by maximizing over a set of stopping times (which is the essential problem related to American claims). One may think of applying the following rule

$$\tau^+ = \inf\,\{\tau \mid \tilde{g}_\tau(S_\tau) \geq \tilde{g}_\sigma(S_\sigma), \qquad \sigma, \tau \in [1, T]\}$$

However, this would require using more information than that available at the time, since the decision is based on a comparison of exact payoffs at each time step. τ^+ is hence not a stopping time or can it be used as an exercise rule.

Still, restricting τ to be an ordinary stopping time would require that each event $\{\tau = t_j\}$ be determined by $S_1, S_2, \dots, S_j$. In allowing randomized stopping times we are allowing such an event to be dependent also on other random variables independent of $S_{j+1}, \dots, S_n$. This extension is needed to accommodate stopping rules by simulation.

In this regard, any policy assigning a value $\hat{V}_j(s)$ to each state s and time t_j, with $\hat{V}_n = \tilde{g}_n$ specifies a stopping rule

$$\hat{\tau} = \min\{\, j \in \{1, \dots, n\} \mid \tilde{g}_j(S_j) \geq \hat{V}_j(S_j)\,\}$$

If the approximations additionally satisfy

$$\hat{V}_j(s) = \max\{\,\tilde{g}_j(s), \hat{C}_j(s)\}\,, \qquad j = 1, \dots, n$$

the rule is essentially the same as the following defined by $\hat{C}_j^E$

$$\hat{\tau} = \min\{\, j \in \{1, \dots, n\} \mid \tilde{g}_j(S_j) \geq \hat{C}_j(S_j)\,\}$$

Which will be used to derive the lower bound on the option price. For a single path, it is determined as the payoff resulting from stopping the first time the immediate exercise value is as great as the estimated continuation value

$$\hat{V}_0 = \tilde{g}_{\hat{\tau}}(S_{\hat{\tau}}) = D_{0,\hat{\tau}}\, g_{\hat{\tau}}(S_{\hat{\tau}})$$

The exercise region determined by $\hat{C}_j$ at the j-th exercise date is the set

$$\{\, s : \tilde{g}_j(s) \geq \hat{C}_j(s)\,\}$$

and the continuation region, the complement of this set. The stopping rule $\hat{\tau}$ can thus also be described as the first time the Markov chain S_j enters an exercise region. The value determined by using $\hat{\tau}$ does not in general coincide with $\widehat{V_0}$, though the two would be equal if we started with the optimal value function. In the following section we review a numerical method to approximate the iterations and the optimal stopping time.

4.2 ADP and Linear Least Squares Regression

In a dynamic programming algorithm, even though the value of exercising is normally easy to determine, both value iteration and Q-value iteration is not feasible in practice and will have to be approximated in some way. Luckily, in recent years efficient ADP algorithms have been developed to address this issue. Longstaff and Schwartz [21] exploit the fact that the continuation values can be described by a linear combination of basis functions. They rely on the ability to simulate paths of the underlying process and use linear regression to transfer the cross-sectional information onto such set of functions. The fitted value from the regression provides a direct estimate of the conditional expectation function. By estimating the conditional expectations at each time step, we obtain a complete specification of the optimal exercise strategy along each path and with-it American options can be valued accurately by simulation.

The Least Squares Monte Carlo (LSM) approach is both fast and precise, but more importantly, it helps assess path-dependent American options. The effect of the approach is to reduce the pricing problem to one of selecting an adequate set of basis functions. The choice of these basis functions is what determines how close to the true value one can get with increased computational effort.

4.2.1 Example – step by step

The LSM approach is best illustrated on a numerical example. The one here specified is taken from Longstaff and Schwartz and will be used to evaluate a fixed strike Asian put. A floating strike Asian or a plain vanilla claim can be handled in a similar way.

Consider a three-year American-Asian put option on a share that can be exercised at the end of year 1, year 2 and year 3. The current stock price is 1.00 and the fixed strike price is 1.10. The risk-free rate is 6% per annum (continuously compounded). For simplicity, we illustrate the algorithm using only eight sample paths (generated in a risk-neutral world). The sample paths are laid out in the following table

Table 4.1 - Stock price paths

Path	$t = 0$	$t = 1$	$t = 2$	$t = 3$
1	*1.05*	*1.09*	*1.08*	*1.34*
2	*1.07*	*1.16*	*1.26*	*1.54*
3	*1.02*	*1.22*	*1.07*	*1.03*
4	*0.99*	*0.93*	*0.97*	*0.92*
5	*1.01*	*1.11*	*1.56*	*1.52*
6	*0.91*	*0.76*	*0.77*	*0.90*
7	*1.00*	*0.92*	*0.84*	*1.01*
8	*0.95*	*0.88*	*1.22*	*1.34*

Dictated by the dynamic programming procedure, the approximate LSM method is also recursive in nature. Working backwards from maturity, we maximize the value of the option one step at a time by comparing the current value of stopping with some estimated (based on simulated paths) value of waiting. To make this work, we first need to compute a number of intermediate matrices.

Conditional on not exercising the option before maturity, the realized cash flows at time $t = 3$ are calculated according to the expression $(K - A_3)^+$and are given in Table 4.2.

The last column shows the average stock prices along the entire path. For example, the average stock price from $t = 0$ to $t = 3$ of Path 7 is 0.94, calculated according to the specified stock prices in Table 4.1

$$\frac{1.00 + 0.92 + 0.84 + 1.01}{4} = 0.9425 \approx 0.94$$

Table 4.2 - Cash flow matrix at time 3

Path	$t = 1$	$t = 2$	$t = 3$	K	A_3
1	-	-	*0.00*	*1.10*	*1.14*
2	-	-	*0.00*	*1.10*	*1.26*
3	-	-	*0.01*	*1.10*	*1.09*
4	-	-	*0.05*	*1.10*	*0.95*
5	-	-	*0.00*	*1.10*	*1.30*
6	-	-	*0.26*	*1.10*	*0.84*
7	-	-	*0.16*	*1.10*	*0.94*
8	-	-	*0.00*	*1.10*	*1.10*

These cash flows, since calculated at expiry, are identical to those received if the option were European rather than American. If the option is in-the-money at the previous time step ($t = 2$), the option holder must decide whether to exercise immediately or continue the options life until the final expiration date. Of course, if the option is out of the money it would make no sense to exercise it.

Table 4.3 – Intrinsic values at time 2

Path	$t = 1$	$t = 2$	K	A_2
1	-	*0.02*	*1.10*	*1.08*
2	-	*0.00*	*1.10*	*1.16*
3	-	*0.00*	*1.10*	*1.10*
4	-	*0.14*	*1.10*	*0.96*
5	-	*0.00*	*1.10*	*1.23*
6	-	*0.29*	*1.10*	*0.81*
7	-	*0.18*	*1.10*	*0.92*
8	-	*0.08*	*1.10*	*1.02*

From Table 4.3 we see that there are five in-the-money paths. Path 1, 4, 6, 7 and 8, have a positive intrinsic value at time step $t = 2$. Since we want to estimate the value of continuing, it is only natural to consider the paths where the exercise decision is relevant.

Now let X denote the vector of stock prices at time $t = 2$ for these five paths and Y the corresponding discounted cash flows received at time $t = 3$ if the option is not exercised. The corresponding values are given in the table entries below.

Table 4.4 - Regression at time 2

Path	$Y = g(S_3)e^{-r\Delta t}$	$X = S_2$
1	*0.00*0.94176=0.00*	*1.08*
2	-	-
3	-	-
4	*0.05*0.94176=0.04*	*0.97*
5	-	-
6	*0.26*0.94176=0.24*	*0.77*
7	*0.16*0.94176=0.15*	*0.84*
8	*0.00*0.94176=0.00*	*1.22*

We use least squares regression to estimate $\mathbb{E}[Y|X]$, the expected cash flow from continuing to keep the option alive (conditional on the stock prices at time $t = 2$). Specifically, we regress Y on (X^0, X^1, X^2), a second-degree polynomial based on the stock prices at time $t = 2$. Later on we will show how the approximation can be improved. The resulting estimate of the conditional expectation function is

$$\mathbb{E}[Y|X] = 1.39 - 2.1X + 0.78X^2$$

The algorithm compares the value of immediate exercise at time $t = 2$, given in the first column below, with the value from continuation, given in the second.

Table 4.5 - Optimal early exercise decision at time 2

Path	*Immediate exercise at time 2*	*Expected continuation value*	*Decision*
1	*0.02*	*0.03*	*Hold*
2	-	-	-
3	-	-	-
4	*0.14*	*0.09*	*Exercise*
5	-	-	-
6	*0.29*	*0.24*	*Exercise*
7	*0.18*	*0.17*	*Exercise*
8	*0.08*	*–0.01*	*Exercise*

The immediate exercise value is just the intrinsic value at time $t = 2$ for in-the-money paths. The continuation value is obtained by substituting X in the conditional expectation function. For a specified path, it is calculated by plugging in the corresponding stock price at time $t = 2$.

Take for example the first path. The payoff from exercising at time $t = 2$ is 0.02, while the expected payoff of continuing is based on the approximate expectation function and equals $1.39 - 2.1 \cdot 1.08 + 0.78 \cdot 1.08^2 \approx 0.03$. Thus, the optimal decision of the option holder wanting to maximize his value at each time step, is to hold the option and wait for a better chance to exercise it. A likewise comparison in the other 4 in-the-money paths suggests that the option should be exercised. Following is the revised cash-flow matrix at time $t = 2$.

Table 4.6 - Cash flow matrix at time 2

Path	$t = 1$	$t = 2$	$t = 3$
1	-	*0.00*	*0.00*
2	-	*0.00*	*0.00*
3	-	*0.00*	*0.01*
4	-	*0.14*	*0.00*
5	-	*0.00*	*0.00*
6	-	*0.29*	*0.00*
7	-	*0.18*	*0.00*
8	-	*0.08*	*0.00*

Note that, if the option is exercised at time $t = 2$, the cash flow in the final column becomes zero. This is because the options here considered can only be exercised once, and if the decision to exercise is made, there are no further cash flows.

We proceed with an analogue procedure at the earlier exercise date. According to Table 4.7 below, we have 6 in-the-money paths at time $t = 1$.

Table 4.7 - Intrinsic values at time 1

Path	$t = 1$	K	A_1
1	*0.03*	*1.10*	*1.07*
2	*0.00*	*1.10*	*1.12*
3	*0.00*	*1.10*	*1.12*
4	*0.14*	*1.10*	*0.96*
5	*0.04*	*1.10*	*1.06*
6	*0.27*	*1.10*	*0.83*
7	*0.14*	*1.10*	*0.96*
8	*0.19*	*1.10*	*0.91*

For paths 1,4,5,6,7,8 we define X in the same way, as the stock price vector at time $t = 1$, and Y as the discounted value of subsequent cash flows. Note that we use the actual realized cash flows at time $t = 2$, and not the estimated expected cash flow from continuing.

As already stated, the option can only be exercised once, so future cash flows occur either at time $t = 2$ or $t = 3$, but not at both. To obtain their current value, cash flows received at time $t = 2$ are discounted back one period, and cash flows received at time $t = 3$ are discounted back two periods to time $t = 1$.

The vectors X and Y are given in the columns below.

Table 4.8 - Regression at time 1

Path	$Y = g(S_2)e^{-r\Delta t}$	$X = S_1$
1	*0.02*0.94176=0.01*	*1.09*
2	-	-
3	-	-
4	*0.14*0.94176=0.13*	*0.93*
5	*0.00*0.94176=0.00*	*1.11*
6	*0.29*0.94176=0.27*	*0.76*
7	*0.18*0.94176=0.16*	*0.92*
8	*0.08*0.94176=0.07*	*0.88*

Once again, we regress Y on a constant, X and X^2. The linear regression approximates the conditional expectation function to be

$$\mathbb{E}[Y|X] = 1.43 - 2.08X + 0.71X^2$$

Substituting values of X into the conditional expectation function we get the estimated continuation values at time $t = 1$ which are then compared to immediate exercise values in the table below.

Table 4.9 - Optimal early exercise decision at time 1

Path	*Immediate exercise at time 1*	*Expected continuation value*	*Decision*
1	*0.03*	*0.01*	*Exercise*
2	-	-	-
3	-	-	-
4	*0.14*	*0.11*	*Exercise*
5	*0.04*	*0.00*	*Exercise*
6	*0.27*	*0.26*	*Exercise*
7	*0.14*	*0.12*	*Exercise*
8	*0.19*	*0.15*	*Exercise*

A comparison of the columns indicates that the optimal decision for all in-the-money paths is to exercise, which concludes our analysis.

Having identified the exercise strategy at all the relevant times $t = 1,2,3$, we have implicitly defined a stopping rule. If ones denote the dates at which the option is exercised, the rule can be described in the following way

Table 4.10 – "Optimal" exercise strategy (Stopping rule)

Path	$t = 1$	$t = 2$	$t = 3$
1	*1*	*0*	*0*
2	*0*	*0*	*0*
3	*0*	*0*	*1*
4	*1*	*0*	*0*
5	*1*	*0*	*0*
6	*1*	*0*	*0*
7	*1*	*0*	*0*
8	*1*	*0*	*0*

The computation of the cash flows for each path is now straightforward. We simply follow the specified stopping rule and exercise the option at indicated dates in the matrix. The obtained values are identified in table 4.11

Table 4.11 - Option cash flow matrix

Path	$t = 1$	$t = 2$	$t = 3$
1	*0.03*	*0.00*	*0.00*
2	*0.00*	*0.00*	*0.00*
3	*0.00*	*0.00*	*0.01*
4	*0.14*	*0.00*	*0.00*
5	*0.04*	*0.00*	*0.00*
6	*0.27*	*0.00*	*0.00*
7	*0.14*	*0.00*	*0.00*
8	*0.19*	*0.00*	*0.00*

The approximate price for a single path is obtained by discounting the corresponding cash flow in the above matrix back to time zero (one, two or three steps back, depending on the specified exercise time of a path). Finally, the Monte Carlo estimate of the option price is obtained by averaging these over all the paths.

American $$0.03 * e^{-r} + 0.00 + 0.01 * e^{-3r} + 0.14 * e^{-r} + 0.04 * e^{-r} + 0.27 * e^{-r} + 0.14 * e^{-r} + 0.19 * e^{-r} = \frac{0.7644822611}{8} = 0.0955602826$$

This yields a value of 0.095 for the American Asian put option, which is almost twice as much as the corresponding European option obtained simply by discounting back the cash flows at maturity.

European $$e^{-3r} * (0.01 + 0.05 + 0.26 + 0.16) = \frac{0.4009297015}{8} = 0.0501162127$$

4.2.2 Estimating the continuation values

As previously indicated, Q-value iteration can be used as an alternative to value iteration. Rather than approximating value functions directly, the following ADP algorithm will focus on approximating the Q-value functions (continuation values). As we now explain, there are a number of reasons to why this is preferable than a direct approximation.

