



MONTE CARLO AND QUASI-MONTE CARLO IN EUROPEAN AND ASIAN CALL OPTION PRICING

Lappeenranta–Lahti University of Technology LUT

Degree Programme in Computational Science and Artificial Intelligence, Bachelor's thesis

2026

Artturi Lehtinen

Examiners: D.Sc. (Tech.) Vesa Kaarnioja

Abstract

Lappeenranta–Lahti University of Technology LUT

LUT School of Engineering Sciences

Degree Programme in Computational Science and Artificial Intelligence

Artturi Lehtinen

Monte Carlo and Quasi-Monte Carlo in European and Asian Call Option Pricing

Bachelor's thesis

2026

32 pages, 9 figures, 1 table

Examiners: D.Sc. (Tech.) Vesa Kaarnioja

Keywords: Asian options, European options, option pricing, Monte Carlo, Quasi-Monte Carlo, convergence rate

This thesis examines the Monte Carlo numerical integration method and its counterpart, Quasi-Monte Carlo in the context of option pricing. Two common option types and a standard asset pricing model are introduced. The effect of the pricing problem dimension on the convergence rates of the methods is investigated. The results are analyzed numerically and through some statistical measures.

The numerical experiments are conducted using Python. Asset price paths, the corresponding payoffs, and the option prices obtained by discounting the mean payoff are computed. The experiments analyze payoff distributions, convergence rates across different option types, and the convergence behavior of Quasi-Monte Carlo in varying dimensions. All experiments are conducted using a single set of parameters and the random walk model of the stochastic process.

According to the results, both numerical methods produce similar payoff distributions. The convergence behavior of Monte Carlo stays consistent across option types, whereas the convergence of Quasi-Monte Carlo trails off in the case of Asian call option.

Tiivistelmä

Lappeenrannan–Lahden teknillinen yliopisto LUT

LUTin insinööritieteiden tiedekunta

Laskennallisen tieteen ja tekoälyn koulutusohjelma

Artturi Lehtinen

Monte Carlo- ja kvasi-Monte Carlo -menetelmät eurooppalaisten ja aasialaisten osto-optioiden hinnoittelussa

Kandidaatintyö

2026

32 sivua, 9 kuvaa, 1 taulukko

Tarkastajat: TkT Vesa Kaarnioja

Avainsanat: aasialaiset optiot, eurooppalaiset optiot, optioiden hinnoittelu, Monte Carlo, kvasi-Monte Carlo, suppenemisnopeus

Tämä tutkielma tarkastelee numeerista integroimismenetelmää Monte Carloa ja sen vastinetta kvasi-Monte Carloa optioiden hinnoittelussa. Kaksi optiotyyppiä ja standardi kohde-etuuden hinnoittelumalli esitellään. Hinnoitteluongelman dimension vaikutusta menetelmien suppenemisnopeuksiin tutkitaan. Tuloksia analysoidaan numeerisesti ja tiettyjen tilastollisten suureiden avulla.

Numeeriset esimerkit on toteutettu Python-ohjelmointikielellä. Omaisuuserien hintapolut, näihin polkuihin perustuvat tuotot sekä keskimääräisen tuoton diskonttaamalla saadut option hinnat lasketaan. Kokeet analysoivat tuottojakautumia, suppenemisnopeutta erilaisten optiotyyppien välillä ja kvasi-Monte Carlon suppenemiskäyttäytymistä eri dimensioissa. Kaikki kokeet toteutetaan käyttäen yhtä parametrijoukkoa ja stokastisen prosessin satunnaiskävely-mallia.

Tulosten perusteella molemmat numeeriset menetelmät tuottavat samankaltaiset tuottojakautumat. Monte Carlon suppenemiskäyttäytyminen pysyy tasaisena optiotyyppien välillä, kun taas kvasi-Monte Carlon konvergenssi heikkenee aasialaisen osto-option tapauksessa.

Declarations

Turnitin

The originality of this thesis has been reviewed with the Turnitin similarity checking service.

AI usage

The author of the thesis used the following AI-tools during the preparation of the thesis:

1. ChatGPT
 - a. Purpose of use: Assistance with code implementation, recommending suitable reference material and proofreading.
 - b. Explanation of the use of the tool: The tool was used to assist code implementation. The tool was used to request recommendations for reference material and also to proofread this thesis.
2. GitHub Copilot Completions
 - a. Purpose of use: Assistance with code implementation.
 - b. Explanation of the use of the tool: The tool was used to assist and accelerate code implementation for the numerical experiments, for example by generating proof-of-concept code and autocompleting repetitive lines of code.
3. Overleaf Writefull Integration
 - a. Purpose of use: Improving clarity of text and spelling.
 - b. Explanation of the use of the tool: The tool provided suggestions on how to alter the text. Suggestions were considered critically.

Responsibility

The author, Artturi Lehtinen, takes full responsibility for the content of this thesis and has reviewed and edited the content generated by the use of AI tools.

Table of contents

Abstract	
Declarations	
Lists of figures and tables	6
1 Introduction	7
1.1 Objectives and delimitation	8
1.2 Structure of the thesis	8
2 Options	9
2.1 Option basics	9
2.2 Asset price model	9
2.3 European call Options	10
2.4 Asian call Options	12
3 Numerical methods	14
3.1 Monte Carlo	14
3.2 Quasi-Monte Carlo	16
3.3 Simulated asset price paths	20
4 Numerical experiments	23
5 Conclusions	31
References	32

List of Figures

- 1 128 sampled points with MC & QMC
- 2 Asset price paths simulated with MC
- 3 Asset price paths simulated with QMC
- 4 Payoff distribution histograms of MC & QMC with European call options
- 5 Payoff distribution histograms of MC & QMC with Asian call options
- 6 The convergence rates of MC & QMC with European call options
- 7 The convergence rates of MC & QMC with Asian Call options
- 8 Convergence of QMC using varying dimensions with Asian call options
- 9 Convergence rates of QMC versus dimensions as a line diagram

List of Tables

- 1 Statistical quantities of payoffs from 2^{20} simulations

1 Introduction

Derivative securities constantly play a bigger role in today's world. Derivative instruments, such as options, are used for various purposes such as hedging and to amplify profits. Accurate pricing of these instruments is essential for efficient markets, because derivative markets improve pricing efficiency by contributing to the effective transmission of information to market prices (Hsiao & Tsai, 2018). Their pricing often requires advanced numerical methods such as Monte Carlo (MC) and Quasi-Monte Carlo (QMC). With a vast amount of computing power, these simulation methods are now faster than ever and executable on personal computers within minutes.

Regular MC is a widely used method in a variety of engineering problems due to its ease of implementation. Rather than sampling random points from, for example, a unit circle, MC in derivative security valuation involves simulating stochastic process paths, which represent underlying asset values, interest rates along with other model parameters. Although having a notoriously slow convergence rate, the value of MC as a computational tool stems from having a consistent convergence rate (the estimator's error with respect to the number of simulations) in all dimensions. Dimension refers to the number of integration variables in this work. (Glasserman, 2004.)

Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous function. MC and QMC methods are suitable for approximating integrals given by

$$I_d(f) = \int_{\mathbb{R}^d} f(x_1, \dots, x_d) \prod_{i=1}^d \frac{e^{-\frac{1}{2}x_i^2}}{\sqrt{2\pi}} dx_1 \cdots dx_d \approx \frac{1}{n} \sum_{i=1}^n f(\mathbf{t}_i), \quad (1)$$

where $\{\mathbf{t}_i\}_{i=1}^n \subset \mathbb{R}^d$ is a point set. In MC the point set is independently and identically distributed and comes from normal distribution $\mathbf{t}_1, \dots, \mathbf{t}_n \sim N(0, \mathcal{I})$, where \mathcal{I} is a $d \times d$ identity matrix. In QMC the point set $\mathbf{t}_1, \dots, \mathbf{t}_n$ is deterministically chosen.

QMC methods have been introduced to address the slow convergence of standard MC methods. Instead of sampling randomly, QMC methods are equally weighted numerical integration methods that rely on sets of deterministic points that cover the sample space uniformly (L'Ecuyer, 2009). This often leads to a higher convergence rate, but might lose advantage in higher dimensions (Glasserman, 2004).