The defining equations for approximate Q-value and value iterations can respectively be written as

$$\hat{C}_j^E(s) = \mathbb{E}\left[D_{j,j+1} \max\left\{ g_{j+1}\left(S_{j+1}\right),\ \hat{C}_{j+1}^E\left(S_{j+1}\right) \right\} \middle| S_j = s\right]$$

$$\hat{V}_j(s) = max\left\{ g_j\left(S_j\right),\ \mathbb{E}\left[D_{j,j+1}\hat{V}_{j+1}\left(S_{j+1}\right) \middle| S_j = s\right] \right\}$$

The above functional forms suggest that the continuation values are smoother and as such easier to approximate. More importantly, since the holder of an American option has to choose at each time step whether to exercise or continue to hold the option, the decision will be based on a comparison of the realized payoff and the continuation value. The unknown quantity of interest is therefore the continuation value and if we only have $\hat{V}_j$ available, such a comparison will be hard to make. For example, when $\hat{V}_j$ is only slightly greater than the value of immediate exercise one can make the mistake of assuming it is optimal to continue, when it is in fact optimal to exercise. This is especially the case when there are relatively few exercise opportunities. Having a direct approximation of C_j^E avoids these problems and still yields an estimate to the price.

Each continuation value $C_j^E(s)$ is the regression of $V_{j+1}\left(S_{j+1}\right)$ on the current state s. We refer to it as "regression now". The name will become clear eventually when introducing the upper bound. The suggested approximation procedure is the following: approximate the continuation value by a linear combination of known functions of the current state and use (least squares) regression to estimate the best coefficients in this approximation.

For the moment suppose that the conditional expectation function is an element of L^2. As a Hilbert space, L^2 has a countable orthonormal basis and we can represent the conditional expectation as a linear function of its elements. For an underlying Markovian process S, a common choice of basis functions would be orthogonal polynomials like power or (weighted) Laguerre suggested by Longstaff and Schwartz. In any case, by assuming certain technical conditions we can represent the continuation values C_j^E at time t_j as a linear combination of a countable set of $\mathcal{F}_j$-measurable basis functions $\{\psi_k\}_{k\in\mathbb{N}}$ of the underlying state. We have the following representation (denominated in time-0 euros)

$$C_j^E(s) = \mathbb{E}\left[\tilde{V}_{j+1}\left(S_{j+1}\right) \middle| S_j = s\right] = \sum_{l=0}^{\infty} \beta_{jl}\, \psi_{jl}\left(S_j\right)$$

To implement the LSM approach, we first approximate the continuation value by truncating the above expression using the first $K < \infty$ basis functions. This finite term representation $\overline{C_j^E}$ is in general not exact and generates an error ϵ. In particular, at time t_j we have

$$C_j^E(s) = \sum_{k=1}^{K} \beta_{jk}\psi_{jk}(s) + \epsilon = \overline{C_j^E} + \epsilon$$

We call ψ_{jk} basis functions and constants β_{jk} regression coefficients. Our goal now is to determine the coefficients β_{jk} by trying to fit the model to data, while minimizing the error ϵ in the least squares sense.

To that end, we rewrite the approximation $\overline{C_j^E}$ in vector form

$$\overline{C_j^E} = \beta_j^\top \psi(s)$$

with $\beta_j^\top = (\beta_{j1}, \beta_{j2}, \dots, \beta_{jK}),\qquad \psi_j(s) = \left(\psi_{j1}(s), \psi_{j2}(s), \dots, \psi_{jK}(s)\right)^\top$

The corresponding square error is defined as

$$SE_j = \mathbb{E}\left[\left(C_j^E(S_j) - \overline{C_j^E}(S_j)\right)^2\right]$$

$$= \mathbb{E}\left[\left(\mathbb{E}[\tilde{V}_{j+1}(S_{j+1}) |\, S_j] - \beta_j^\top \psi_j(S_j)\right)^2\right]$$

To minimize it, we differentiate with respect to β_{jk} for $k = 1, \dots, K$ and set the derivative equal to zero

$$\mathbb{E}\left[\left(\mathbb{E}[\tilde{V}_{j+1}(S_{j+1}) |\, S_j] - \overline{C_j^E}(S_j)\right)\psi_j(S_j)\right] = 0$$

Rewriting the above, we obtain the following expression

$$\mathbb{E}[\tilde{V}_{j+1}(S_{j+1})\psi_j(S_j)] = \mathbb{E}\left[\overline{C_j^E}(S_j)\psi_j(S_j)\right]$$

$$= \mathbb{E}\left[\psi_j(S_j)\psi_j(S_j)^\top\right]\beta_j$$

Where, using simple matrix algebra, we can easily disclose the regression vector

$$\beta_j = \left(\mathbb{E}\left[\psi_j(S_j)\psi_j(S_j)^\top\right]\right)^{-1} \mathbb{E}[\psi_j(S_j)\tilde{V}_{j+1}(S_{j+1})]$$

Having specified the subset of basis functions and the form of the regression vector minimizing the squared error, we integrate the approach in our simulation scheme. In particular, we use sample values to estimate the defining expectations in the regression vector.

The described procedure assumes that payoffs and value estimates are denominated in time-0 euros, which are in practice usually denominated in current-time euros. For implementational purposes, all quantities will be recorded in current time. Making the switch will require including explicit discounting. The present value of continuing is thus refined as

$$C_j^E = \mathbb{E}\left[D_{j,j+1} V_{j+1}\left(S_{j+1}\right) \mid S_j \right]$$

It will also be convenient to use the following notation

$$B_\psi = \mathbb{E}\left[\psi_j\left(S_j\right)\psi_j\left(S_j\right)^\top\right]$$

$$B_{V\psi} = \mathbb{E}\left[\psi_j\left(S_j\right)D_{j,j+1}V_{j+1}\left(S_{j+1}\right)\right]$$

where B_ψ is the above indicated $K \times K$ matrix (assumed nonsingular) and $B_{V\psi}$ the indicated vector of length K denominated in time-t_j euros. Recall that in a risk neutral world values are discounted at the risk-free rate r, so the one-step discount factor $D_{j,j+1}$ is of the form $e^{-r\Delta t}$.

The variables $\left(S_j, S_{j+1}\right)$ in the above expectations have the joint distribution of the state of the underlying Markov chain at dates t_j and t_{j+1}. Note that for Markovian problems, only current values of the state variables are necessary and we won't need to include any past realizations in the basis functions or regression. In particular, having simulated m independent paths $(S_{i1}, S_{i2}, \dots, S_{in})$, $i = 1, \dots, m$, the coefficients β_{jk} could be estimated from observations of pairs $\left(S_{ij}, V_{j+1}\left(S_{i,j+1}\right)\right)$, $i = 1, \dots, m$ each consisting of the state at time t_j and the corresponding option value at time t_{j+1}. Suppose for a moment that the values $V_{j+1}\left(S_{i,j+1}\right)$ are known beforehand. The least squares estimate of β_j is then given by replacing the expectations with their sample counterparts

$$\hat{\beta}_j = \hat{B}_\psi^{-1}\hat{B}_{V\psi}$$

In particular, $\hat{B}_\psi \approx \mathbb{E}\left[\psi_j\left(S_j\right)\psi_j\left(S_j\right)^\top\right]$ is the $K \times K$ matrix with qr entry given by

$$\frac{1}{m}\sum_{i=1}^{m}\psi_{jq}\left(S_{ij}\right)\psi_{jr}\left(S_{ij}\right)$$

And $\hat{B}_{V\psi} \approx \mathbb{E}\left[\psi_j\left(S_j\right)D_{j,j+1}V_{j+1}(S_{j+1})\right]$ is the K-vector with r-th entry

$$\frac{1}{m}\sum_{i=1}^{m}\psi_{jr}\left(S_{ij}\right)D_{j,j+1}V_{j+1}\left(S_{i,j+1}\right)$$

All of the above quantities can be calculated from function values at pairs of consecutive states $(S_{ij}, S_{i,j+1})$, $i = 1, \dots, m$. However, V_{j+1} is in practice unknown and must be replaced by estimated values $\hat{V}_{j+1}$ at downstream nodes. The exact procedure will be described in the following section.

As there is no continuation at maturity, we have $\hat{C}_n^E \equiv 0$. Starting with the initialization $\hat{V}_n = g_n(S_n)$ and working backwards in time, we suppose that the value estimates $\hat{V}_{j+1}, \hat{V}_{j+2}, \dots, \hat{V}_n$ have already been determined. This allows us to calculate the regression vector $\hat{\beta}_j$ which in turn defines the approximation continuation value $\hat{C}_j^E(s)$ at time t_j in state s through

$$\hat{C}_j^E(s) = \hat{\beta}_j^\top \psi_j(s)$$

It can be shown that the fitted value of the regression $\hat{C}_j^E$ converges in mean square and in probability to the true value C_j^E as the number of paths increases.

4.2.3 Exercise decision

Once we have estimated the continuation value at time t_j, for an in-the-money path, we can determine whether it is optimal to exercise the option or continue to hold it. The decision is based on the comparison of the estimated continuation value $\hat{C}_j^E$ and the immediate exercise value $g(S_j)$. We will have to work backwards in time, since the intermediate cash flows generated by the option are defined recursively. Specific implementations can still vary. Tsitsiklis and Van Roy [30] construct a high estimate by maximizing between the immediate exercise value and the continuation value. In order to construct the low estimate, Longstaff and Schwartz dissociate the exercise decision from the reward collected in the event of exercise. Broadie and Glasserman [6] further suggest combining the two techniques. Once the trajectories have been simulated, using backward induction at each node we can derive two estimates of the price. One biased high, the other biased low. Even though they arrive at the conclusion by a different argument and using a stochastic mesh, if the weights are produced by a least squares procedure, the methods coincide. However, we will not pursue this idea any further, and instead propose an alternative algorithm for the derivation of the upper bond. By imposing some restrictions on the regression, it too can be used in conjunction with the method of Longstaff and Schwartz. The details are given in subsequent chapters.

Assume the option holder follows the optimal exercise strategy at all times and denote by $\hat{Y}_j$ the path of cash flows generated by the option, conditional on it not being exercised prior to time t_j. The notation is motivated by the theoretical formula, and for a fixed path, the function is an analogue representation of a single column (corresponding to a time step) of the intermediate cash flow matrices derived in the numerical example of the previous section.

4.2.3.1 Interleaving estimator

Having identified the optimal decision, we can assign a value at node S_j in the following way

$$\hat{Y}_j = \begin{cases} g_j(S_j)\,, & g_j(S_j) \geq \hat{C}_j^E(S_j)\,, \quad exercise \\ D_{j,j+1}\hat{Y}_{j+1}\,, & g_j(S_j) < \hat{C}_j^E(S_j)\,, \quad hold \end{cases}$$

If the immediate exercise value $g_j(S_j)$ is at least as great as the estimated continuation value $\hat{C}_j^E$, we simply set $\hat{Y}_j = g_j(S_j)$. Since there should be no further cashflows after the decision to exercise is made, this additionally requires changing all subsequent flows along the realized path to zero.

However, if it is optimal to hold the option, instead of assigning the high value $\hat{C}_j^E$ to the current node, we consider the cashflows evolving the current path into the future. In particular, starting from the present node S_j we record the payoff received by exercising according to the following stopping rule

$$\hat{\tau}_j = min\{l \in \{j, \ldots, n\} : g_l(S_l) \geq \hat{C}_l^E(S_l)\}$$

The obtained payoff is then discounted to current time t_j and assigned as the value at $\hat{Y}_j$. The two cases for assigning a value to $\hat{Y}_j$ can now be combined into the rule

$$\hat{Y}_j = D_{j,\hat{\tau}_j} h_{\hat{\tau}_j}(S_{\hat{\tau}_j})$$

4.2.3.2 In the money paths

Longstaff and Schwartz additionally recommend omitting states S_{ij} with $g_j(S_{ij}) = 0$ in estimating the regression coefficients β_j. At each time step, instead of projecting on all m simulated paths, they only use in-the-money paths in the regression. This makes sense since the exercise decision is only relevant when the option is in the money. By limiting the region over which the conditional expectation must be estimated, far fewer basis functions are needed to obtain an accurate approximation. Specifically, more than three times as many functions to accomplish the same level of accuracy as the estimator based only on in-the-money paths. This not only reduces bias, but significantly improves the efficiency of the algorithm.

4.2.3.3 Separating regression from pricing

To further improve the efficiency of the algorithm we can separate the regression (where we estimate the coefficients β) from pricing. Instead of generating m number of trajectories once and for all, during each iteration we simulate independent sets of paths for the purpose of pricing (assigning a value at $\hat{Y}_j$). In particular, each time the optimal decision is to continue, we generate further values of the underlying Markov chain $\{\check{S}_j, \check{S}_{j+1}, \dots, \check{S}_n\}$ starting from the current state $\check{S}_j = S_j$. This is imposed to ensure that the newly simulated path really extends the original path. We then proceed with the same procedure described above. More precisely, we define a new stopping time and apply it to the path starting from $\check{S}_j$

$$\check{\tau}_j = min\{\sigma \in \{j, j+1, \dots, n\} : h_l(\check{S}_l) \geq \hat{C}_l^E(\check{S}_l)\}$$

Note that this requires once more approximating the continuation $\hat{C}_l^E(\check{S}_l)$ for values $\check{S}_l$, and repeating the procedure each subsequent time step until the optimal decision is to exercise. This can be done in a "forward manner" since the regression coefficients $\hat{\beta}_l$, $l = j, \dots, n$ have already been determined.

An alternative implementation would be to first simulate a small number of regression paths to estimate the coefficients $\hat{\beta}_j$ at each time step t_j. We then go ahead by simulating a large set of pricing paths whose values we use to compute the lower bound. Holding the coefficients fixed, the continuation values are easily determined. Starting at t_1 we work forward in time, comparing the immediate exercise value with the approximated value of continuing and exercise as soon as the optimal decision is to stop. In the worst case, we evaluate the option at the final exercise date. The obtained payoff is then discounted to initial time and assigned as the estimate of the optimal value for the specified path.

4.2.3.4 Convergence results

With the number of paths increasing to infinity, it can be shown (see (Longstaff & Schwartz, 2001)) that the interleaving estimator of Longstaff and Schwartz converges to any desired degree of accuracy. For it to actually coincide with the true price $V_0(S_0)$, the representation of the conditional expectation through a set of basis functions

$$\hat{C}_j^E(s) = \sum_{k=1}^{K} \beta_{jk}\psi_{jk}(s)$$

would have to hold exactly at each time step $j = 1, \dots, n-1$. However, that is in general not the case, and as we now explain, the algorithm will produce low biased estimates.