Recent research has continued investigating the performance of QMC and MC methods in financial applications. In option pricing and sensitivity analysis, QMC was found to outperform MC's convergence rate even in the higher dimensional simulations (Scoleri, Bianchetti

& Kucherenko, 2021). Similarly, another study that conducted numerical experiments on geometric basket and Asian call options found that QMC offered great improvements over MC in simulations with moderate dimensionality (Case, 2025).

1.1 Objectives and delimitation

The goal of the thesis is to compare traditional Monte Carlo and lattice rule Quasi-Monte Carlo methods in option pricing, more specifically in European and Asian call option pricing. The results will be compared with the analytical solutions which will be used as benchmarks for option prices. The deviations of the simulated option prices from analytical solutions are utilized in the convergence plots. Convergence behavior across the methods in two option types is the main object of study.

The thesis considers the following research questions:

- What are the differences between the MC and QMC methods used in this thesis?
- How do the MC and QMC methods perform in Asian and European option pricing?
- How can these methods be applied to option pricing?
- What part does dimension play in the convergence of the methods?

1.2 Structure of the thesis

This thesis contains five sections. The basics of options and the underlying asset price model together with European and Asian options are presented in the second section. The third section covers the numerical methods used in this thesis to simulate option prices. In the fourth section the results from the numerical experiments are introduced and reviewed. The last section discusses the conclusions of the thesis and possible further research.

2 Options

This chapter introduces the underlying theory of options and option pricing used in this thesis. First, basic financial options are defined, including call and put options. After this, a standard asset price model used in this thesis is presented. Finally, the two types of options and their pricing under the stochastic pricing model considered in this thesis are described.

2.1 Option basics

Options are rights to either sell or purchase the underlying asset at a future date for a fixed price. There exist two types of options. Call is an option that grants the holder the right to acquire the asset at a predefined price and date. A put option provides the right to sell the underlying asset at some date and price. The price is referred to as the exercise price or the strike price, while the date is called the expiration date or maturity. (Hull, 2017)

The owner of a call option benefits from the price increases of the underlying asset. In case of a prolonged downward price action, the owner's maximum loss is limited to the price of the option. Inversely, the holder of a put benefits from a decrease in price and can also lose at most the cost of the option in the case of an upsurge. (Hull, 2017)

The main difference between options and forward or futures contracts is that with options the holder is not required to exercise the right the contract gives. Unlike forward or futures contracts, options also come with a cost to acquire. For more information on options and derivatives, see Hull, 2017.

2.2 Asset price model

To model uncertainty, we assume that the underlying stock price $S(t)$ follows the Black–Scholes model

$$dS(t) = \mu S(t)dt + \sigma S(t)dW(t), \quad (2)$$

where μ is the drift (expected rate of return), $\sigma > 0$ is the volatility and $W(t)$ represents a standard Brownian motion. Volatility is a measure of uncertainty about the asset's returns. It corresponds to the standard deviation of returns over a time period. The relationship between volatility and the standard deviation over time is defined as

$$\sigma = \frac{\sigma_{SD}}{\sqrt{t}}, \quad (3)$$

where σ_{SD} is the standard deviation of returns during the time t . (Holtz, 2011)

To describe the risk-neutral dynamics of the stock price, the key to the effective option pricing, parameter μ is set equal to the risk-free rate of return r . The stochastic differential equation (2) has the solution given by

$$S(t) = S(0)e^{\left(r - \frac{\sigma^2}{2}\right)t + \sigma W(t)}, \quad (4)$$

where $S(0)$ is the present stock price and $W(t)$ is the standard Brownian motion. Since $W(t) \sim N(0, t)$ follows that $W(t) = \sqrt{t}Z$, where $Z \sim N(0, 1)$. Thus, the stock price at time t can be represented as

$$S(t) = S(0)e^{\left(r - \frac{\sigma^2}{2}\right)t + \sigma\sqrt{t}Z}. \quad (5)$$

Therefore, the logarithm of the stock price is normally distributed meaning that the distribution of the stock price is lognormal. This particular equation is useful in Monte Carlo simulations, where stock price paths are generated using standard normal random variables. (Glasserman, 2004)

Brownian motion $\{W(t_n)\}_{n \geq 0}$ is a Markov process that can be discretized recursively as follows

$$W(t_n) = W(t_{n-1}) + \sqrt{\Delta t}Z_n, \quad Z_n \sim N(0, 1), \quad (6)$$

where Z_n represents standard normal random variables and the initial value is $W(0) = 0$ (MörTERS et al., 2010). This Brownian motion construction method, often referred to as the random walk, is the most commonly used and simplest to implement. The path is generated by drawing one normally distributed random number after another, so that each number determines the next step independently (Holtz, 2011). In this thesis, Brownian motion is simulated using this random walk approximation.

2.3 European call Options

European options can be exercised solely at maturity. Call options entitle the holder to buy the underlying asset at a strike price K at a specified maturity T (measured in years). Commonly, European options are characterized in terms of the payoff to the purchaser of the options. The payoff formula of a European call does not account for the premium paid:

$$\max(S(T) - K, 0). \quad (7)$$

Here, $S(T)$ is the asset price at maturity and K denotes the strike price. If $S(T) > K$, the holder exercises the option and purchases the asset below market value, resulting in a payoff

of $S(T) - K$. Conversely, if $S(T) \leq K$, the option has no intrinsic value at expiration and it is not optimal to exercise it. (Hull, 2017)

The payoff can be viewed as a measure of the feasibility to exercise the call option. If payoff has a positive value, exercising the option at least covers part of the costs of acquiring the option. In the case where payoff exceeds the option premium, the investment is profitable overall.

Since the losses are limited to the premium paid for the option and the gains increase as the asset price rises above the strike price, the payoff structure is nonlinear. Nonlinearity is one of the primary reasons why numerical methods such as MC simulation are useful in option pricing.

The payoff of the call option has now been defined. The next step is to determine its current value at time $t = 0$. To get the present value of this payoff, the payoff is multiplied by e^{-rT} , the discount factor corresponding to a continuously compounded risk-free interest rate r . The European call option's (EC) expected present value at time $t = 0$ is given by

$$EC(0) = e^{-rT} \mathbb{E}[\max(S(T) - K, 0)], \quad (8)$$

where \mathbb{E} is the expectation under the risk-neutral measure. This representation leaning on expectation is fundamental for MC and QMC methods, since these methods estimate option prices by numerical approximation of expectations. (Glasserman, 2004)

For European call options there exists a closed-form analytical solution to the differential equation (5) called the Black–Scholes–Merton formula. The analytical option price is (Hull, 2017):

$$EC(0) = S(0)\Phi(d_1) - Ke^{-rT}\Phi(d_2), \quad (9)$$

where

$$d_1 = \frac{\ln\left(\frac{S(0)}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}, \quad d_2 = d_1 - \sigma\sqrt{T}. \quad (10)$$

Here Φ denotes the cumulative distribution function (CDF) of the standard normal distribution, defined by (Linde, 2016):

$$\Phi(x) = \int_{-\infty}^x \frac{e^{-\frac{1}{2}t^2}}{\sqrt{2\pi}} dt. \quad (11)$$

This particular formula (9) serves as an exact benchmark for European call option prices under the Black–Scholes model assumptions. In this thesis, it serves as a reference value against which numerical estimates obtained using MC and QMC methods can be compared.

Since the price of a European call option can be written as an expectation (see equation 8), MC methods can be used to numerically approximate said prices. The underlying idea is to simulate a large number of theoretical asset prices at maturity $S(T)$, calculate the corresponding option payoffs and then take the average discounted payoff. The Monte Carlo estimator for the European call option price is:

$$\hat{C}^{MC} = e^{-rT} \frac{1}{n} \sum_{i=1}^n \max(S_i(T) - K, 0), \quad (12)$$

where n corresponds to the number of simulated asset prices and $S_i(T)$ signifies the i -th simulated asset price at maturity T (Glasserman, 2004). This representation provides the foundation and enables the use of MC and QMC methods in option pricing, introduced in Chapter 3.