The LSM algorithm provides an exercise policy that is a pathwise approximation of the optimal stopping rule maximizing the option value. Since no policy is better than the optimal exercise policy, following that, or any other (suboptimal) exercise strategy results in values that underestimate the true price. On the other hand, applying backward induction over a finite set of paths results in high bias. Instead of keeping the two sources of bias separate, the described procedure interleaves elements of the high and low estimators, alternating between the two by applying backward induction to estimate a continuation value and then applying a suboptimal stopping rule starting from each node. This has the potential to produce a more accurate value, but even when using an independent set of paths for the regression, the results are still priced low. The reason for this is that the high bias resulting from Jensen's inequality tends to be less pronounced than the low bias resulting from a suboptimal stopping rule.

The low biasedness can be particularly useful as it provides an objective convergence criterion. For example, to determine the number of basis functions needed for an accurate approximation, we can simply increase N until the value implied by the algorithm no longer increases. Note that this property is not shared by other regression-based methods.

4.2.3.5 Variations

An algorithm developed around the same time by Tsitsiklis and Van Roy [30] proposes the following alternative. First apply backward induction to estimate the continuation values and then use the stopping rule implicitly defined by the procedure, thus keeping the sources of bias separate. To assign a value at each node S_j set

$$\hat{Y}_j = \max\{g(S_j), \hat{C}_j^E(S_j)\}$$

The intermediate cashflows are obtained by taking the maximum of the immediate exercise value and the estimated expected value of continuing. While Longstaff and Schwartz use the actual realized cash flows in defining $\hat{Y}_j$ along each path, here we use the conditional expected value estimated at time t_{j+1} to assign a value at time t_j. Due to the convexity of the maximum operator, discounting back this conditional expectation can result in an upward bias in the value of the option. The resulting estimate is indeed biased high because it either overestimates the true continuation value (in which case $\hat{Y}_j$ is biased high regardless if it is optimal to exercise or not) or it underestimates it, in which case it is replaced by a larger quantity, the exercise payoff.

Tsitsiklis and Van Roy prove convergence of $\hat{V}_0$ to V_0 as the number of paths increases to infinity $m \longrightarrow \infty$. An interesting observation brought forth by Glasserman [11] is that this particular method actually corresponds to using a specific set of weights in a stochastic mesh, a method of Broadie and Glasserman [6] proven to be very effective in pricing high dimensional American options.

4.2.4 Least Squares Monte Carlo scheme

While specific implementations can vary, the scheme described in this section will follow the procedure developed by Longstaff and Schwartz. Note that path dependence of option payoffs does not pose any difficulties to the simulation based LSM algorithm. Cashflows of American Asian options for example additionally depend on the stock price over a certain averaging window. By introducing the average to date as a second state variable, we can transform the path dependent problem to a Markovian one and follow the same steps to derive an estimate to the option price.

1. Simulate m independent paths of the Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$, $i = 1, \dots, m$ in a risk neutral world
2. At maturity set $\hat{V}_{in} = g_n(S_{in})$ for each computed path
3. Apply backward induction until $\hat{V}_{i1}$ is reached for every path

   For $j = n - 1, \dots, 1$ do the following

   a. Account for in-the-money paths at the current time step t_j
   b. Calculate the continuation values $\hat{C}_j^E(s)$ for each in-the-money path by regressing $\hat{V}_{j+1}(S_{j+1})$ on the current state s
      i. Given estimated values $\hat{V}_{j+1}$ from the previous time step, calculate the regression coefficients $\hat{\beta}_j = \hat{B}_\psi^{-1}\hat{B}_{V\psi}$ using only in-the-money paths
      ii. Estimate the continuation value from $\hat{C}_j^E(S_j) = \hat{\beta}_j^\top \psi(S_j)$
   c. Following an optimal exercise decision for each path $i = 1, \dots, m$ set

$$\hat{V}_{ij} = \begin{cases} g_j(S_{ij}), & g_j(S_{ij}) \geq \hat{C}_{ij}^E(S_{ij}), \quad exercise \\ D_{j,j+1}\hat{V}_{i,j+1}, & g_j(S_{ij}) < \hat{C}_{ij}^E(S_{ij}), \quad hold \end{cases}$$

4. Discount the obtained values for each path one step back to initial time, and then calculate the average to obtain a lower bound on the option price

$$\underline{\hat{V}_0} = \frac{1}{m} D_{0,1}\left[\hat{V}_{11} + \hat{V}_{21} + \cdots + \hat{V}_{m1}\right]$$

Unlike the pricing scheme for European options, the regression part of the algorithm requires storing all m simulated paths at once, and this is only for plain vanilla options. If the problem were that of evaluating path dependent options, the computational effort further intensifies. Consider for example, a fixed strike American Asian option. To compute the intrinsic value at each time step, we need more than just the average of prices at maturity. In fact, the algorithm will require a complete specification of the running average for each path and exercise date. Better yet, for a floating strike Asian, both a matrix of spot values and running averages will have to be available during the course of the algorithm, which further clarifies the high dimensionality of the problem.

4.2.4.2 Minimizing an error

The computation of all steps in our scheme seems more less straightforward. However, having simulated paths, it is not completely transparent how the regression coefficients are to be determined. In this section we describe the implementational part of the least-squares regression. It will require projecting the discounted (future) cash flows onto the basis functions of the current state (for the paths where the option is in the money).

Consider an m-dimensional vector Y and an $m \times K$ matrix X. We assume there is linear dependence between the two, which we specify by multiplying X with a K-dimensional vector B. Through the choice of the appropriate linear coefficients, we aim to express Y in terms of X, but since this is in general not exact, the representation is subject to an error. Denoting the m-dimensional error vector by $\mathcal{E}$ we have the following

$$Y = XB + \mathcal{E}$$

We want to reduce the incurred error in the least squares sense. Since all the other data is otherwise determined, we do this by varying the coefficients in the regression. To that end, let SE denote the squared error incurred by choosing a specific B. We have

$$\begin{aligned} SE(B) &= min\, \mathcal{E}^{\top}\mathcal{E} \\ &= min(Y - XB)^{\top}(Y - XB) \\ &= min(Y^{\top}Y + B^{\top}X^{\top}XB - 2B^{\top}X^{\top}Y) \end{aligned}$$

The latter representation implies that the squared error will be minimal if the regression coefficients have the following form

$$B = (X^{\top}X)^{-1}X^{\top}Y$$

In our specific problem, for each time step $j = n - 1, \dots, 1$, the above vectors and matrices assume the following roles

K – truncation factor

m – number of simulated paths

Y – $m \times 1$ column vector of intermediate cashflows

B – $K \times 1$ vector of regression coefficients

X – $m \times K$ matrix of basis functions (polynomials of the relevant variables)

The approximate expected continuation value is determined as

$$C = XB$$

Since the algorithm of Longstaff and Schwartz includes only in-the-money paths in the regression, the dimension of matrices at each time step need not be the same. Instead of considering all m simulated paths, we would first have to account for the ones where the intrinsic value of the option at the given time step is non-negative, and then replace m with the obtained number.

4.3 Analysis of the method and improvements

It usually takes 10 000 to 50 000 paths to obtain a precise estimate of the option value using the LSM algorithm. The main idea is to employ regression in approximating the continuation values, and since this technique is usually very fast in practice, the resulting iterations are also very fast. Another advantage is that the method is well suited for parallel programming. For example, we could generate 5 000 paths on a single computer, or we could generate 100 paths on 50 computers. This can be useful in many large-scale applications like risk management where the computational speed is far more important than the cost of hardware. Still, there are further opportunities for simulation or regression techniques to significantly improve the efficiency of ADP algorithms. To identify them, we begin by analyzing the errors originating from the approach.

1. Discretization error

   The error resulting from restricting the exercise opportunities to a finite set of dates. Due to the complexity of the problem, we have focused our attention solely on the case of Bermudan options, as a discrete approximation of American, and worked under the assumption that the underlying Markov chain can be simulated exactly. This means that the joint distribution of the simulated values coincides with the joint distribution of the continuous-time process on the simulation time grid. Still, we can accommodate continuously exercisable American options by considering a sufficiently large number of exercise dates, but since exact simulation is in reality infeasible, some discretization error is usually inevitable.

2. Approximation error of the continuation value

The error arises in consequence of using a finite number of basis functions to estimate the expected continuation value. A better estimate results in the smaller norm of the standard error vector SE, so any regression-based approach clearly depends on the choice of basis functions. Through Taylor expansion, we can use polynomials to approximate sufficiently smooth value functions, which makes them a popular choice even with Longstaff and Schwartz. They recommend analyzing the payoff and figuring out which types of functions actually lead to a good fit.

3. Stochastic error

Just like in the European option setting, the quality of the simulation can be improved by the use of variance reduction techniques. In addition to antithetic variates, a control variate can be used to "control the error" of the Monte Carlo estimate. For the options of American type, Rasmussen [25] suggests not only controlling the value estimate, but the additional use of a control for the approximating continuation values. Extending the approach requires a few steps. First it is necessary to sample the control variate whenever the American option is exercised. This of course adds no extra cost if the control variate is already used to improve the Monte Carlo valuation. Second, we need the known conditional expectation of the control for each observation where the option is in the money. The last step is to include four rather than one dependent variables in the least squares regression. The computational cost is therefore considered fixed and, in his article, Rasmussen proves the method improves the efficiency of the least squares projection without changing the bias.

4. Foresight bias

Another important error is known as the foresight bias. For example, when we constructed the stopping rule, we used information not fitting the information content reflected by the underlying filtration $\mathcal{F}_t$. Actually, we used more information than available at a given exercise opportunity. This stopping rule leads to a high foresight biased estimate of the price. First, we could have based the option pricing on the same paths used for regression. But since the option value is based on the regression values it depends on the realized continuation values of the paths. To keep this dependence as small as possible we have chosen to use another set of paths, the pricing paths, to compute the option value. This approach also has practical reasons since the error arising from the pricing is much bigger than that from determining the exercise rule. To this end, fewer paths are used for regression than are used for pricing.

Note that increasing the number of paths only reduces variance. What improves the estimate is the selection of basis functions. The remainder of the chapter is dedicated to a more detailed discussion on the possible refinements.

4.3.1 Improving the quality of the regression

Consider the error from applying an approximation of the expected continuation value using regression. The only point where the accuracy matters is the point where the option gets in-the-money, which is around the strike level. Here, the accuracy directly influences the decision of exercising. At all other points the accuracy does not matter, and that makes the considered error significantly smaller than the one arising from pricing. Still, if the approximation of the exercise boundary, which is done by computing the expected continuation values, is slightly biased, the construction of the stopping time is slightly wrong. It is slightly wrong because in the neighborhood of the true exercise boundary the decision of holding the option or exercising leads to very similar results, but far away a wrong decision leads to very different results. Not exercising a deep in-the-money option leads to losing a lot of money but not exercising a nearly at-the-money option does not.

So, when accuracy is of importance, we seek to improve the quality of the estimate either by considering a different polynomial family for the basis functions, or by a better regression method or algorithm.

In the procedures described, we use ordinary least squares regression to estimate the conditional expectation function. However, in some cases a more efficient approach would be to use weighted least squares, generalized least squares or even a generalized method of moments. As this is not particularly relevant to the options considered in the paper, we no longer pursue such alternatives and instead turn our attention to basis functions.

4.3.1.1 Selecting basis functions

Several functions have been suggested as basis functions. Longstaff and Schwartz argue for the use of orthogonal polynomials such as (weighted) Laguerre, Hermite, Jacobi, or Legendre. Another popular choice is to use $1, x, x^2$ ... and variants including the payoff itself as a basis function. In any case, the choice of basis functions is not universal and may require experimentation or good information about the structure of the problem. Luckily, one existing study will apply particularly well to our situation. In the following section we specify the results brought forth by de Lima and Samanez [7] who investigate the efficiency of the Least-Squares Monte Carlo approach by using different polynomial basis. Since the main focus of their paper are Asian options, this will give us a good intuition on which basis functions to use in our approximations.

De Lima and Semanez asses the leas squares performance in pricing both fixed and floating strike Asian options of either geometric or arithmetic average. The polynomials considered in the study are Power, Legendre, Laguerre and Hermite A. The format used to write a polynomial is chosen in such a way that the pricing procedure becomes operationally more practical. They suggest combining the following three expressions with terms specified in the succeeding tables

a. Explicit form

$$f_n = d_n \sum_{m=0}^{N} c_m g_m(x)$$

Table 4.0.12 – Explicit expressions of the basis functions

	$f_n(x)$	N	d_n	c_m	$g_m(x)$
Power	$W_n(x)$	*0*	*1*	*1*	x^n
Legendre	$P_n(x)$	$\frac{n}{2}$	2^{-n}	$(-1)^m \binom{n}{m}\binom{2n-2m}{n}$	x^{n-2m}
Laguerre	$L_n(x)$	n	1	$\frac{(-1)^m}{m!}\binom{n}{n-m}$	x^m
Hermite A	$H_n(x)$	$\frac{n}{2}$	$n!$	$(-1)^m \frac{(1)}{m!\,(n-2m)!}$	$(2x)^{n-2m}$

b. Rodrigues' formula

$$f_n(x) = \frac{1}{a_n \rho(x)} \frac{\partial^n}{\partial x^n} [\rho(x) g(x)^n]$$

Table 4.0.13 – Coefficients using the Rodrigues' formula

	$f_n(x)$	N	d_n	c_m	$g(x)$
Power	$W_n(x)$	0	$\frac{(2n)!}{n!}$	x^{2n}	1
Legendre	$P_n(x)$	$\frac{n}{2}$	$(-1)^n 2^n n!$	1	$1-x^2$
Laguerre	$L_n(x)$	n	$n!$	e^{-x}	x
Hermite A	$H_n(x)$	$\frac{n}{2}$	$(-1)^n$	e^{-x^2}	1

c. Recurrence law

$$a_{n+1}f_{n+1}(x) = (a_n + b_n x)f_n(x) - a_{n-1}f_{n-1}(x)$$

Table 4.14 – Specific terms of the recurrence law

	$f_n(x)$	a_{n+1}	a_n	b_n	a_{n-1}	$f_0(x)$	$f_1(x)$
Power	$W_n(x)$	1	0	1	0	1	x
Legendre	$P_n(x)$	$n+1$	0	$2n+1$	n	1	x
Laguerre	$L_n(x)$	$n+1$	$2n+1$	-1	n	1	$1-x$
Hermite A	$H_n(x)$	1	0	2	$2n$	1	$2x$

A mixture of above implementations has proven to be very effective in their study. Particularly, they suggest using the explicit form for Power and Laguerre polynomials, Rodrigues' formula for Legendre, and the recurrence law to express the Hermite A polynomial.

Using the above specifications, they consider three different sets of parameters to determine which choice of basis functions will be the best fit. The underlying asset is modelled as a geometric Brownian motion with $S_0 = 100$, $r = 0.05$, $\sigma = 0.2$, $q = 0$. For a fixed option type, an optimal choice of the polynomial can vary depending on the choice of strike. In the case of an arithmetic Asian option, the optimal choice for the polynomial basis is presented in the table below.