2.4 Asian call Options

European option payoffs depend only on the asset price at maturity, whereas Asian options are path-dependent derivatives whose payoffs are determined by the average price of the underlying asset over the option's lifetime (Hull, 2017). Asian options can be categorized according to the averaging method used. Two common forms are arithmetic average Asian options and geometric average Asian options. The difference between these forms is the way the payoff is calculated. In arithmetic average Asian options the payoff is calculated using arithmetic average (general way of calculating averages), and in discretely monitored geometric average case the payoff of an Asian call AC is:

$$AC = \max \left(\left(\prod_{n=1}^d S(t_n)^{\frac{1}{d}} \right) - K, 0 \right), \quad (13)$$

where d is the number of observation dates, $S(t_n)$ is the price of the asset at date t_n and K is the strike price. There does not exist a closed-form solution for Asian options with arithmetic average, thus they will not be considered in the numerical experiments. However, geometric average Asian options do have a closed-form solution:

$$AC(0) = S(0)\gamma\Phi(\beta + \sigma\sqrt{T_1}) - Ke^{-rT}\Phi(\beta), \quad (14)$$

where

$$T_1 = T - \frac{d(d-1)(4d+1)}{6d^2} \Delta t, \quad (15)$$

$$T_2 = T - \frac{(d-1)}{2} \Delta t, \quad (16)$$

$$\gamma = e^{-r(T-T_2) - \sigma^2 \frac{(T_2-T_1)}{2}}, \quad (17)$$

$$\beta = \frac{\ln\left(\frac{S(0)}{K}\right) + \left(r - \frac{1}{2}\sigma^2\right) T_2}{\sigma\sqrt{T_1}}, \quad (18)$$

which will be used in this thesis as a benchmark to evaluate the accuracy and convergence of the simulation methods used. (Holtz, 2011)

The expected present value for a geometric average Asian call option can be calculated similarly to the European call option:

$$AC(0) = e^{-rT} \mathbb{E} \left[\max \left(\left(\prod_{n=1}^d S(t_n) \right)^{\frac{1}{d}} - K, 0 \right) \right]. \quad (19)$$

The difference is that to calculate the payoff, intermediate dates t_n are also considered, which differs from European call option considering only the current price and price at maturity T of the underlying asset. (Glasserman, 2004)

Since the payoff of an Asian call option is dependent on the evolution of the stock price $S(t)$ between now and maturity, it is necessary to simulate asset price paths over multiple dates between those points. All observation dates contribute to the payoff, which means that the Asian option pricing problem has a higher dimension than that of a European. Monte Carlo methods are seen as more attractive in approximating higher dimensional integrals, which makes comparing the option pricing of geometric average Asian options with MC and QMC interesting. (Glasserman, 2004)

3 Numerical methods

This chapter introduces the numerical methods used in this thesis for option pricing. First, the standard MC method is presented as a reference. This is followed by the QMC method based on low-discrepancy sequences as an alternative approach to numerical integration. Pseudocode representations of the methods are shown in the respective chapters. Finally, the visual difference of simulated price paths is provided and examined.

3.1 Monte Carlo

The integral of f over the unit interval can be written as

$$I(f) = \int_0^1 f(x) dx. \quad (20)$$

The integral I could be estimated by expressing it as an expectation $\mathbb{E}[f(X)]$, with X uniformly distributed between the limits of the integral $[0, 1]$. By drawing points X_1, \dots, X_n independently and identically distributed (i.i.d.) from the unit interval, evaluating f at n randomly sampled points and averaging the results yields the MC estimate

$$\hat{I}_n(f) = \frac{1}{n} \sum_{i=1}^n f(X_i). \quad (21)$$

In multiple dimensions the interval is a d -dimensional unit cube and the integral is the form:

$$I_d(f) = \int_0^1 \cdots \int_0^1 f(x_1, \dots, x_d) dx_1 \cdots dx_d. \quad (22)$$

Now, the MC estimate is of the form

$$\hat{I}_{n,d}(f) = \frac{1}{n} \sum_{i=1}^n f(\mathbf{X}_i), \quad (23)$$

where $\mathbf{X}_1, \dots, \mathbf{X}_n$ are now uniform random sample vectors (instead of scalars) from $[0, 1]^d$ (Dick, Kuo & Sloan, 2013). If f is integrable on $[0, 1]$, the law of large numbers indicates that $\hat{I}_n \rightarrow I$ with probability of 1 as $n \rightarrow \infty$. The same property holds for arbitrarily many dimensions. (Glasserman, 2004)

Following from the property of the samples being i.i.d., the expectation of the estimator satisfies the following equality

$$\mathbb{E}[\hat{I}_{n,d}(f)] = I_d(f), \quad (24)$$

which leads to the MC method being unbiased. The estimator's variance is then

$$\text{Var}[\hat{I}_{n,d}(f)] = \frac{\sigma^2(f)}{n}, \quad (25)$$

where

$$\sigma^2(f) = I_d(f^2) - (I_d(f))^2. \quad (26)$$

A proof of the theorem that leads to the above mentioned results is presented in for example Dick, Kuo & Sloan, 2013. The variance decreases proportionally to $\frac{1}{n}$, thus the standard deviation $\sigma(f)$ (the square root of variance) of the estimator decreases proportionally to $\frac{1}{\sqrt{n}}$.

According to the central limit theorem, the distribution of the MC estimator approaches a normal distribution as the number of simulations approaches infinity, provided that $0 < \sigma(f) < \infty$. In a form of equation it is expressed as

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(|I_d(f) - \hat{I}_{n,d}(f)| \leq c \frac{\sigma(f)}{\sqrt{n}} \right) = \frac{1}{\sqrt{2\pi}} \int_{-c}^c e^{-\frac{1}{2}x^2} dx. \quad (27)$$

In consequence, the MC estimate has an error with a convergence rate $O(n^{-\frac{1}{2}})$. This convergence rate holds regardless of the number of dimensions and is a fundamental property of MC. As a result, high-dimensional problems such as path-dependent option pricing are well suitable for the MC method. However, the convergence rate is relatively slow, especially if variance reduction methods are not used. To reduce the error by a factor of two, approximately four times as many simulations are required. (Dick, Kuo & Sloan, 2013; Glasserman, 2004)

In financial applications, including option pricing, the quantities of interest are often expectations with respect to standard normal multivariate random variables. For a function f of d variables, this expectation can be written as a multidimensional Gaussian integral:

$$I_d(f) = \int_{\mathbb{R}^d} f(x_1, \dots, x_d) \prod_{i=1}^d \frac{e^{-\frac{1}{2}x_i^2}}{\sqrt{2\pi}} dx_1 \cdots dx_d. \quad (28)$$

In the implementation of MC introduced next, the integral is approximated by sampling independent standard normal random variables using a pseudorandom number generator and computing the corresponding sample average of f . (Glasserman, 2004)

In Algorithm 1 the generation of asset price paths and option pricing based on payoffs is given. First, the time to maturity is partitioned into time steps, and these steps are used to create a vector. Independent standard normal variables are then generated and used to create

Brownian motion paths. Then, the risk-neutral prices of the asset are calculated. Risk-neutral prices are used to calculate the geometric averages, which in turn are used to calculate the payoffs. Finally, the MC estimator for the fair Asian call option price is derived from the average of payoffs.

Algorithm 1: MC estimator for geometric Asian call option

Input: Current asset price S_0 , volatility σ , risk-neutral return r , time to maturity T , strike price K , number of time steps d , number of simulations n

Output: MC estimate of option price \hat{C}^{MC} , payoffs POs

- 1: $\Delta t \leftarrow \frac{T}{d}$
 - 2: $\mathbf{t} \leftarrow [\Delta t, 2\Delta t, \dots, d\Delta t]$
 - 3: Generate $\mathbf{Z} = [Z_{i,j}]_{\substack{1 \leq i \leq n \\ 1 \leq j \leq d}} \sim N(0, 1)$ // Independent standard normal variables
 - 4: $\mathbf{W} \leftarrow \text{cumsum}(\sqrt{\Delta t} \mathbf{Z})$ // Brownian motion paths
 - 5: $\mathbf{S} \leftarrow S_0 \exp\left(\left(r - \frac{\sigma^2}{2}\right)\mathbf{t} + \sigma \mathbf{W}\right)$ // Row-wise sum of matrix and vector
 - 6: **for** $i = 1$ **to** n **do**
 - 7: $G_i \leftarrow \left(\prod_{j=1}^d S_{i,j}\right)^{\frac{1}{d}}$
 - 8: $POs_i \leftarrow \max(G_i - K, 0)$
 - 9: $\hat{C}^{MC} \leftarrow e^{-rT} \frac{1}{n} \sum_{i=1}^n POs_i$ // Discounted payoff average
-

For the European call option, the structure of the algorithm is simpler. Instead of simulating an entire price path over multiple time steps, only the terminal asset price is required. Following this, the dimension reduces to one and the payoff is computed directly from the terminal price as in equation (7). The need to compute geometric averages and time discretization is removed, while the algorithm remains the same otherwise.