Strike	Fixed Strike Call	Floating Strike Call	Fixed Strike Put	Floating Strike Put
95	Power	Power	Legendre	Power
100	Power	Power	Hermite A	Hermite A
105	Legendre	Legendre	Legendre	Hermite A

Further analysis of Asian options of geometric average implies certain homogeneity between the two. Namely, the results suggest that the pricing can be done by selecting a single polynomial basis for all types of American Asian options and that the main concern in fact is not with the type of the option, but whether we are pricing a put or a call. In the case of Asian American call option, a preferable choice is Power polynomial basis. For an Asian American put, the recommended Hermite A polynomials provide a better performance.

However, the difference between the efficiency of a particular choice becomes negligible when a greater number of simulated trajectories is used. This leads to the conclusion that, in the case of orthogonal polynomial basis, all of the considered polynomials lead to approximately same results and thus present a good fit for pricing.

4.3.1.2 State variables

When choosing the underlying state variables of the basis functions, one should consider which information influences the possibility of exercise. For instance, at time step t_j, we can approximate the conditional expectation function of a plain vanilla option by two-variable quadratic polynomials, taken on a basis of spot values. For simplicity we assume there is only one path ($m = 1$). The regression coefficients $B = (\beta_0, \beta_1, \beta_2)^\top$ are then taken as the least squares estimates to the regression line

$$\hat{C}_j^E(s) = \begin{pmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{pmatrix} \times \begin{pmatrix} S_j^0 & S_j^1 & S_j^2 \end{pmatrix} = \beta_0 + \beta_1 S_j + \beta_2 S_j^2$$

The regression equation should obviously take all relevant variables into account. When there are two state variables X and Y, the set of basis functions should include terms in both X and Y, as well as the cross-product of the two. So, when pricing average options, it is necessary to consider both the S_t and the average value A_t of the option. We can include the prefix arithmetic average into the regression equation by modifying the above equations in the following way

$$B = (\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6)^\top$$

$$X = \left(1, S_j, S_j^2, S_j A_j, A_j, A_j^2\right)$$

$$\hat{C}_j^E(s) = \beta_1 + \beta_2 S_j + \beta_3 S_j^2 + \beta_4 S_j A_j + \beta_5 A_j + \beta_6 A_j^2$$

This seems to suggest that the number of basis functions needed to obtain convergence grows exponentially with the dimensionality of the problem. In reality, however, the number grows at a much slower rate, and may be very manageable for high-dimensional problems.

4.3.1.3 Number of basis function

When applying regression, each basis function adds another degree of freedom in the form of a corresponding regression coefficient. This additional degree of freedom might increase the efficiency of the regression and in turn improve the approximation of the exercise boundary. None the less, an excessive increase in the number of basis functions can in some cases reduce the efficiency, making the procedure computationally expensive. According to Longstaff and Schwartz, the structure used in regressions, comprises a constant, the two first degrees of the chosen polynomial basis and their crossed product up to the third degree, totaling eight basis. Yet, the truncation number additionally depends on the dimension of the underlying state vector, and in the tested examples it actually ranges from 5 to 20 basis functions. Since the resulting estimate to the option price is low-biased, a criterion specific to Longstaff and Schwartz's algorithm is to increase the number of basis functions until the value implied by the algorithm no longer increases.

The next section provides details for the computation of the upper bound on the option price. The procedure will be used in addition to the one described so far.

4.4 The Dual Approach to Optimal Stopping Problems

The regression approach we have outlined tends to underprice American-style options since the value is computed using an approximation to the optimal stopping rule. To determine how far the approximate solution is from optimality, we introduce a dual-based procedure developed by Rogers [27] and Haugh and Kogan [15]. It can be used in conjunction with any algorithm that generates a lower bound and pinpoints the true value of an American-style option more precisely than the algorithm does by itself.

In contrast to ADP methods, the approach makes no attempt to determine an approximately optimal exercise policy, and always comes up with an answer that is an upper bound for the true price. Though it says little about how the option should be exercised, it does give guidance on how the option should be hedged. Thus, it should be of value to the party writing the option.

Haugh and Kogan suggest using low discrepancy sequences in computing the upper bound. Because they are evenly dispersed, this usually result in a very fast rate of convergence. However, the bound can also be computed using Monte Carlo simulation. This is made feasible by the representation of the American option price as a solution of a properly defined dual minimization problem. In particular, the price of the American option may be expressed as the infimum of a family of expectations, the infimum being taken over the class of (super) martingales. The expression immediately suggests how one might try to estimate the price of an American option. To state the theoretical results of the paper we again suppress explicit discounting and assume that g_t and V_t are in fact the discounted payoff and discounted value functions denominated in time-zero euros. We address the details of discounting when implementing the algorithm.

4.4.1 Dual formulation

The problem of pricing an American option, the primal problem, is that of computing

$$V_0 = \sup_{\tau \in \mathcal{T}} \mathbb{E}[g_\tau(S_\tau)]$$

Under certain regularity conditions, the associated Snell envelope process (discounted value process)

$$V_t = \operatorname*{ess\,sup}_{\tau \in \mathcal{T}_{[t,T]}} \mathbb{E}\left[g_\tau(S_\tau) | \mathcal{F}_t\right]$$

Is a supermartingale, and as such has the following Doob-Meyer decomposition

$$V_t = V_0^* + M_t^* - A_t^*$$

where M is a martingale vanishing at zero, and A a predictable integrable increasing process, also vanishing at zero. For an arbitrary adapted supermartingale M_t, the value of an American option satisfies

$$\begin{aligned} V_0 = \sup_{\tau \in \mathcal{T}} \mathbb{E}[g_\tau] &= \sup_{\tau \in \mathcal{T}} \mathbb{E}[g_\tau - M_\tau + M_\tau] \\ &\leq \sup_{\tau \in \mathcal{T}} \mathbb{E}[g_\tau - M_\tau] + \sup_{\tau \in \mathcal{T}} \mathbb{E}[M_\tau] \\ &\leq \sup_{\tau \in \mathcal{T}} \mathbb{E}[g_\tau - M_\tau] + M_0 \\ &\leq \mathbb{E}\left[\max_{t \in \mathcal{T}} (g_t - M_t)\right] + M_0 \end{aligned}$$

where the second inequality follows from the optional sampling theorem for supermartingales. Taking the infimum over all supermartingales M_t proves that V_0 is bounded above in the following way

$$V_0 \leq U_0 \coloneqq \inf_M \mathbb{E}\left[\max_{t \in \mathcal{T}} (g_t - M_t)\right] + M_0$$

On the other hand, it is a known fact that the discounted value process is itself a supermartingale. In fact, it is the smallest supermartingale that dominates the discounted payoff of the option at all exercise periods i.e. $V_t \geq g_t$ for all t, which implies

$$U_0 \leq \inf_M \mathbb{E}\left[\max_{t\in\mathcal{T}}(g_t - V_t)\right] + V_0$$

Therefore, the primal and the dual problem coincide, i.e. $V_0 = U_0$.

The theorem suggests that an upper bound on the price of the American option can be constructed simply by evaluating the dual function over an arbitrary supermartingale M_t. In particular, if such a supermartingale satisfies $M_t \geq g_t$, the option price V_0 is bounded from above by M_0. When the supermartingale M_t coincides with the discounted option value process V_t, the upper bound equals the true price of the American option. This suggests that a tight upper bound can be obtained by using an accurate approximation $\hat{V}_t$ to define M_t, in the sense that when the approximate option price $\hat{V}_t$ coincides with the exact price V_t, M_t equals the discounted option price V_t.

4.4.2 Upper bound on the option price

The proposed method now consists of picking a suitable supermartingale M, and evaluating the expectation $\mathbb{E}\left[\max_{t\in\mathcal{T}}(g_t - M_t)\right]$ using simulation. Later on we show how to find a "suitable" martingale which is of course a task of similar complexity to finding the optimal exercise policy, but we still often find simple martingales that provide remarkably good (and quick) bounds. Haugh and Kogan then suggest the use of the following definition of M_t

$$\hat{M}_0 = \hat{V}_0$$

$$\hat{M}_{t+1} = \hat{M}_t + \hat{V}_{t+1} - \hat{V}_t - \max\left(\mathbb{E}\left[\hat{V}_{t+1} - \hat{V}_t \middle| \mathcal{F}_t\right], 0\right)$$

By construction, $\mathbb{E}[M_{t+1} - M_t|\mathcal{F}_t] \leq 0$, so the defined process is indeed an adapted supermartingale for any approximation $\hat{V}_t$. Furthermore, when the approximation $\hat{V}_t$ equals the true value process V_t, we have that $M_t = V_t$ because the process is a supermartingale, and the positive part of the above expectation equals zero.

Note also that the upper bound can be tightened further by omitting the positive part in the definition of M_t. The resulting process is a martingale (the martingale part of $\hat{M}_t$) and therefore also a supermartingale, so that it too can be used to construct an upper bound. It coincides with $\hat{M}_t$ at $t = 0$ and is always greater than or equal to $\hat{M}_t$ for $t > 0$. Therefore, it leads to a lower value of the upper bound. So, by using the martingale component of a supermartingale we can in general produce tighter upper bounds. For the remainder of the chapter we thus define M_t as

$$M_0 = \hat{V}_0$$

$$M_{t+1} = M_t + \hat{V}_{t+1} - \hat{V}_t - \mathbb{E}\left[\hat{V}_{t+1} - \hat{V}_t \middle| \mathcal{F}_t\right]$$

Now let $\overline{V}_0$ denote the upper bound we get by using the above defined (super) martingale M_t. It is easy to see that the bound is explicitly given by the following expression

$$\overline{V}_0 = \hat{V}_0 + \mathbb{E}\left[\max_{t\in\mathcal{T}}\left(g_t - \hat{V}_t + \sum_{k=1}^{t} \mathbb{E}\left[\hat{V}_k - \hat{V}_{k-1} \middle| \mathcal{F}_{k-1}\right]\right)\right]$$

Haugh and Kogan relate the worst case performance of the above denoted upper bound to the accuracy of the original approximation $\hat{V}_t$. The approximation error $\left|\hat{V}_t - V_t\right|$ is under certain conditions and for certain ADP algorithms bounded above by a constat $\sqrt{n}$ where n is the number of exercise periods. This suggests that the quality of the upper bound should deteriorate with n, however, not in a linear fashion. In their paper, Haugh and Kogan provide further evidence to support this claim by pricing options with as many as 100 exercise periods.

4.4.3 Computing the upper and lower bounds

The derivation of the upper bound $\overline{V}_0$ assumes that an approximation to the price $\hat{V}_0$ is available. We can thus combine the described dual method with the regression based approach developed in previous chapters by setting $\hat{V}_0 = \underline{V}_0$, where $\underline{V}_0$ is the estimated lower bound using ADP algorithms. Because $\underline{V}_0 \geq g_0$, this new definition satisfies $\hat{V}_t \geq g_t$ for all t, and the upper bound in this case is given by

$$\overline{V}_0 = \underline{V}_0 + \mathbb{E}\left[\max_{t\in\mathcal{T}}\left(g_t - \hat{V}_t + \sum_{k=1}^{t} \mathbb{E}\left[\hat{V}_k - \hat{V}_{k-1} \middle| \mathcal{F}_{k-1}\right]\right)\right]$$

Which gives a natural decomposition of the upper bound into a sum of two components. The first component is the estimated lower bound, while the second component in some sense measures the extent to which the discounted approximate value function fails to be a supermartingale.

We estimate $\overline{V}_0$ by simulating m sample paths of the state variables $\{S_{i1}, S_{i2}, \dots, S_{in}\}$, and evaluating

$$\max_{t_j}\left(g_j(S_{ij}) - \hat{V}_j(S_{ij}) + \sum_{k=1}^{j} \mathbb{E}\left[\hat{V}_k(S_{ik})\middle|S_{i,k-1}\right] - \mathbb{E}\left[\hat{V}_{k-1}(S_{i,k-1})\middle|S_{i,k-1}\right]\right)$$

along each path $i = 1,2 \dots m$ and taking the average over all paths. This can be rather time consuming because we also need to accurately estimate the conditional expectation

$$\mathbb{E}\left[\hat{V}_k(S_{ik}) - \hat{V}_{k-1}(S_{i,k-1}) \mid S_{i,k-1}\right]$$

at each time period along each simulated path, which means using Monte Carlo simulations within the simulation.

Alternatively, if the initial value function approximation $\underline{V}_0$ comes from an ADP algorithm (like the one described in the previous section) a method by Glasserman and Yu [13] suggests that it might be possible to choose the basis functions in such a way that the above conditional expectations can be compute analytically. In that case, the need for conducting nested simulations would not arise. The procedure is closely related to a technique called "regression later" which we now describe.

4.4.4 Regression now vs. Regression later

In this section we contrast methods that, at time t_j regress option values from time t_{j+1} against basis function values at time t_j (regression now) with the methods that regress against basis function values at time t_{j+1} (regression later). To better contrast the two methods, we briefly summarize the necessary tools discussed in previous chapters.

Unknown quantity of interest (to be estimated) $\quad V_0(S_0) = \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[D_{0,\tau}\, g_\tau(S_\tau)\right]$

Dynamic programming equations

$$V_n(s) = g_n(s)$$

$$V_j(s) = \max\left\{ g_j(s)\,,\ \mathbb{E}\left[D_{j,j+1}V_{j+1}(S_{j+1}) \mid S_j = s\right] \right\}, \qquad j = 0,\dots,n$$

In terms of continuation value $\quad C_j(s) = \mathbb{E}\left[D_{j,j+1}V_{j+1}(S_{j+1}) \mid S_j = s\right]$

$$C_n(s) = 0$$

$$C_j(s) = \mathbb{E}\left[D_{j,j+1}\, max\left\{ g_{j+1}(S_{j+1})\,,\ C_{j+1}(S_{j+1}) \right\} \mid S_j = s\right], \qquad j = 1,\dots,n-1$$

Option values satisfy $V_j(s) = \max\left\{ g_j(s), C_j(s) \right\}$, $j = 0, \dots, n$

Approximate Dynamic Programming

Our procedure comprises of the following steps

1. Approximate the conditional expectations by least squares regression

Regression now	Regression later
$\overline{C}_j = \beta_j^\top \psi_j$	$\overline{C}_j^+ = \gamma^\top \psi_j$

2. Approximate the projections by fitting the model to data

Regression now	Regression later
$\hat{C}_j = \hat{\beta}_j^\top \psi_j$	$\hat{C}_j^+ = \hat{\gamma}_j{}^\top \psi_j$

The ADP method estimates value or Q-value iterations by projecting onto a linear combination of basis functions ψ. The important point is that the basis variables have finite variance, are linearly independent and uncorrelated. When using "regression later" we will additionally require that the basis functions satisfy the martingale property.

The finite term representation does not need to be exact and the generated error arising from "regression now" and "regression later" is respectively given by

Regression now $\quad \epsilon_{j+1} = V_{j+1}(S_{j+1}) - \sum_{k=1}^{K} \psi_{jk}(S_j)\beta_{jk}$

Regression later $\quad \epsilon_{j+1}^+ = V_{j+1}^+(S_{j+1}) - \sum_{k=1}^{K} \psi_{j+1,k}(S_{j+1})\gamma_{jk}$

If however, at each time step the residual errors in t_{j+1} are zero conditional on the current state in t_j, the ADP relations would hold exactly.

If this does not apply, we can still provide an approximation through simulation. By simulating m independent paths of the underlying Markov chain, we approximate the Q-value iterations by a sample projection through ordinary least-squares regression.

When the basis functions are martingales, and the residuals are uncorrelated with the martingale differences

$$\mathbb{E}\left[\, \epsilon_{j+1}^+ \left(\psi_{j+1,k}(S_{j+1}) - \psi_{jk}(S_j)\right) \right] = 0\,, \qquad \forall j = 0, \dots, n-1, k = 1, \dots, K$$

it can be shown that the sample projections $\hat{C}_j^+$ coincide with the projections $\overline{C}_j^+$ in the ADP algorithm, and so the solution satisfies the ADP problem exactly.

We now explain the difference between "regression now" and "regression later" in detail.

Regression now – projecting on the current state

$[0, T]$ option lifetime divided into equidistant time steps $0 < t_1 < t_2 < \cdots < t_n = T$, $\Delta t = \frac{T}{n}$ representing exercise opportunities and dates at which we simulate the underlying process

m number of simulated paths $\quad$ n number of time steps $\quad$ K size of truncation

Independent replications of the Markov chain $(S_{i1}, S_{i2}, \dots, S_{in})$, for $i = 1,2, \dots, m$

Basis functions $\quad \psi_j = (\psi_{j1}, \psi_{j2}, \dots, \psi_{jK})$, $\quad j = 1,2 \dots, n-1$

Regression coefficients $\quad \hat{\beta}_j^\top = (\hat{\beta}_{j1}, \hat{\beta}_{j2}, \dots, \hat{\beta}_{jK})$, $\quad j = 1,2, \dots, n-1$

$$\hat{\beta}_j = \left(\sum_{i=1}^{m} \psi_j(S_{ij})^\top \psi_j(S_{ij}) \right)^{-1} \left(\sum_{i=1}^{m} \psi_j(S_{ij})^\top D_{j,j+1} \hat{V}_{j+1}(S_{i,j+1}) \right)$$

Approximate continuation $\quad \hat{C}_n(s) = 0$

$$\hat{C}_j(s) = \sum_{k=1}^{K} \hat{\beta}_{jk} \psi_{jk}(s) , \qquad j = 1,2, \dots, n-1$$

Value iterations $\quad \hat{V}_{j+1} = \max\{g_{j+1}, \hat{C}_{j+1}\}$, $\quad j = 0,1, \dots, n$

Estimate $\quad \hat{C}_0(S_0) = \frac{1}{m} \sum_1^m D_{0,1} \hat{V}_{i1} \qquad \hat{V}_0 = \max\{g_0(S_0), \hat{C}_0(S_0)\}$

If we do not allow exercise at the initial time $t_0 = 0$, the value is obtained simply by discounting the continuation $\hat{C}_0(S_0)$ value to the initial state.

Regression later – projecting on the later state

Each continuation value $\hat{C}_j(s)$ in "regression now" is the projection of $\hat{V}_{j+1}(S_{j+1})$ on the current state s, with the corresponding coefficients $\hat{\beta}_j$ defined by regressing against $\psi_j(S_{ij})$. The coefficients $\hat{\gamma}_j$ in "regression later" project on $\psi_{j+1}(S_{i,j+1})$ instead. In other words, "regression now" uses current basis functions, while "regression later" uses later basis functions. The crucial point when using "regression later" is that the basis functions have to be martingales. We thus impose the following condition

Martingale property $\quad \mathbb{E}[\, \psi_{j+1}(S_{j+1}) \mid S_j \,] = \psi_j(S_j)$, $\quad j = 0, \dots, n-1$

This implies that $\hat{C}_j$ and $\hat{C}_j^+$ are in fact linear combinations of the same basis functions. They only differ in the estimates of the coefficients they use.

Basis functions $\quad \psi_j = (\psi_{j1}, \psi_{j2}, \dots, \psi_{jK}), \quad j = 1,2 \dots, n-1$

Regression coefficients $\quad \hat{\gamma}_j^\top = (\hat{\gamma}_{j1}, \hat{\gamma}_{j2}, \dots, \hat{\gamma}_{jK}), \quad j = 1,2, \dots, n-1$

$$\hat{\gamma}_j = \left(\sum_{i=1}^{m} \psi_{j+1}(S_{i,j+1})^\top \psi_{j+1}(S_{i,j+1})\right)^{-1} \left(\sum_{i=1}^{m} \psi_{j+1}(S_{i,j+1})^\top \hat{V}_{j+1}^+(S_{i,j+1})\right)$$

Approximate continuation $\quad \hat{C}_n^+(s) = 0$

$$\hat{C}_j^+(s) = \sum_{k=1}^{K} \hat{\gamma}_{jk} \psi_{jk}(s), \quad j = 1,2, \dots, n-1$$

Value iterations $\quad \hat{V}_{j+1}^+ = \max\{g_{j+1}, \hat{C}_{j+1}^+\}, \quad j = 0,1, \dots, n$

Estimate $\quad \hat{C}_0^+(S_0) = \frac{1}{m} \sum_1^m \hat{V}_{i1} \qquad \hat{V}_0^+ = \max\{g_0(S_0), \hat{C}_0^+(S_0)\}$

Note that the discounting of the estimated continuation value is done implicitly using the martingale condition on the basis functions.

Intuitively, we expect "regression later" to give better results than "regression now" because the option values at time t_{j+1} should be more highly correlated with the basis functions at time t_{j+1} than with the basis functions at time t_j. And while $\hat{C}_j$ uses simulation to estimate both the conditional expectations and their projections, $\hat{C}_j^+$ takes advantage of the martingale property to compute the conditional expectations exactly, using simulation only to estimate the projections. Glasserman and Yu [13] provide a formal proof to support these observations, but only for a single-period problem, where the "regression later" yields a better fit and less variable estimates of coefficients than "regression now".

4.4.5 Choice of a suitable martingale basis

We want to find basis functions that satisfy the martingale property. The first thing that comes to mind is a simple constant, as it obviously satisfies the martingale condition.

- $\psi(x) = 1$

To specify other basis functions, we take a look at the model at hand. When the underlying process is a geometric Brownian motion, a set of martingale basis functions can be obtained from the following equation (see [9])

- $$\psi_{j,k}(x) = x^{k-1} exp\left(-\left((k-1)(r-q) + (k-1)(k-2)\frac{\sigma^2}{2}\right)\right)(t_j - t_0)$$

Note that the first derived martingale basis function, the constant, is obtained when $k = 1$.

Many studies have also shown that including the payoff as a basis function improves the accuracy of regression methods. The martingale part of the corresponding European option is thus a common choice. A general way of constructing martingales from payoff functions that does not require solving the full option pricing problem is taking

$$\psi_{jk}(S_j) = e^{-r(t_k - t_j)} \mathbb{E}[g(S_k)|S_j]$$

When it comes to pricing exotic options, the choice of a suitable martingale basis becomes more involving. For example, the discounted exercise value of an Asian option has no martingale part, which means we cannot simply follow the recipe and take the martingale part of the payoff. To illustrate the complexity of the procedure, and the martingale basis used, considers a weighted average Asian call. In his paper, Rogers [27] examines the problem at hand and concludes a good choice would be the following three Lagrangian martingales

$$dM_1(t) = I_{\{G_t<0, M_0(t)>0\}} dM_0(t)$$

where G_t indicates the time when there would never be exercise

$$G_t \equiv e^{-rt}\left[\frac{S_t - A_t}{t+\delta} - r(A_t - K)\right]$$

$$dM_2(t) = I_{\{M_0(t)>0\}} dM_0(t)$$

and M_0 denotes the positive part of the payoff of a corresponding European option

$$M_0(t) \equiv \mathbb{E}[e^{-rT}(A_T - K)|\mathcal{F}_t] = e^{-rT}\left\{\frac{\int_{-\delta}^{t} S_u du + S_t\left(e^{r(T-t)} - 1\right) - r}{T+\delta} - K\right\}$$

And for the last martingale, we consider the European style problem, whose value at time t will be

$$M_3'(t) = \mathbb{E}[e^{-rT}(A_T - K)^+|\mathcal{F}_t] = \frac{\mathbb{E}\left[e^{-rT}\left\{\int_t^T S_u du - \left(K(T+\delta) - \int_{-\delta}^{t} S_u du\right)\right\}\middle|\, \mathcal{F}_t\right]}{T+\delta}$$

Since there is no closed-form expression for this, we approximate the conditional distribution with a lognormal with matching first two moments, use the Black Scholes formula, and take the martingale part of the resulting expression as the third martingale M_3.

4.4.6 Connection to the dual problem

Having specified the preliminaries, we now show that with "regression later" a dual estimate can be computed with minimal additional effort. We use a separate set of paths for the regression, and then compute another set for pricing. For illustration, the procedure is demonstrated on a single pricing path.

Fix the original b paths $(S_{i1}, S_{i2}, \dots, S_{in})$, $i = 1, \dots, b$ used to estimate the regression coefficients $\hat{\gamma}_j$, $j = 1, \dots, n-1$ and simulate a new path $S_1, \dots, S_n$ independent of the other paths. We consider the coefficients fixed and proceed conditional on the original set of paths. Set $\hat{\gamma}_0 \equiv 0$ and define

$$\hat{C}_j^+(\,.\,) = \sum_{k=1}^{K} \hat{\gamma}_{jk}\psi_{jk}(\,.\,)\ , \qquad j = 1,2, \dots, n-1$$

$$\breve{V}_{j+1}^+(\,.\,) = \sum_{k=1}^{K} \hat{\gamma}_{jk}\psi_{j+1,k}(\,.\,)\ , \qquad j = 0,1, \dots, n-1$$

Conditional on the estimated coefficients $\hat{\gamma}_0, \hat{\gamma}_1, \dots, \hat{\gamma}_{m-1}$ we view $\hat{C}_j^+(s)$ and $\breve{V}_{j+1}^+$ as deterministic functions. Define the first time t_j at which the payoff from immediate exercise $g_j(S_j)$ exceeds the continuation value estimated by regression as

$$\hat{\tau} = \min\{j = 0, \dots, n : g_j(S_j) \geq \hat{C}_j^+(S_j)\}$$

Further define

$$M_0 = 0$$

$$M_u = \sum_{j=0}^{u-1} \left[\breve{V}_{j+1}^+(S_{j+1}) - \hat{C}_j^+(S_j)\right], \qquad u = 1, \dots, n$$

where each summand is simply

$$\breve{V}_{j+1}^+(S_{j+1}) - \hat{C}_j^+(S_j) = \sum_{k=1}^{K} \hat{\gamma}_{jk}\left[\psi_{j+1,k}(S_{j+1}) - \psi_{jk}(S_j)\right]$$

Let now $\underline{V_0}$ denote the low biased estimator, and $\overline{V_0}$ the high biased estimator. If the basis functions used are martingales, Glasserman and Yu [13] prove the following statement

$$\underline{V_0}(S_0) = \mathbb{E}[g_{\hat{\tau}}(S_{\hat{\tau}})] \leq V_0(S_0) \leq \mathbb{E}\left[\max_{\mathrm{u}=1,2,\dots,\mathrm{n}} g_u(S_u) - M_u\right] = \overline{V_0}$$

Thus, the true value V_0 is bounded above and below by terms that can be estimated through simulation. As we assume the basis functions are martingales and use “regression later” we have all the required information to calculate the summands in M_u. The upper bound can be estimated by simulating an independent set of paths, computing the summands at each time step, adding them to get M_u and then taking the maximum of $g_u(S_u) - M_u$ along the path. The key point is that using “regression later”, the martingale terms are available at almost no cost.

And while the upper bound should work very well in practice, the lower bound should be superior, as experience suggests it is closer to the true price. The lower bound can be estimated from these same independent paths each stopping according to the rule defining $\hat{\tau}$. The stopping rule $\hat{\tau}$ is not a stopping time with respect to the history $S_1, \dots, S_j$ because it depends on the estimated coefficients $\hat{\gamma}_j$. It is however a randomized stopping time because the event $\{\tau = t_j\}$ is contained in the sigma algebra generated by $S_1, \dots, S_j$ and $\hat{\gamma}_j^\top = (\hat{\gamma}_1, \dots, \hat{\gamma}_{n-1})$

Note that $0 = M_0, M_1, \dots, M_n$ is indeed a martingale conditional on $\hat{\gamma}$

$$\mathbb{E}\left[\, M_{j+1} \mid S_1, S_2, \dots, S_j, \hat{\gamma} \,\right] = M_j$$

Which is evident since the summands defining M_j have conditional expectation zero. While any martingale would provide an upper bound, the one specified is close to optimal. The martingale

$$M_u = \sum_{j=0}^{u-1} \left[V_{j+1}(S_{j+1}) - C_j(S_j)\right]$$

constructed from the true value and continuation functions turns the upper bound into an equality. Thus, the specified martingale is in a sense a best approximation to the optimal martingale, given the choice of basis functions. The drawback to the procedure is that it places more restrictive conditions on the available (martingale) basis functions.

4.4.7 Implementation

Perhaps the obvious way to compute the lower and upper bounds is in a sequential fashion so that after estimating the Q-value functions, we simulate a number of sample paths to compute the lower bound and another set to compute the upper bound. One difficulty with this strategy is that the difference between the upper and lower bound might be significant, so there is a large duality gap. We therefore specify the regression on a set of regression paths, and then estimate both bounds on a new set of independent pricing paths. We use only in-the-money paths in the regression.