3.2 Quasi-Monte Carlo

QMC methods estimate integrals in a similar manner to MC methods. However, instead of pseudorandom numbers, deterministic point sequences, designed to form a greater uniformity, are used. The following is a 2D example figure that illustrates the difference of sampling between these two methods. MC generated pseudorandom points tend to cluster, whereas quasi-random points are distributed rather evenly. (Holtz, 2011)

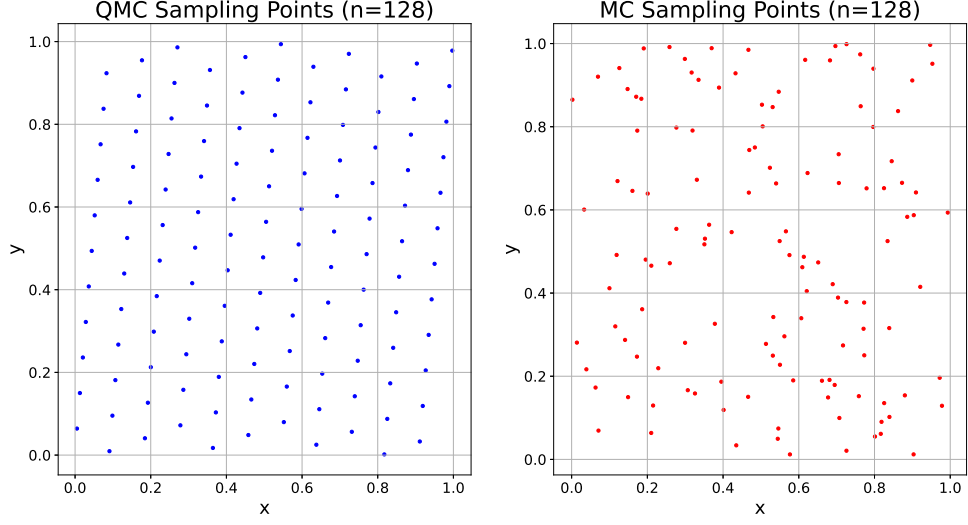


Figure 1: 128 sampled points with MC & QMC

Discrepancy measures play a key role in bounding the approximation error in (23). The Koksma–Hlawka inequality is a known result bounding the integration error by the product of the variation V of the integrand f and the star discrepancy D^* of the point set $\mathbf{x}_1, \dots, \mathbf{x}_n \in [0, 1]^d$ (Glasserman, 2004):

$$\left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{x}_i) - \int_{[0,1]^d} f(\mathbf{x}) \, d\mathbf{x} \right| \leq V(f) D^*(\mathbf{x}_1, \dots, \mathbf{x}_n). \quad (29)$$

The variation of f by Hardy and Krause is defined as:

$$V(f) = \sum_{k=1}^d \sum_{1 \leq i_1 < \dots < i_k \leq d} V^{(k)}(f; i_1, \dots, i_k), \quad (30)$$

where $V^{(k)}(f; i_1, \dots, i_k)$ presents the k -dimensional variation of f with respect to the coordinates i_1, \dots, i_k . The discrepancy of the point set $\mathbf{x}_1, \dots, \mathbf{x}_n$ relative to a collection \mathcal{A} of Lebesgue measurable subsets of $[0, 1]^d$ is defined as:

$$D(\mathbf{x}_1, \dots, \mathbf{x}_n; \mathcal{A}) = \sup_{A \in \mathcal{A}} \left| \frac{\#\{\mathbf{x}_i \in A\}}{n} - \text{vol}(A) \right|, \quad (31)$$

where $\#\{\mathbf{x}_i \in A\}$ is the number of points \mathbf{x}_i in A and $\text{vol}(A)$ is the volume of A . Letting \mathcal{A} be the collection of rectangles of the form:

$$A = \prod_{j=1}^d [0, u_j], \quad u_j \in [0, 1]. \quad (32)$$

The notation $[0, \mathbf{u}]$ denotes the Cartesian product:

$$[0, \mathbf{u}] := [0, u_1] \times \cdots \times [0, u_d], \quad \mathbf{u} \in [0, 1]^d. \quad (33)$$

Then the star discrepancy satisfies:

$$D^*(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sup_{\mathbf{u} \in [0, 1]^d} \left| \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{[0, \mathbf{u}]}(\mathbf{x}_i) - \prod_{j=1}^d u_j \right| = \mathcal{O} \left(\frac{(\log n)^d}{n} \right). \quad (34)$$

Here $\mathbf{1}_{[0, \mathbf{u}]}(\mathbf{x})$ is the indicator function defined as:

$$\mathbf{1}_{[0, \mathbf{u}]}(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} \in [0, \mathbf{u}] \\ 0, & \text{otherwise.} \end{cases} \quad (35)$$

Thus for functions of bounded Hardy–Krause variation, the QMC integration error is:

$$\left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{x}_i) - \int_{[0, 1]^d} f(\mathbf{x}) d\mathbf{x} \right| = \mathcal{O} \left(\frac{(\log n)^d}{n} \right). \quad (36)$$

Since $V(f)$ is fixed for a given integrand and does not depend on n , it is omitted in the big \mathcal{O} -notation. The logarithmic term grows more slowly than any positive power of n , allowing a looser bound:

$$\frac{(\log n)^d}{n} = \mathcal{O} \left(\frac{1}{n^{1-\varepsilon}} \right) \quad \forall \varepsilon > 0. \quad (37)$$

Discrepancy is the measure of irregularity of QMC points distribution in the unit cube. There exists a wide range of discrepancy measures and approaches to generate low-discrepancy sequences. These sequences are based mostly on digital nets or on lattices. Popular QMC methods include Faure, Sobol and Halton together with lattice rules grounded in Korobov rules or component-by-component (CBC) algorithms. (Holtz, 2011)

The sequences used in the numerical experiments of this thesis are based on lattice rules. The generating vectors for the lattice rules are obtained from the publicly available lattice rule collection maintained by Kuo, 2007. The vectors are constructed using CBC algorithms designed for low-discrepancy lattice point sets suitable for high-dimensional integration problems.

The generating vectors used are designed for rank-one lattice rules. The n -point rank-one lattice rule in d dimensions is defined by:

$$\mathbf{x}_i = \left\{ \frac{i\mathbf{g}}{n} \right\}, \quad i = 1, 2, \dots, n, \quad (38)$$

where $\mathbf{g} \in \mathbb{Z}^d$ is the generating vector whose components share no common factor with n , and the braces indicate fractional part of all components of the vector (Dick, Kuo & Sloan, 2013). For non-periodic integrands, randomly shifted lattice rules are often used. In this approach:

$$\mathbf{x}_i = \left\{ \frac{i\mathbf{g}}{n} + \Delta \right\}, \quad (39)$$

the generation of lattice points includes a uniformly distributed random shift $\Delta \in [0, 1)^d$ (Holtz, 2011).

The lattice points are now approximately uniform in $[0, 1)^d$. The same multidimensional integral (28) can also be estimated using QMC methods. To use these lattice points in the construction of Brownian motion, they are transformed into standard normal variables. Φ denotes the CDF of the standard normal distribution, defined in equation (11). To transform uniform lattice points in the d -dimensional unit cube to standard normal variables, the following transformation is done:

$$\mathbf{Z}_i = \Phi^{-1}(\mathbf{x}_i), \quad \mathbf{Z}_i \sim N(0, 1), \quad (40)$$

where Φ^{-1} is the inverse of the standard normal CDF (Glasserman, 2004). Now the integral is of the form:

$$I_d(f) = \int_{[0,1]^d} f(\Phi^{-1}(x_1), \dots, \Phi^{-1}(x_d)) dx_1 \cdots dx_d. \quad (41)$$

Although this transformation results in an unbounded integrand at the boundaries 0 and 1, the transformation is highly effective as it eliminates the Gaussian weight (Holtz, 2011).