Regression

1. Simulate b independent regression paths of the Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$
2. Initialize the recursion – at maturity for each path set set

   $\hat{V}_{in} = g_n(S_{in})$ used to compute the lower bound

   $\hat{V}_{in}^{+} = g_n(S_{in})$ used to compute the upper bound
3. Apply backward induction until $\hat{V}_{i1}$ and $\hat{V}_{i1}^{+}$ is reached for every path

   At each time step $j = n-1, \dots, 1$ do the following (calculation for all paths, fixed j)

   a. Regression now

      Calculate $\hat{C}_j(s)$ by regressing $\hat{V}_{j+1}(S_{j+1})$ on the current state s (in-the-money paths)

      i. Discount the estimated option value to current time and set $Y = D_{j,j+1}\hat{V}_{j+1}$

      ii. Set the $b \times K$ matrix of basis functions $X_{ik} = \psi_{jk}(S_{ij})$

      iii. Compute the regression coefficients $\hat{\beta}_j = (X^{\top}X)^{-1}X^{\top}Y$

      iv. Estimate the continuation value from $\hat{C}_j(S_j) = X\hat{\beta}_j$

      Following an optimal exercise decision for each path set

$$\hat{V}_j = \begin{cases} g_j(S_{ij}), & g_j(S_{ij}) \geq \hat{C}_j(S_{ij}), \quad exercise \\ D_{j,j+1}\hat{V}_{j+1}, & g_j(S_{ij}) < \hat{C}_j(S_{ij}), \quad hold \end{cases}$$

   b. Regression later

      Calculate $\hat{C}_j^{+}(s)$ by regressing $\hat{V}_{j+1}^{+}(S_{j+1})$ on the later state (in-the-money paths)

      i. Set the column vector $Y = \hat{V}_{j+1}^{+}$

      ii. Set the $b \times K$ matrix of basis functions $X_{ik} = \psi_{j+1,k}(S_{i,j+1})$

      iii. Compute the regression coefficients $\hat{\gamma}_j = (X^{\top}X)^{-1}X^{\top}Y$

iv. Matrix of basis functions at current time step $X_{ik} = \psi_{jk}(S_{ij})$

v. Estimate the continuation value from $\hat{C}_j^+(S_j) = X\hat{\gamma}_j$

Following an optimal exercise decision for each path set

$$\hat{V}_j^+ = \begin{cases} g_j(S_{ij}), & g_j(S_{ij}) \geq \hat{C}_j^+(S_{ij}), \quad exercise \\ D_{j,j+1}\hat{V}_{j+1}^+, & g_j(S_{ij}) < \hat{C}_j^+(S_{ij}), \quad hold \end{cases}$$

4. Discount the obtained values for each path one step back to initial time, and then calculate the average to obtain the estimates

Estimate obtained using regression now $\underline{\hat{V}_0} = \frac{1}{b} D_{0,1}[\hat{V}_{11} + \hat{V}_{21} + \cdots + \hat{V}_{b1}]$

Estimate obtained using regression later $\underline{\hat{V}_0^+} = \frac{1}{b} D_{0,1}[\hat{V}_{11}^+ + \hat{V}_{21}^+ + \cdots + \hat{V}_{b1}^+]$

We store the obtained coefficients and proceed with pricing. As the only recursive component was the calculation of coefficients, we no longer need to work backwards in time.

Pricing

1. Simulate m independent pricing paths of the Markov chain $\{S_{i1}, S_{i2}, \dots, S_{in}\}$

   Requires martingale basis functions, and the computed coefficient estimates $\hat{\gamma}_{jk}$

2. Along each path $i = 1$ to m do

   For each time step $j = 1$ to n do

   $MG = 0$ $\quad U = 0$

   Immediate exercise value $g_{ij} = g_j(S_{ij})$

   Estimated continuation value $\hat{C}_{ij}^+ = \hat{\gamma}_{jk}\psi_{jk}(S_{ij})$

   Estimated future option value $\tilde{V}_{ij}^+ = \hat{\gamma}_{jk}\psi_{j+1,k}(S_{ij+1})$

   If $(g_{ij} > \hat{C}_{ij}^+)$ Set $\underline{\hat{V}_{i1}^+} = D_{0,1}g_{ij}$ (do not repeat)

   $MG = MG + \tilde{V}_{ij}^+ - \hat{C}_{ij}^+$

   $U = \max(g_{ij}, MG)$

   $\overline{V}_{i1}^+ = D_{0,1}U$

3. Low biased estimator $\underline{\hat{V}_0^+} = \frac{1}{m}\left[\underline{\hat{V}_{11}^+} + \underline{\hat{V}_{21}^+} + \cdots + \underline{\hat{V}_{m1}^+}\right]$

   High biased estimator $\overline{V}_0^+ = \frac{1}{m}\left[\overline{V}_{11}^+ + \overline{V}_{21}^+ + \cdots + \overline{V}_{m1}^+\right]$

Chapter 5

Numerical Analysis

After introducing all the necessary concepts, we now turn to specific examples. In this section we present the computational results of the aforementioned algorithms in pricing both plain vanilla and exotic path dependent options. While simple Monte Carlo for European options won't prove to be as involving, the early exercise feature of American options requires the use of additional procedures to estimate the price. We employ the regression-based method of Longstaff and Schwartz [21] in conjunction with the dual method of Rogers [27] and Haugh and Kogan [15] to produce a lower and an upper bound. The implementation of the algorithm follows along the lines of Glasserman and Yu [13] who employ "regression later" to avoid nested simulations and provide tight bounds with minimal computational cost. The purpose of these numerical experiments is to evaluate the duality relations derived at the very beginning. Using Monte Carlo simulation, we approximate the left- and the right-hand side of the symmetries and compare the outcomes.

The reported results are therefore divided into four categories. To illustrate the accuracy of the implementation we begin by evaluating European puts and calls whose prices can be exactly calculated. The estimates of the corresponding American options and that of European Asian options are cross referenced to results of other pricing techniques like the one based on partial differential equations and their numerical solution via finite differences. However, when it comes to pricing Asian options of American type, the backing literature is not that extensive. The proposed method is in fact Monte Carlo simulation, and we report these results in the last section.

All simulation estimates are derived within the now well-known Black Scholes Merton framework, where the underlying asset distributes a continuous dividend yield q. The dynamics of the process are given by

$$dS_t = (r - q)S_t dt + \sigma S_t dW_t$$

To check the symmetries, we do not simulate processes under the dual measure, but rather estimate the price of the option with accordingly changed parameters in the auxiliary economy.

In the analysis we use several sets of parameters and evaluate each option on a large number of paths and exercise opportunities. We use antithetic variates and a suitable control variable to reduce the variance in the simulation. As is common with Monte Carlo simulations, along with the estimates, we report the respective standard error.

For options of European exercise, we only need to estimate the outermost expectation in the definition of the price. The added difficulty of determining the optimal stopping problem of American options requires the additional approximation of the conditional expectations in the defining ADP equations. For options of American type, we therefore use a different set of paths for the purpose of regression. This number will be smaller as the error from regression is usually much smaller than that arising from approximating the conditional expectation function.

Longstaff and Schwartz argue that using only in the money paths (and regression on the current state) yields the best linear unbiased estimator (of the continuation values) based on mean square metric. The approach does not assume further restrictions on the basis functions which allows us to implement various orthogonal polynomial basis discussed by de Lima and Semanez [7]. For reference, we report this estimate alongside the definite upper and lower bound resulting from the dual approach. To obtain a sufficiently small duality gap, we follow the suggestion of Haugh and Kogan and repeat the whole procedure for several batches. Another alternative is to use more training points or a more flexible approximation structure. Further problem-specific improvements will be detailed in the text.

5.1 Put Call duality

The arbitrage prices of a European call resp. put are given by the following expression

$$c(S_0, K, r, q, 0, T) = e^{-rT}\mathbb{E}(S_T - K)^+$$

$$p(S_0, K, r, q, 0, T) = e^{-rT}\mathbb{E}(K - S_T)^+$$

And that of the corresponding American with

$$C(S_0, K, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau}\mathbb{E}(S_\tau - K)^+$$

$$P(S_0, Kr, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau}\mathbb{E}(K - S_\tau)^+$$

We now specify the set of parameters used in the simulation. Further examples are introduced by varying one parameter while holding the others fixed.

Parameter set Consider an option written on a single asset modeled as a $GMB(r - q, \sigma^2)$ with interest rate $r = 3\%$, dividend yield $q = 2\%$, and the volatility $\sigma = 0.20$. The initial price is set at $S_0 = 100$ and the option's strike price K is also 100. The option expires in $T = 0.5$ years and we do not allow for exercise at the initial time $t_0 = 0$.

5.1.1 European exercise

With the prices of a European call and put defined as above, we examine the following symmetry relation

$$c(S_0, K, r, q, 0, T) = p(K, S_0, q, r, 0, T)$$

According to the suggested parameters, we simulate a Call option with initial price S_0, strike K, interest rate r and dividend yield q. The same random numbers are used to construct paths for a Put option with initial price K, strike S_0, interest rate q, and dividend yield r. Note how the roles of both the strike and initial price, and the interest rate and dividend yield have interchanged.

The Monte Carlo simulation was conducted on 1 000 000 paths with 1000 gridpoints. We consider the parameter set and vary the strike to produce more examples with estimates reported in the table below. As we want to asses whether the equality holds, it will also be convenient to calculate the discrepancies between the two options. We report the standard error and compare the estimates to the Black Scholes price.

Table 5.1.1 European Put Call Duality – varying the strike

Strike	BS Price	CALL Estimate	CALL SE	PUT Estimate	PUT SE	Discrepancy
70	30.06668	30.06951	0.2351739	30.06507	0.1631699	0.00444
80	20.47692	20.48026	0.2276801	20.47554	0.1765214	0.00472
90	11.99107	11.99897	0.2004293	11.99422	0.1655235	0.00474
100	5.817557	5.828771	0.1517544	5.826008	0.1284881	0.00267
110	2.316143	2.323614	0.0982206	2.322965	0.0832154	0.00065
120	0.766436	0.769109	0.0558244	0.769034	0.0466241	0.00008

To demonstrate the rate of convergence, we fix $K = 100$ and report the simulation average for every 2^n paths. The results are plotted in the graph below.

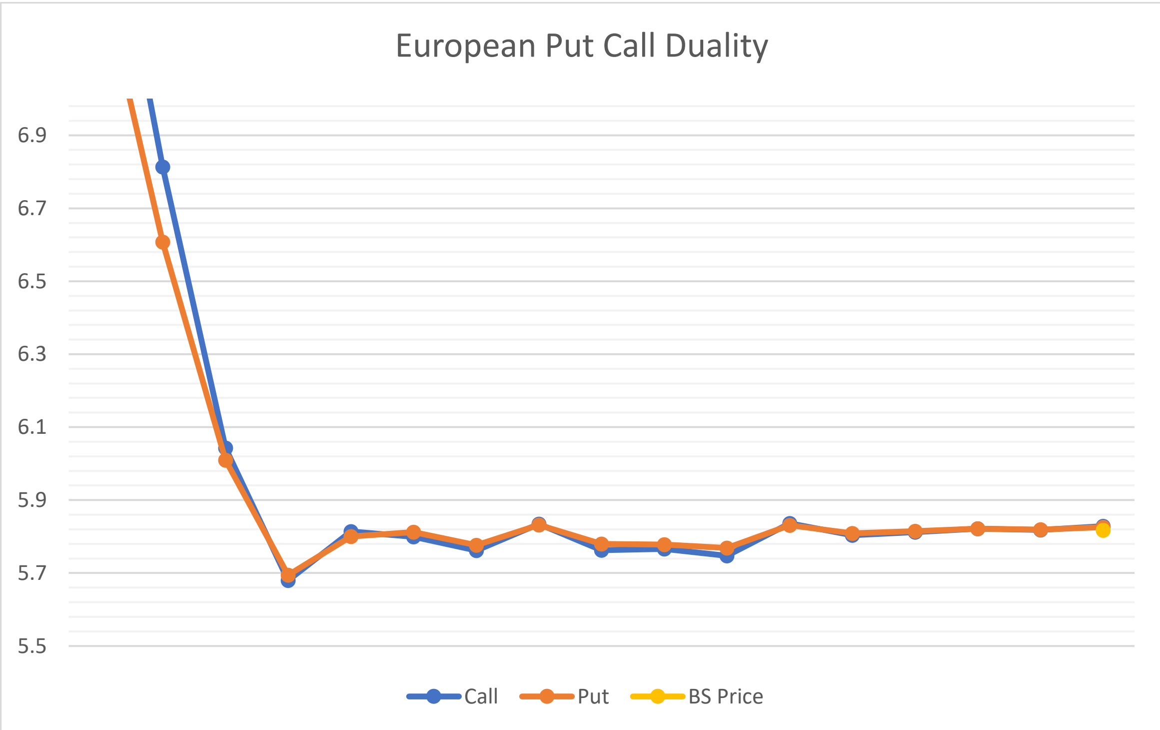

As can be seen, the estimates converge to the corresponding Black Scholes prices. The discrepancies between the two options are barely noticeable.

5.1.2 American exercise

The put call duality for American options states the following

$$C(S_0, K, r, q, 0, T) = P(K, S_0, q, r, 0, T)$$

The parameter exchange in the symmetry is the same as that for European options, so we simulate the trajectories in an equal manner. Due to the early exercise feature, the price of the American option should be superior to that of the European. Therefore, situations of nontrivial exercise will be of particular interest.

In addition to antithetic variates, we include the corresponding European option as a control variable. To price the option using regression we need to specify the set of basis functions. For the Longstaff and Schwartz estimate, which uses regression on the current state, we use the payoff function and the first three power polynomials of the underlying spot prices. The method is contrasted with regression later and the dual approach, providing both a lower and an upper bound on the price. The set of basis functions in this case must be martingales. We follow the suggestion of Rogers [27] and use the discounted value of the corresponding European option. Only in-the-money paths were used in the regression.

The simulation was conducted on $200\,000$ pricing paths and 7000 regression paths, with 60 time steps each and the procedure repeated for 5 batches. All three estimates, and their corresponding standard errors are noted in the table below. We cross-check the result with those obtained using a finite difference technique. The standard error for each estimate is reported in a separate table.

Table 5.1.2 American Put Call Duality – varying the strike

K		CALL			PUT		
	FD	LSM	Lower	Upper	LSM	Lower	Upper
70	30.08291	30.0629	30.0545	30.5082	29.9227	29.9852	30.3639
80	20.47882	20.4612	20.3454	20.7726	20.3700	20.2954	20.6760
90	11.99124	11.9629	11.7416	12.1658	11.9181	11.7432	12.1078
100	5.817571	5.73412	5.51557	5.90867	5.66059	5.52277	5.88627
110	2.316144	1.37476	2.06011	2.41819	1.51671	1.97768	2.52396
120	0.766436	0.00749	0.59763	1.05468	0.13482	0.56641	1.10841

Table 5.1.3 Corresponding standard error

K	CALL			PUT		
	SE LSM	SE Lower	SE Upper	SE LSM	SE Lower	SE Upper
70	0.0013	0.0165	0.0020	0.0007	0.0105	0.0010
80	0.0023	0.0047	0.0031	0.0013	0.0039	0.0003
90	0.0133	0.0087	0.0099	0.0086	0.0048	0.0042
100	0.0169	0.0135	0.0107	0.0095	0.0106	0.0076
110	0.0135	0.0097	0.0106	0.0058	0.0066	0.0087
120	0.0003	0.0033	0.0052	0.0062	0.0045	0.0035

Note how the estimated price of the American option is consistently higher than that of the corresponding European option. In the case when the dividend yield is zero, the price of the American call corresponds to that of the European call, which is not the case for a put option. When the dividend yield is small, this distinction can decay asymptotically, so we will additionally consider varying q and document the results below.