In Algorithm 2 the QMC estimator using lattice rules for geometric Asian call option is given. The discretization of time steps, the calculation of the risk-neutral price, and the payoffs is implemented similarly to Algorithm 1. The difference is that instead of pseudorandom standard normal variables, deterministic low-discrepancy lattice points are used. First, a generating vector obtained from (Kuo, 2007) is loaded for the lattice rule. A random shift is then applied to the QMC points in the unit cube $[0, 1]^d$. The shifted lattice points are then transformed into standard normal variables. After the transformation, the variables are

used to construct Brownian motion paths and simulate the asset price line diagrams as in Algorithm 1.

Algorithm 2: QMC estimator using lattice rules for geometric Asian call option

Input: Current asset price S_0 , volatility σ , risk-neutral return r , time to maturity T , strike price K , number of time steps d , number of simulations n

Output: QMC estimate of option price \hat{C}^{QMC} , payoffs POs

```

1:  $\Delta t \leftarrow \frac{T}{d}$ 
2:  $\mathbf{t} \leftarrow [\Delta t, 2\Delta t, \dots, d\Delta t]$ 
3:  $\mathbf{g} \leftarrow \text{LoadLatticeVector}(d)$  // Generating vector
4:  $\Delta \sim U([0, 1]^d)$  // Random shift
5: for  $i = 1$  to  $n$  do
6:    $\mathbf{x}_i \leftarrow \left( \frac{i\mathbf{g}}{n} + \Delta \right) \bmod 1$ 
7:    $\mathbf{Z}_i \leftarrow \Phi^{-1}(\mathbf{x}_i)$  // Inverse normal transform
8:    $\mathbf{W} \leftarrow \text{cumsum}(\sqrt{\Delta t}\mathbf{Z})$ 
9:    $\mathbf{S} \leftarrow S_0 \exp\left(\left(r - \frac{\sigma^2}{2}\right)\mathbf{t} + \sigma\mathbf{W}\right)$  // Row-wise sum of matrix and vector
10: for  $i = 1$  to  $n$  do
11:    $G_i \leftarrow \left(\prod_{j=1}^d S_{i,j}\right)^{1/d}$ 
12:    $POs_i \leftarrow \max(G_i - K, 0)$ 
13:  $\hat{C}^{QMC} \leftarrow e^{-rT} \frac{1}{n} \sum_{i=1}^n POs_i$ 

```

For the European call option, the QMC algorithm simplifies in a similar manner as in the MC case. The asset price is again evaluated only at maturity, reducing the dimension of the integral to one. As a consequence, the lattice lattice points are used to form a single uniform variate per simulation, which is then transformed into a standard normal variable as in Algorithm 2. The payoff is calculated directly following equation (7) and using the terminal asset value.

3.3 Simulated asset price paths

This section illustrates simulated asset price paths using MC and QMC methods. The figures presented provide a visual comparison of the sampling structures possessed by the two approaches before analyzing their convergence behavior and simulated payoff structure.

The parameters used in the following figures are the same as in the experiments presented in Chapter 4. Although numerical experiments use larger sample sizes, only 2^7 price paths were shown to conserve visual clarity. The price paths are those used in the calculation of an Asian call option price. The time frame from now to maturity is divided into 16 time steps, and the price paths are line diagrams. In European options, the simulated paths would be line segments from the initial price $S(0)$ to the price at maturity $S(T)$.

The black dotted line denoted “Breakeven Price” in the legends of the plots is the same as the initial price $S(0)$. In the calculations of the option prices, the cases where $S(T)$ is below the set strike price K (set at 100 in the experiments), the payoff equals 0.

In Figure 2 the asset prices at maturity range approximately from 67 to 180. Seven simulated paths end above the price level of 160, while five terminate below 80. The concentration of paths increases around the strike price level. Overall, the distribution of the terminal asset prices resembles the shape of a lognormal probability density function, with a mean slightly above the initial price. The finding is consistent with the theory behind the underlying asset price introduced in Chapter 2.2.

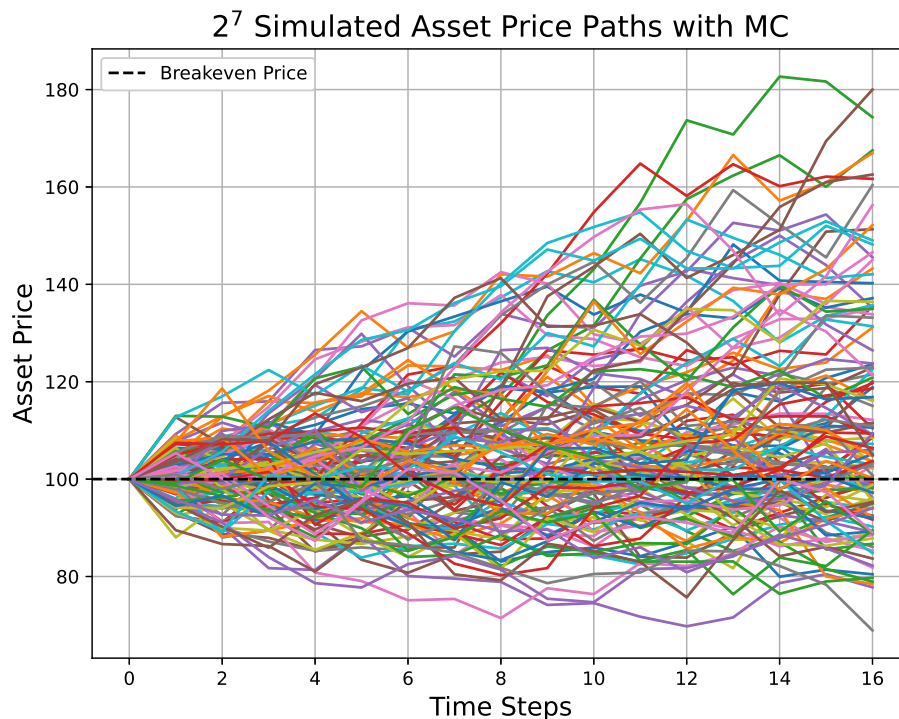


Figure 2: Asset price paths simulated with MC

In Figure 3 terminal asset prices range approximately from 60 to 183. Only path crossing the price levels of 160 and 180 is the highest asset price in the figure. Nine paths end below the price level of 80. Although the most concentrated level of the paths is at the initial price level, the amount of concentration appears to be uniform from 100 to 140 and 90 to 100. The distribution of the terminal asset prices in this figure resemble the shape of a lognormal probability density function with less variance than in Figure 2.