Table 5.1.4 American Put Call Duality – varying the dividend yield

q		CALL			PUT		
	FD	LSM	Lower	Upper	LSM	Lower	Upper
0.02	5.81757	5.7341	5.5155	5.6424	5.6605	5.5227	5.8862
0.025	5.68642	5.6191	5.4115	5.5126	5.5545	5.4199	5.7677
0.03	5.56794	5.5105	5.3167	5.3849	5.4449	5.3314	5.6507
0.035	5.45693	5.4015	5.2262	5.2593	5.3399	5.2373	5.5353
0.04	5.35053	5.2932	5.1339	5.1358	5.2337	5.1487	5.4218
0.045	5.24795	5.1868	5.0426	5.6424	5.1329	5.0563	5.3096

Table 5.1.5 Corresponding standard error

q	CALL			PUT		
	SE LSM	SE Lower	SE Upper	SE LSM	SE Lower	SE Upper
0.02	0.0169	0.0132	0.0107	0.0095	0.0106	0.0076
0.025	0.0143	0.0133	0.0107	0.0105	0.0107	0.0076
0.03	0.0126	0.0138	0.0108	0.0093	0.0094	0.0076
0.035	0.0103	0.0131	0.0111	0.0109	0.0091	0.0076
0.04	0.0102	0.0131	0.0112	0.0089	0.0125	0.0076
0.045	0.0091	0.0132	0.0111	0.0087	0.0099	0.0077

Like in the previous section, we illustrate the rate of convergence for a fixed strike price of 100 by recording the “running ”estimate during the course of the approximation (for a single batch). The LSM estimate and the lower and upper bound are reported in separate graphs below

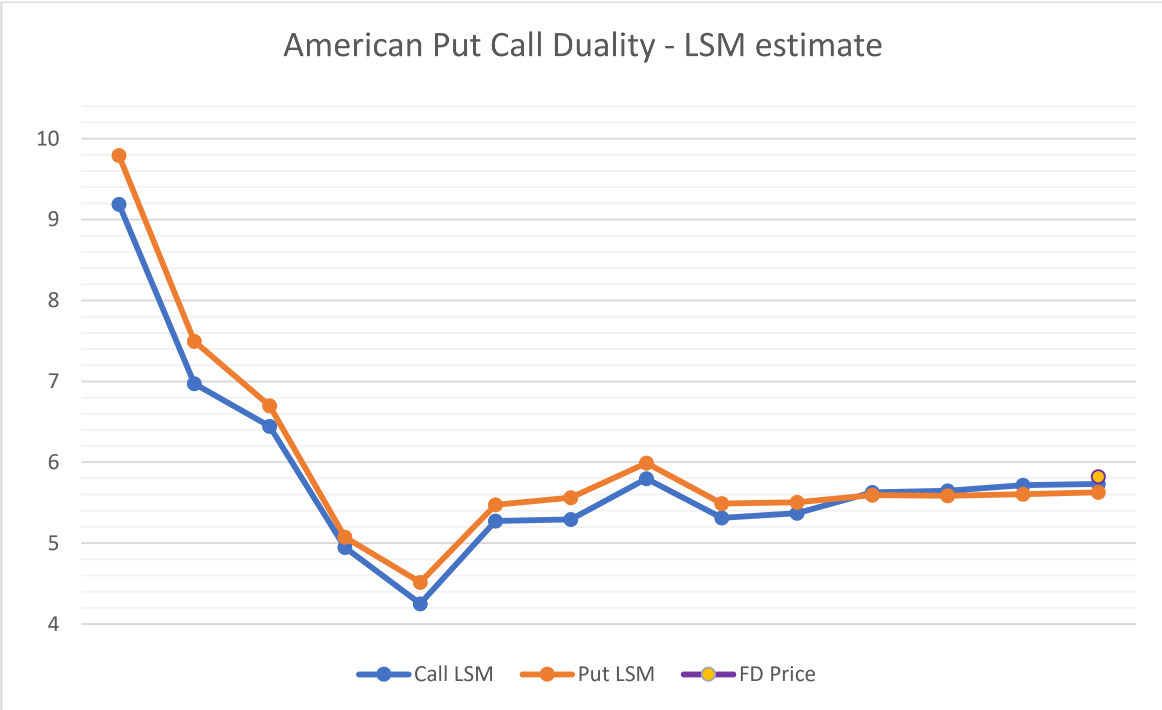

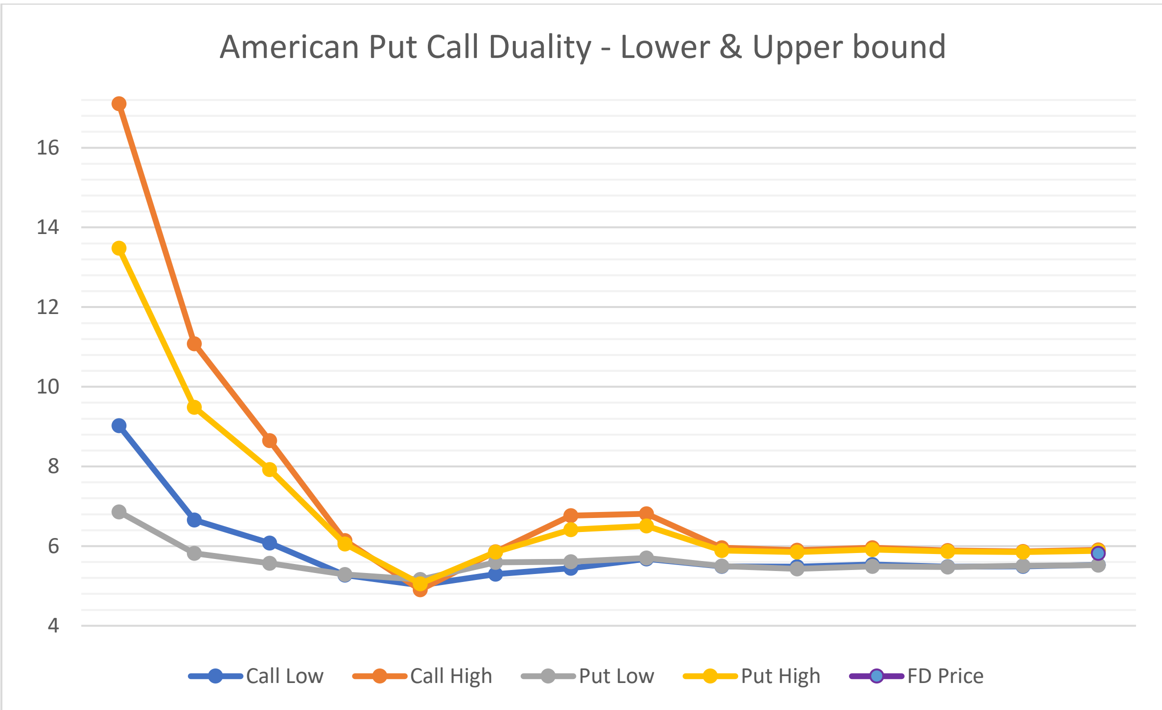

Note how the LSM estimates are consistently biased low, while the upper and lower bounds straddle the option price. With the number of batches increasing, the bounds grow closer together, thus providing a better estimate. The results still suggest minor discrepancies, which can be attributed to the errors arising from approximation. The LSM estimate is indeed closer to the true price of the option.

5.2 Asian option duality

The arbitrage prices of European Asian options of floating and fixed strike are

$$c_x(K, S_0, r, q, 0, T) = e^{-rT}\mathbb{E}(A_T - K)^+$$

$$p_x(K, S_0, r, q, 0, T) = e^{-rT}\mathbb{E}(K - A_T)^+$$

$$c_f(S_0, \lambda, r, q, 0, T) = e^{-rT}\mathbb{E}(\lambda S_T - A_T)^+$$

$$p_f(S_0, \lambda, r, q, 0, T) = e^{-rT}\mathbb{E}(A_T - \lambda S_T)^+$$

And the corresponding American prices read

$$C_x(K, S_0, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau}\mathbb{E}(A_\tau - K)^+$$

$$P_x(K, S_0, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau}\mathbb{E}(K - A_\tau)^+$$

$$C_f(S_0, \lambda, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau}\mathbb{E}(\lambda S_\tau - A_\tau)^+$$

$$P_f(S_0, \lambda, r, q, 0, T) = \sup_{\tau \in \mathcal{T}_{[0,T]}} e^{-r\tau}\mathbb{E}(A_\tau - \lambda S_\tau)^+$$

We begin with the same parameter set specified at the beginning of the section. The primary case of concern will be to evaluate the duality relations where $\lambda = 1$, but we additionally consider varying the multiplier parameter and document the results.

5.2.1 European exercise

The symmetry relation between the two types of Asian options states

a. $c_f(S_0, \lambda, r, q, 0, T) = p_x(\lambda S_0, S_0, q, r, 0, T)$

b. $c_x(K, S_0, r, q, 0, T) = p_f\left(S_0, \frac{K}{S_0}, q, r, 0, T\right)$

The first equation relates a floating strike Asian call to a fixed strike Asian put. The dual measure is implied in the changed parameters, so the fixed strike put is assumed to have a strike value of λS_0, with the simulated trajectories following the dynamics of a process in the auxiliary economy with initial spot S_0 and reversed roles of q and r. Similarly, the right-hand side of the second relation is the value of a floating strike put in the auxiliary economy with interest rate q, initial spot S_0 and the multiplier parameter equal to S_0/K.

The simulation was carried out for $2\,000\,000$ paths each with 60 time steps. The following tables report the obtained estimates for the two noted duality relations. For reference, we also document an estimate resulting from a Turnbull and Wakeman's approximation procedure. Since varying the strike does not really affect the average strike Asian option, we report only the results for the second duality relation.

Table 5.2.1 American Fixed Call – Floating Put Asian Duality – varying strike

Strike		Fixed Call		Floating Put		
	T&W	Estimate	SE	Estimate	SE	Discrepancy
80	19.82685	19.95289	0.305599	19.95411	0.219891	0.00122
90	10.43317	10.40941	0.286282	10.41144	0.236408	0.00203
100	3.803244	2.464029	0.124441	2.463314	0.124719	0.00071
110	0.503149	0.547107	0.078274	0.549127	0.067386	0.00202
120	0.034318	0.047546	0.021569	0.047536	0.017815	0.00001

Due to the averaging mechanism the Asian options are less volatile, which, as can be seen in the above table, makes them less expensive than their European counterparts. The following table documents the estimates for the first duality relation and a varying multiplier parameter λ.

Table 5.2.2 American Floating Call – Fixed Put Asian Duality – varying multiplier

λ		Floating Call		Fixed Put		
	T&W	Estimate	SE	Estimate	SE	Discrepancy
0.8	0.008074	0.005684	0.002866	0.005668	0.002296	0.00002
0.9	0.363363	0.356586	0.026801	0.355913	0.022831	0.00067
1	3.163138	2.362014	0.117873	2.362009	0.118138	0.10201
1.1	10.53749	10.63192	0.328544	10.63183	0.270544	0.00009
1.2	19.88598	20.08686	0.401955	20.08875	0.301211	0.00184

Concluding, we illustrate the rate of convergence for $K = 100,\ \lambda = 1$. The first duality is reported on the primary vertical axis to the left, and the second on the axis to the right.

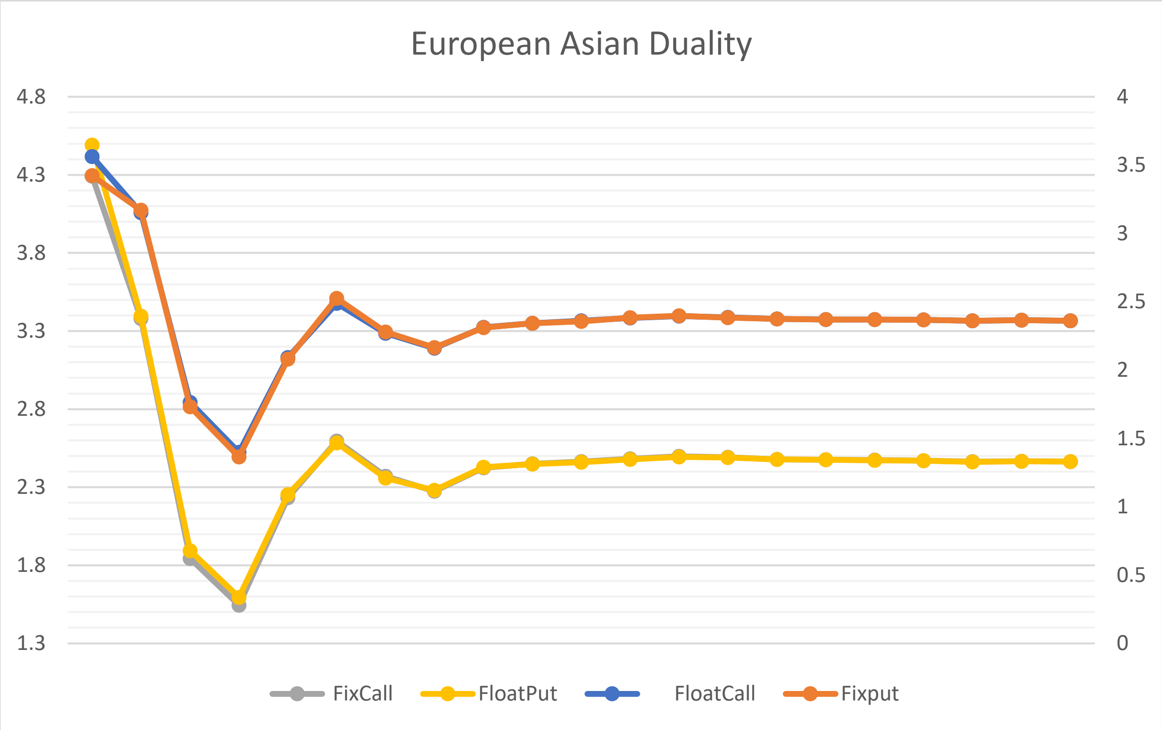

Just like in the European case, the discrepancies between the price estimates reported are barely noticeable.

5.2.2 American exercise

With the notation as above, Gounden and O'Hara claim the following relations hold

a. $C_f(S, \lambda, r, q, 0, T) = P_x(\lambda S, S, q, r, 0, T)$

b. $C_x(K, S, r, q, 0, T) = P_f\left(S, \frac{K}{S}, q, r, 0, T\right)$

The trajectories used to evaluate the above duality relations are simulated in the same way as that of the European Asian option.