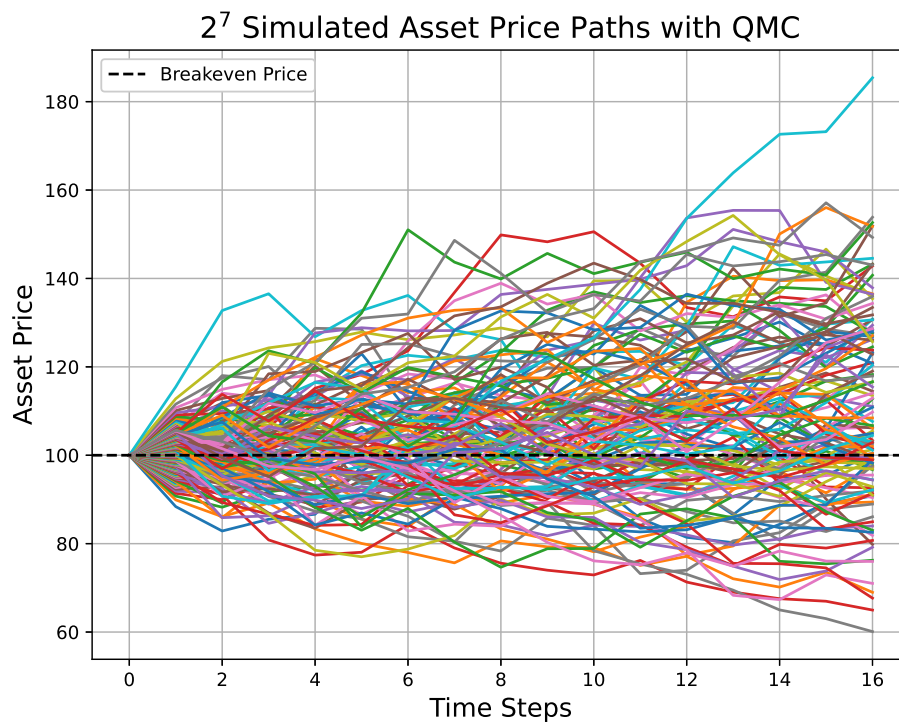


Figure 3: Asset price paths simulated with QMC

The illustrations help explain why the convergence of MC methods is typically slower than that of QMC methods. When the sample space is sampled deterministically and more uniformly, the clustering effect of random sampling is reduced. This leads to lower variance and a more even coverage of the integration domain.

4 Numerical experiments

The purpose of the numerical experiments is to demonstrate how the results of the MC and QMC methods differ from each other in simulated payoffs, and in convergence. Convergence graphs are in log-log scale and show the errors of the estimators from MC and QMC simulations as data points and the fits as least squares estimates (LSE). QMC's convergence versus the number of time steps is shown in a separate plot. Payoff distributions are visualized as histograms. The results are analyzed numerically and by using a table of statistical quantities of payoffs.

Since the theoretical convergence rates of MC and QMC are $\mathcal{O}(n^{-\frac{1}{2}})$ and essentially $\mathcal{O}(n^{-1})$ respectively, the errors of the estimators satisfy a power rule $err = an^b$ where a is some constant and b is the convergence rate exponent. When represented using a log-log plot, the decay rate indicated by b corresponds to the slope of the data. To fit $err = an^b$ to data $(n_1, err_1), \dots, (n_m, err_m)$, the first step is to linearize the model by taking the logarithm of both sides:

$$\begin{cases} err_1 = an_1^b \\ err_2 = an_2^b \\ \vdots \\ err_m = an_m^b \end{cases} \Leftrightarrow \begin{cases} \ln err_1 = \ln a + b \ln n_1 \\ \ln err_2 = \ln a + b \ln n_2 \\ \vdots \\ \ln err_m = \ln a + b \ln n_m. \end{cases} \quad (42)$$

In consequence this can be expressed as a system:

$$\begin{bmatrix} \ln err_1 \\ \vdots \\ \ln err_m \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & \ln n_1 \\ \vdots & \vdots \\ 1 & \ln n_m \end{bmatrix}}_{=:A} \begin{bmatrix} \ln a \\ b \end{bmatrix}. \quad (43)$$

The fit can then be obtained by solving the following equation:

$$A^T A \begin{bmatrix} \ln a \\ b \end{bmatrix} = A^T \begin{bmatrix} \ln err_1 \\ \vdots \\ \ln err_m \end{bmatrix}. \quad (44)$$

For more on data fitting and numerical computation, see Press et al., 2007.

The experiments were conducted in Python 3.10.6 using three libraries. The NumPy 2.2.6 library was used for array creation, generating pseudorandom numbers and mathematical

operators such as logarithms. Reproducibility was ensured with `numpy.random.seed(42)` to initialize the pseudorandom number generator. Statistical library Scipy's version 1.15.3 `scipy.stats` object `norm` and its methods `cdf` and `ppf` were used in sampling from both the CDF and the percent point function (inverse CDF) of the normal distribution. Matplotlib's version 3.10.9 built-in plotting functions were used in all of the plots.

The following parameters were used in all the numerical experiments:

- $S_0 = 100$
- $\sigma = 0.2$
- $r = 0.1$
- $T = 1$
- $K = 100$

The same parameters were used in (Holtz, 2011), where risk-neutral prices of Asian options were calculated with different quadrature and path generating methods. While payoffs were not analyzed, the convergence rates of standard MC and QMC were compared. Using the same Brownian motion construction method as in this thesis, the random walk, the convergence rate of QMC was the closest to that of MC. Using other construction methods covered in (Holtz, 2011), QMC achieved a less oscillatory convergence rate of approximately 1.

Figure 4 shows the payoff distributions for the European call option obtained using 2^{20} MC and QMC simulations. In both histograms, the majority of simulated payoffs is concentrated near zero, reflecting the probability that the option will end up unprofitable under the chosen parameters. A smaller fraction of simulations end up generating strictly positive payoffs.

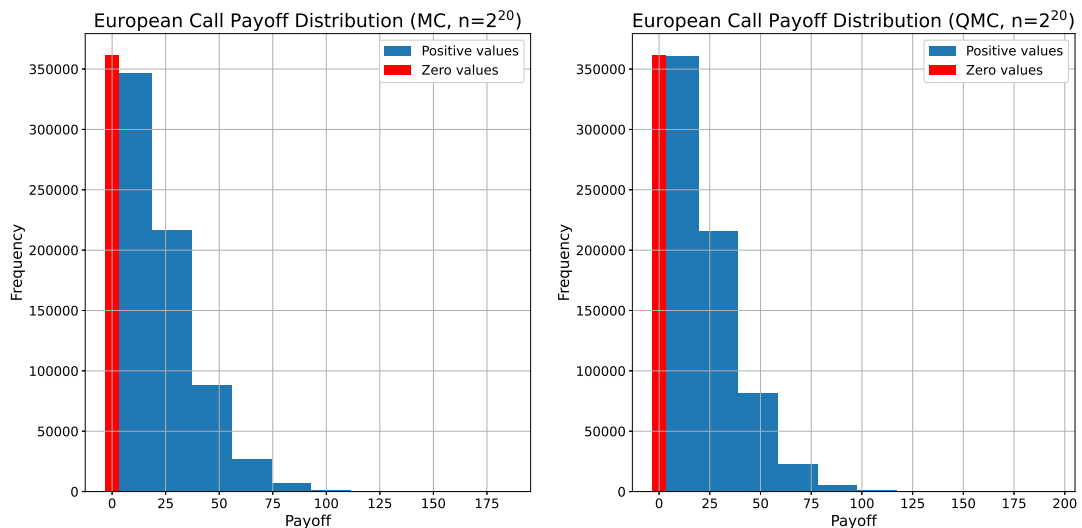


Figure 4: Payoff distribution histograms of MC & QMC with European call options

Both MC and QMC methods produce similar expected payoff levels, indicating consistency in the estimation of the mean. After the payoff level of 100 the frequency of payoffs becomes visually non-existent. However, when carefully examined, the QMC distribution appears slightly more concentrated compared to the MC alternative.

The vertical red bar highlights zero payoffs, which is the highest region of both distributions. This emphasizes that the option frequently expires worthless. Despite this, the determining factor of the option price is the positive tail of the distribution.

Figure 5 shows the simulated payoff distributions of the geometric Asian call option obtained using MC and QMC methods with 2^{20} simulations. The histograms illustrate the frequency distribution of payoffs, including zero and positive values.