For each type of arithmetic Asian option, the corresponding geometric average Asian is used to control the error of the approximation. Following de Lima and Semanez [7], we implement Hermite A polynomials as basis functions for a put option, and Power polynomials as basis for a call option, regardless of the type. The average of asset prices is included as a relevant variable next to the spot price. If we denote by X_j the relevant variable for the j-th time step, the martingale basis consists of the following three martingales

$$\psi_{j,0}(X_j) = 1$$

$$\psi_{j,1}(X_j) = X_j \exp\left(-(r-q)(t_j - t_0)\right)$$

$$\psi_{j,2}(X_j) = X_j^2 \exp\left(-(2(r-q)+\sigma^2)(t_j - t_0)\right) C_{FL}^{LSM}$$

We use 7000 paths in the regression, and another 200 000 paths for pricing. Options can be exercised at 60 equally spaced dates and the procedure is repeated for 5 batches. Both duality relations were evaluated on the same set of paths. For a varying strike, we report the estimates and the corresponding standard error of the second duality relation

Table 5.2.3 American Fixed Call – Floating Put Asian Duality – varying strike

K	Fixed Call			Floating Put		
	LSM	Lower	Upper	LSM	Lower	Upper
70	31.4734	31.5146	30.0923	29.5431	31.5943	30.2452
80	21.6627	21.5944	20.1494	19.5962	21.6314	20.2513
90	11.9879	11.8927	10.5151	9.75869	11.6467	10.5727
100	4.08152	3.97975	3.37442	2.12798	2.99242	3.56726
110	0.41791	0.59054	0.71326	0.00362	0.32293	0.76684
120	0.00243	0.04262	0.08833	0	0.00866	0.08238
130	0	0.00151	0.00567	0	0.00011	0.00389

Table 5.2.4 Corresponding standard error

K	Fixed Call			Floating Put		
	SE LSM	SE Lower	SE Upper	SE LSM	SE Lower	SE Upper
70	0.0302	0.0162	0.0533	0.0507	0.0451	0.0883
80	0.0425	0.0226	0.0631	0.0508	0.0455	0.0882
90	0.0495	0.0281	0.0678	0.0445	0.0458	0.0827
100	0.0352	0.0264	0.0478	0.0252	0.04001	0.0594
110	0.0122	0.0131	0.0193	0.0009	0.0122	0.0233
120	0.0007	0.0028	0.0052	0	0.0012	0.0061
130	0	0.0005	0.0009	0	0.0001	0.0011

To better illustrate the discrepancies, we plot the above results as a function of strike.

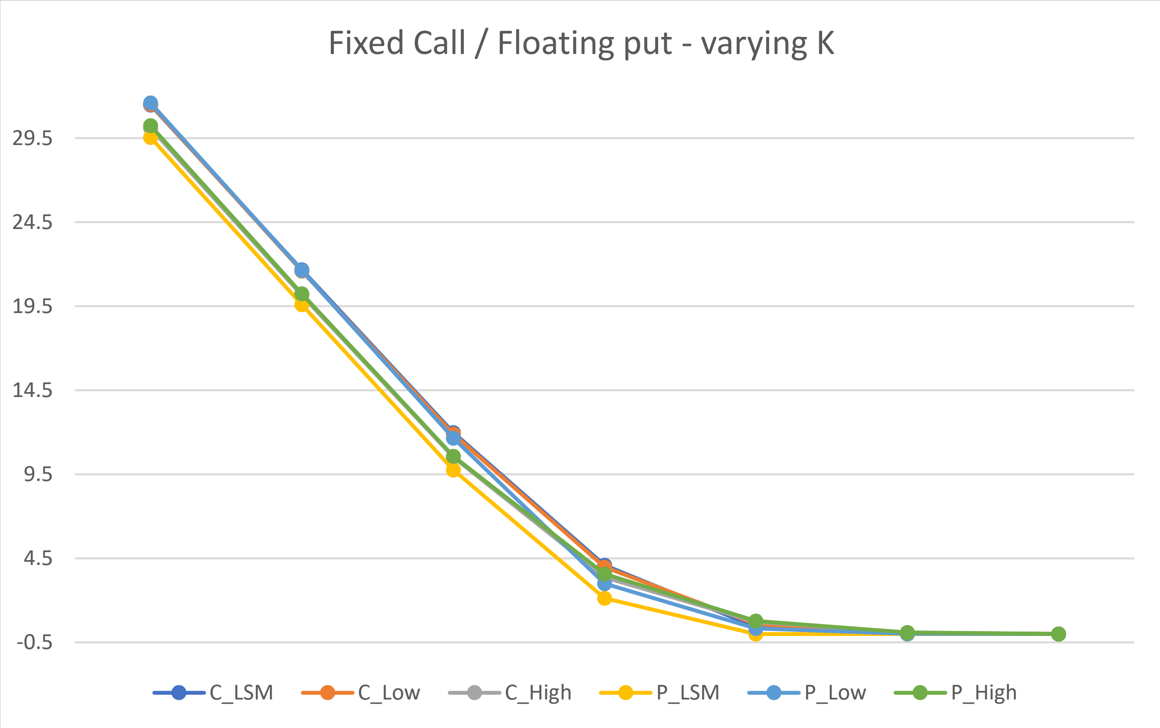

We additionally consider varying the multiplier parameter. The estimates for the first duality relation are reported below.

Table 5.2.5 American Floating Call – Fixed Put Asian Duality – varying multiplier

λ	Floating Call			Fixed Put		
	LSM	Lower	Upper	LSM	Lower	Upper
0.7	0	0	0	0	0	0
0.8	0	0.00482	0.01726	0	0.00041	0.00953
0.9	0.09354	0.38681	0.49964	0.01685	0.15672	0.52121
1	4.15093	3.99387	3.41242	2.12657	2.98534	3.50512
1.1	12.1392	12.0446	10.7933	9.83251	11.6168	10.7382
1.2	21.6222	21.6955	20.3864	19.6081	21.6077	20.2854
1.3	31.4453	31.6276	30.3948	29.5559	31.5691	30.2464

Table 5.2.6 Corresponding standard error

λ	Floating Call			Fixed Put		
	SE LSM	SE Lower	SE Upper	SE LSM	SE Lower	SE Upper
0.7	0	0	0	0	0	0
0.8	0	0.0009	0.0015	0	0.0001	0.0013
0.9	0.0075	0.0118	0.0179	0.0016	0.0057	0.0151
1	0.0418	0.0317	0.0573	0.0229	0.036	0.0502
1.1	0.0774	0.044	0.0947	0.0393	0.0432	0.0777
1.2	0.0997	0.0517	0.1158	0.0477	0.0428	0.0865
1.3	0.1157	0.0591	0.1311	0.0489	0.0423	0.0877

Again, the inconsistencies between the estimates are best observed on a graph. We plot the results as a function of λ.

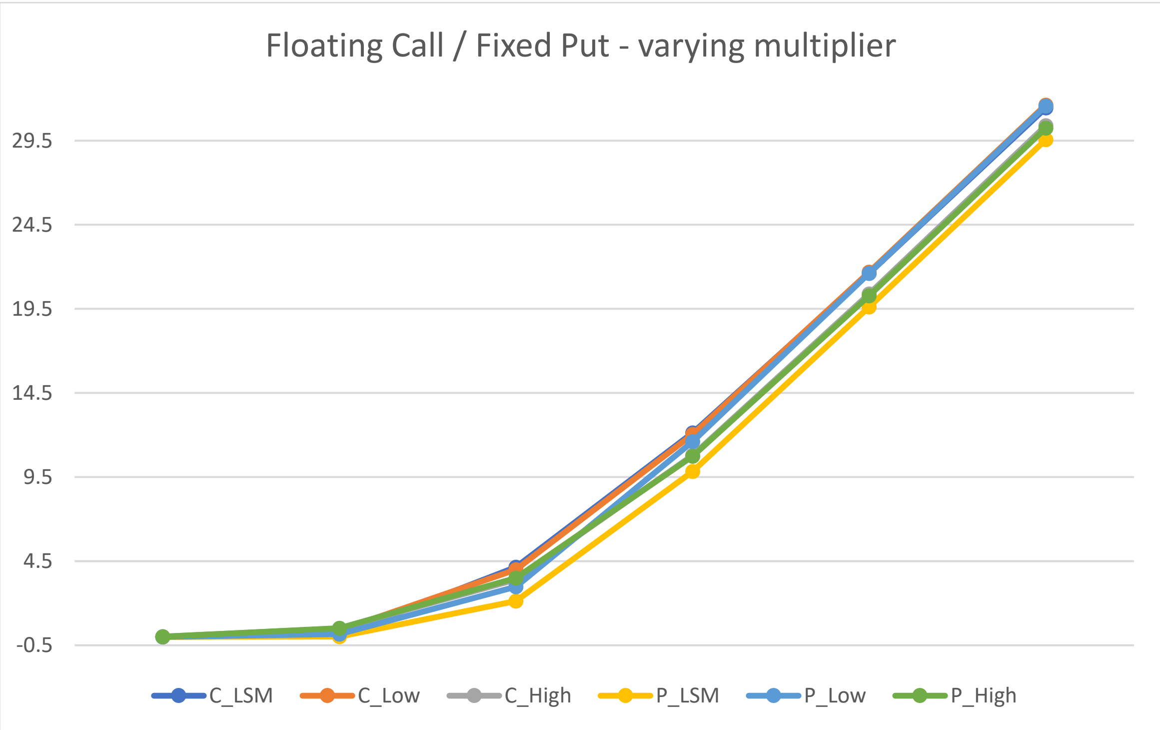

5.3 Conclusion

In view of the given examples, we can conclude the following.

The pricing of European options, both plain vanilla and Asian type, requires much less computational effort and accurate estimates are fairly easy to obtain. For the American type, meeting the desired degree of accuracy requires a good choice of basis functions and therefore good knowledge of the option payoff. An equally large set of paths is sometimes hard to compute due to the higher dimensionality. Variance reduction techniques in this case play a crucial role.

In the estimated duality relations, the American Asian case really does produce larger discrepancies between the price estimates. However, some discrepancies are noted even in the American Put Call duality. As Asian options have particularly complex payoffs, and further refinements regarding the choice of basis functions can still be made, it remains unclear whether the differences stem from approximation error or is it really an indication that the duality relation does not hold. It could also be that the suggested evaluation method is not well suited, or precise enough for the given problem. In any case, without further analysis, the question remains open to investigation.

Bibliography

[1] Baxter, M. & Rennie, A., 2012. *Financial Calculus - An Introduction to Derivative Pricing.* s.l.:Cambridge University Press.

[2] Ben-Ameur, H., Breton, M. & L'Ecuyerhttp, P., 2002. A Dynamic Programming Procedure for Pricing American-Style Asian Options. *Management Science*, 48(5), pp. 625-643.

[3] Bernt, Ø., 2000. *Stochastic Differential Equations: An Introduction with Applications.* s.l.:Springer Berlin Heidelberg.

[4] Björk, T., 2009. *Arbitrage Theory in Continuous Time.* s.l.:Oxford University Press.

[5] Black, F. & Scholes, M., 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy,* 81(3), pp. 637-654.

[6] Broadie, M. & Glasserman, P., 1997. Pricing American-Style Securities Using Simulation. *Journal of Economic Dynamics and Control,* Volume 21, p. 1323–1352.

[7] de Lima, U. S. M. & Samanez, C. P., 2016. Complex Derivatives Valuation: Applying the Least-Squares Monte Carlo Simulation Method with Several Polynomial Basis. *Financial Innovation*, 2(1), p. 1.

[8] Duffie, D., 2001. *Dynamic Asset Pricing Theory.* s.l.:Princeton University Press.

[9] Firth, N. P., 2004. Upper Bounds for American Option Prices Using Regression with Martingale Basis Functions.

[10] Föllmer, H. & Schied, A., 2004. *Stochastic Finance: An Introduction In Discrete Time,* s.l.: de Gruyter.

[11] Glasserman, P., 2003. *Monte Carlo Methods in Financial Engineering.* s.l.:Springer Berlin Heidelberg.

[12] Glasserman, P. & Yu, B., 2004. Number of Paths versus Number of Basis Functions in American Option Pricing. *The Annals of Applied Probability,* 14(4), pp. 2090-2119.

[13] Glasserman, P. & Yu, B., 2004. *Simulation for American Options: Regression Now or Regression Later?.* s.l., Springer Berlin Heidelberg, pp. 213-226.

[14] Gounden, S. & O'Hara, J. G., 2010. An Analytic Formula for the Price of an American Style Asian Option of Floating Strike Type. *Applied Mathematics and Computation,* 217(7), p. 2923–2936.

[15] Haugh, M. B. & Kogan, L., 2004. Pricing American Options: A Duality Approach. *Operations Research,* 52(2), pp. 258-270.

[16] Henderson, V. & Wojakowski, R., 2002. On the Equivalence of Floating and Fixed-Strike Asian Options. *Journal of Applied Probability,* 39(2), pp. 391-394.

[17] Investopedia, LLC, n.d. *www.investopedia.com,* s.l.: s.n.

[18] Kemna, A. G. & Vorst, T. A. C. F., 1990. A Pricing Method for Options Based on Average Asset Values. *Journal of Banking and Finance,* 14(1), pp. 113-129.

[19] Kwok, Y.-K., 2008. *Mathematical Models of Financial Derivatives.* Hong Kong: Springer Berlin Heidelberg.

[20] Lamberton, D., 2009. *Optimal Stopping and American Options.* s.l.:Lectures in Financial Mathematics, Université Paris-Est.

[21] Longstaff, F. A. & Schwartz, E. S., 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies*, 14(1), pp. 113-147.

[22] Marek, M. & Marek, R., 2005. *Martingale Methods in Financial Modelling.* s.l.:Springer Berlin Heidelberg.

[23] Merton, R. C., 1973. The Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science,* Volume 4, pp. 83-141.

[24] Peskir, G. & Shiryaev, A., 2006. *Optimal Stopping and Free-Boundary Problems.* s.l.:Lectures in Mathematics. ETH Zürich.

[25] Rasmussen, N. S., 2005. Control Variates for Monte Carlo Valuation of American Options. *Journal of Computational Finance,* 9(1), pp. 83-118.

[26] Rheinländer, T. & Saxton, J., 2011. *Hedging Derivatives.* Singapore: World Scientific.

[27] Rogers, L. C. G., 2002. Monte Carlo Valuation of American Options. *Mathematical Finance,* 12(3), p. 271–286.

[28] Schmock, U., 2019. *Stochastic Analysis.* s.l.:Lectures in Stochastic Analysis. TU Wien.

[29] Tilley, J. A., 1993. Valuing American Options in a Path Simulation Model. *Transactions of the Society of,* Volume 45, p. 83–104.

[30] Tsitsiklis, J. N. & Van Roy, B., 2001. Regression Methods for Pricing Complex American Style Options. *IEEE Transactions on Neural Networks,* 12(4), p. 694–703.

[31] Wang, H., 2012. *Monte Carlo Simulation with Applications to Finance.* s.l.:Chapman and Hall/CRC.