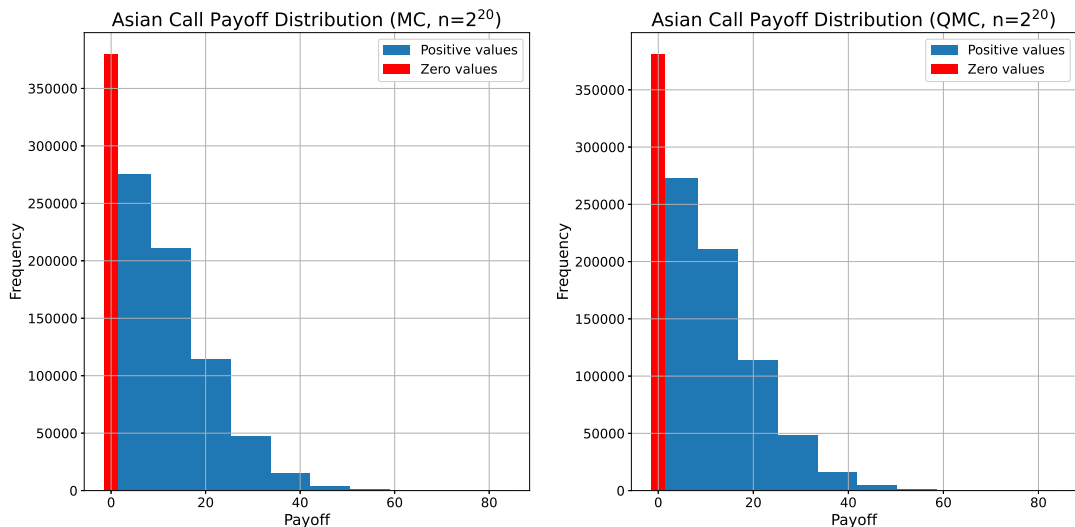


Figure 5: Payoff distribution histograms of MC & QMC with Asian call options

In the case of Asian call option, MC and QMC again produce highly similar payoff distributions, further strengthening the consistency in the expected value estimation. As in the European case, a large proportion of the simulated payoffs are concentrated at zero. The payoff value range is noticeably smaller than that of its European counterpart, which is intuitive based on the geometric averaging factor of Asian calls.

The distribution is strongly right-skewed, with a relatively small number of simulations generating large payoffs. These rare but significant outcomes greatly affect the calculation of option price's expectation. Compared to the MC method, the QMC distribution appears to include a higher concentration of high positive payoff values. The difference is subtle, which indicates that both methods converge to a similar underlying payoff distribution for this problem. The large number of simulations could be an explanatory factor.

In figure 6, the estimation errors relative to the Black–Scholes benchmark solution of the MC and QMC estimators are plotted as a function of the number of simulations. The simulations are carried out for sample sizes ranging from 2^0 to 2^{20} , increasing in powers of two. For each sample size, the simulation is repeated 20 times, and the plotted data points are averages of the repetitions. LSE is used to form a regression line to best fit the data points.

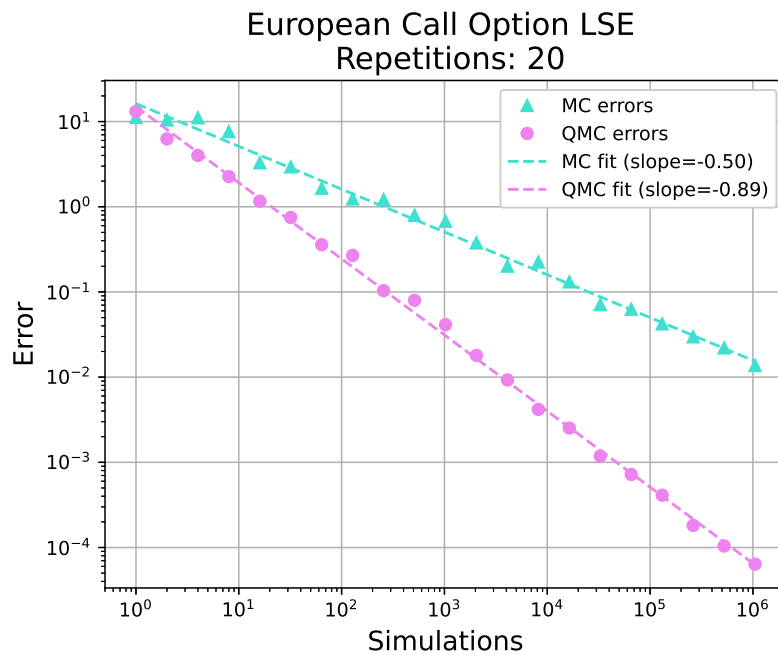


Figure 6: The convergence rates of MC & QMC with European call options

A clear difference between the convergence rates of the two methods can be observed. The MC estimator exhibits an approximately linear relationship on the log-log scale with a LSE slope of approximately -0.50 , which is right at the theoretical convergence rate of standard MC presented in Chapter 3.1. This result is consistent with the theory, where the standard deviation of the estimator decreases proportionally to the inverse square root of the number of simulations.

The QMC estimator demonstrates significantly faster convergence. The fitted slope for QMC is approximately -0.89 , which is close to the ideal of essentially $\mathcal{O}(n^{-1})$. The steeper slope indicates that the pricing error decreases substantially as the number of simulations increases.

At low simulation counts, the errors of MC and QMC are relatively similar. As the number of simulations grows, the gap between the methods widens considerably. For example, at simulation sizes above 10^4 , the QMC method achieves errors over an order of magnitude smaller than the standard MC. At 10^6 simulations the difference is approximately two orders of magnitude, with QMC error being less than 10^{-4} and MC error more than 10^{-2} .

Figure 7 presents the convergence behavior of MC and QMC methods in the pricing of an Asian call option. Similarly to figure 6 the horizontal axis shows the number of simulations on a logarithmic scale, while the vertical axis presents the pricing error relative to the analytical benchmark value. The number of time steps used during the simulations is again constant at $d = 16$. The reported errors are averages over 20 independent repetitions.

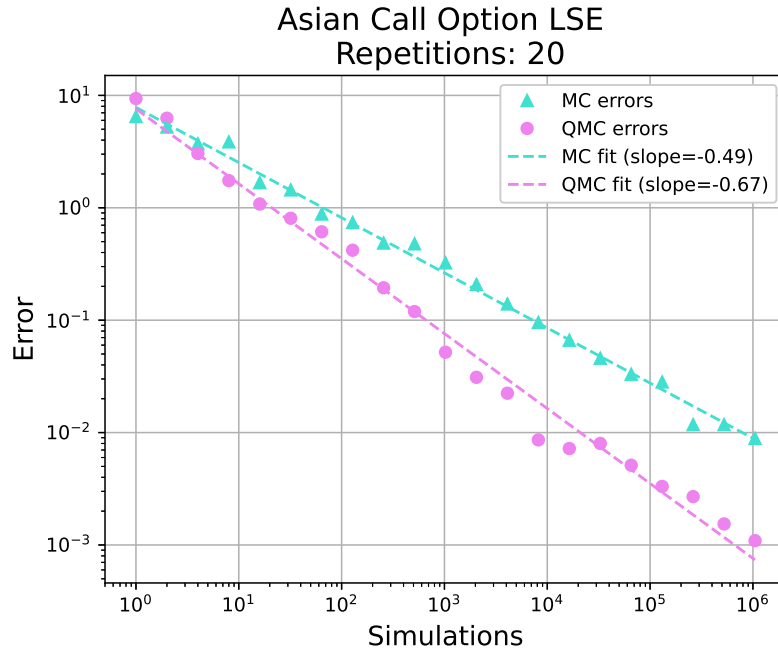


Figure 7: The convergence rates of MC & QMC with Asian Call options

The MC method exhibits a fitted convergence slope of approximately -0.49 , which is again very close and in line with the theoretical convergence rate of MC. This result is consistent with the theoretical properties of MC methods, whose convergence rate is independent of the dimensionality of the integration problem.

The QMC method once more demonstrates a faster convergence than MC, with an observed slope of approximately -0.67 . Although this is clearly superior to MC, the improvement is noticeably smaller than in the European option experiment, where the QMC slope was approximately -0.89 . This indicates that the advantage of QMC decreases as the dimensionality of the problem increases.

The weaker performance improvement can be explained by the path-dependency of Asian options. Unlike European options, Asian options depend on the average asset price over multiple observation dates. This results in the pricing problem requiring simulation of the entire price path, which greatly increases the dimension of the numerical integration problem.

In higher-dimensional settings, low-discrepancy lattice point sets lose part of their uniformity advantages. This loss in uniformity contributes to the reduced convergence of QMC methods. Nevertheless, QMC still consistently outperforms MC throughout the wide simulation range. Lattice rule QMC methods seem to retain practical benefits even in moderately high-dimensional option pricing problems.

In figure 8 the convergence slopes of QMC method in the pricing of Asian call option with varying dimensions are shown. Both simulations and time steps used increase in powers of two with simulations ranging from 2^0 to 2^{20} , and the number of time steps ranging from 2^4 to 2^{10} .

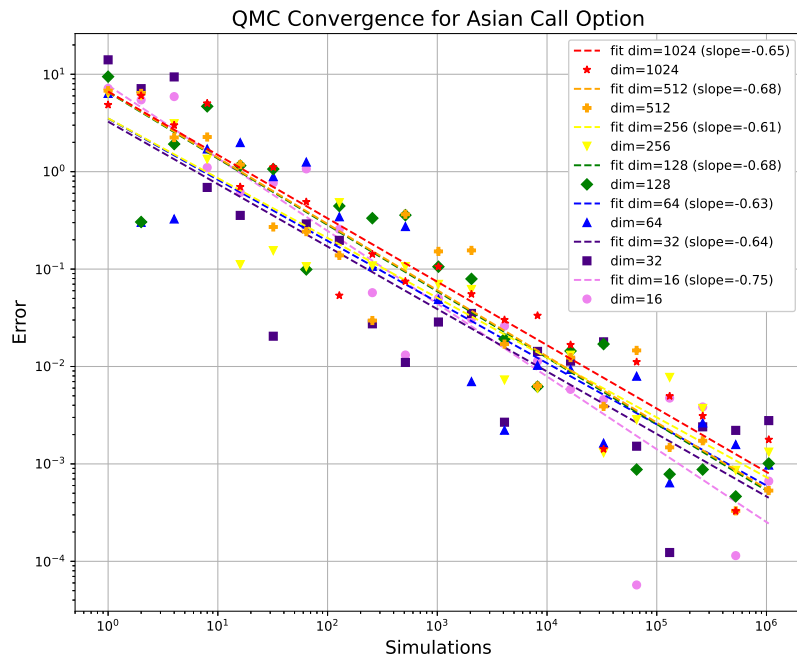


Figure 8: Convergence of QMC using varying dimensions with Asian call options

The convergence rate of QMC is the lowest at approximately -0.75 with a problem dimensionality of 16 and the highest at -0.61 using a dimensionality of 256. After dimensionality of 32 the convergence slopes seem to be very similar. Convergence slopes are sensitive to perturbations in the measured error values. Since slopes are estimated via LSE regression, fluctuations in error measurements may introduce variability in fitted convergence rates.

As the dimensionality increases, the convergence slope is higher and stays higher than in the lowest measured dimension. Nevertheless, even at dimensions as high as 1024, the observed convergence slope remains better than the performance of MC method. By measuring the

convergence at even higher dimensions, there might be a case where QMC has a higher convergence rate than MC. However, computation at such high dimensions could be extremely slow.

Figure 9 shows a line diagram of the convergence slopes as a function of the problem dimensions used in figure 8. The convergence rates of MC and QMC are shown as dotted lines as benchmarks. Although convergence performance weakens with increasing dimensionality, the QMC method consistently exhibits steeper convergence slopes than the MC method.




Figure 9: Convergence rates of QMC versus dimensions as a line diagram

Although some variability is present, no severe degradation in convergence is observed between dimensions 32 and 1024. As previously mentioned, the fluctuations in the estimated slopes are likely caused by perturbations in the error measurements coupled with a limited amount of samples used in the LSE regression. Results suggest that the QMC method maintains stable convergence behavior in moderately high dimensions, while preserving a steeper slope of convergence over standard MC.

Table 1 shows that the MC and QMC methods produce nearly identical sample statistics for Asian and European options. The mean and standard deviation of payoffs are virtually the same across methods, indicating that QMC does not alter the underlying payoff distribution but instead improves the sampling efficiency of the integration space. Minor differences in maximum payoffs are the result of the deterministic structure of low-discrepancy sequences and the finite sample size. It is highly likely that as the sample size grows, the differences between the maximum payoff values shrink.

Table 1: Statistical quantities of payoffs from 2^{20} simulations

Statistical quantity	Asian option	European option
MC Mean Payoff	7.89	14.67
QMC Mean Payoff	7.91	14.67
MC Payoff SD	9.60	17.80
QMC Payoff SD	9.59	17.80
MC Max Payoff	83.31	175.05
QMC Max Payoff	84.97	191.86

Overall, the results confirm that both MC and QMC generate statistically consistent payoff estimates for the options considered. The nearly identical sample mean and standard deviation suggests that the improved convergence behavior observed for QMC is due to a more efficient coverage of the integration domain rather than from changes in the payoff distribution itself.

5 Conclusions

The objective of this thesis was to examine the differences between the MC and QMC methods and to compare their performance in the pricing of European and Asian call options. The study also aimed to investigate how dimensionality affects the convergence rate of both methods. The fundamental difference between MC and QMC is the way the sample points are generated. In MC the points are independent random samples drawn from the normal distribution, whereas in QMC the points are deterministic sequences. The results indicate that the payoff structure produced by the methods is consistent and the higher convergence rate of QMC is due to a more uniform coverage of the integration domain.

MC's convergence rate remained virtually unchanged between European and Asian call options at nearly -0.5 , whereas QMC's decreased from approximately -0.89 to -0.67 . QMC is likely to perform better for European options due to the one-dimensional nature of the integration problem, while for Asian options its convergence appears sensitive to the chosen parameters through their effect on the effective dimension. QMC's convergence rate fluctuated between -0.61 and -0.68 after the dimensionality of 32, which may be explained by the linear least-squares estimation and floating-point errors arising from very small error values. Compared with (Holtz, 2011), the relatively weaker QMC performance for Asian options in this thesis is likely due to the use of the random walk method for the construction of Brownian motion, as alternative methods have been shown to improve the convergence rate.

The Black–Scholes option pricing model applied in this thesis is a standard model published in 1973. The model assumes constant interest rates and volatility, assumptions that often do not hold in real financial markets. Furthermore, only one Brownian motion construction method was considered, and no variance reduction techniques were used.

Future research could examine more advanced pricing models, such as the Heston model which involves stochastic volatility or jump-diffusion models, which take sudden large price movements into consideration. In addition, alternative Brownian motion construction methods could be explored to increase the convergence performance of QMC. Future research could also examine the use of variance reduction techniques in MC simulations and dimension reduction approaches for QMC methods.

References

- Case, G. (2025). Comparative Study of Monte Carlo and Quasi-Monte Carlo Techniques for Enhanced Derivative Pricing. arXiv: 2502.17731.
- Dick, J., Kuo, F. Y. & Sloan, I. H. (2013). High-dimensional integration: The quasi-Monte Carlo way. *Acta Numerica* 22, pp. 133–288.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Applications of Mathematics. New York: Springer.
- Holtz, M. (2011). *Sparse Grid Quadrature in High Dimensions with Applications in Finance and Insurance*. Lecture Notes in Computational Science and Engineering. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Hsiao, Y.-J. & Tsai, W.-C. (2018). Financial literacy and participation in the derivatives markets. *Journal of Banking & Finance* 88, pp. 15–29.
- Hull, J. C. (2017). *Options, Futures, and Other Derivatives*. Harlow: Pearson Education.
- Kuo, F. Y. (2007). *Lattice rule generating vectors*. University of New South Wales, accessed 22 May 2026. URL: <https://web.maths.unsw.edu.au/~fkuo/lattice/>.
- L'Ecuyer, P. (2009). Quasi-Monte Carlo methods with applications in finance. *Finance and Stochastics* 13(3), pp. 307–349.
- Linde, W. (2016). *Probability Theory: A First Course in Probability Theory and Statistics*. 1st edition. De Gruyter Textbook. Berlin: De Gruyter.
- Mörters, P., Peres, Y., Schramm, O. & Werner, W. (2010). *Brownian Motion*. Cambridge Series on Statistical and Probabilistic Mathematics. Cambridge: Cambridge University Press.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (2007). *Numerical Recipes: The Art of Scientific Computing*. 3rd ed. Cambridge University Press. URL: <https://numerical.recipes/book.html>.
- Scoleri, S., Bianchetti, M. & Kucherenko, S. (2021). Application of Quasi Monte Carlo and Global Sensitivity Analysis to Option Pricing and Greeks: Finite Differences vs. AAD. *Wilmott* 2021(116), pp. 66–83.