Proceedings of the 5th International Conference on Financial Technology and Business Analysis (ICFTBA 2026)
DOI: 10.54254/2754-1169/2026.39315

Efficiency Comparison and Interaction Effects of Variance Reduction Methods in Monte Carlo Option Pricing: An Empirical Analysis Based on SHFE Copper Options

Taiheng Chen

Hong Kong Shue Yan University, Hong Kong, China
249A21@hksyu.edu.hk

Abstract. Monte Carlo (MC) simulation is a core numerical method for pricing financial derivatives, but its slow convergence can impose substantial computational costs when high precision is required. Variance reduction techniques, including antithetic variates (AV), control variates (CV), and quasi-Monte Carlo (QMC), have therefore been developed to improve computational efficiency. But these methods have been mainly analyzed individually or in simple combination in previous studies. It is not clear whether QMC and CV have complementary, overlapping or offsetting effects, especially in the Chinese commodity options. This paper systematically compares the standard MC, AV, CV, QMC and QMC+CV for pricing European and arithmetic-average Asian options by using the copper options traded on the Shanghai Futures Exchange. The variance reduction ratio, root mean squared error, and computation time are used to evaluate performance. The results indicate that QMC is much faster than MC and AV when pricing low dimensional European options while CV provides maximum variance reduction for Asian options. The efficiency gain due to QMC+CV is not just the sum of the two constituent methods: the interaction is positive at small numbers of time steps, but decreases and could be negative at high numbers of time steps. This research changes the debate from which is the best method to the interaction between the different variance reduction methods, and offers some useful guidance on the choice of variance reduction methods when pricing commodity options.

Keywords: Monte Carlo option pricing, Variance reduction, Quasi-Monte Carlo, Control variates, Interaction effects

1. Introduction

Monte Carlo (MC) simulation is widely used to price financial derivatives. However, standard MC converges slowly, so achieving high precision requires a large number of simulated paths and substantial computational resources. Compared with mainland Chinese equity options, Shanghai copper options offer more active trading and more continuous data. Copper futures also exhibit both trending behavior and mean reversion, making them a suitable empirical setting [1]. This study therefore asks how MC pricing efficiency can be improved and whether combining variance reduction methods yields additional gains.

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.39315

Existing research relies largely on simulated data and provides limited validation using observations from actual markets. In this study, it is argued that the uniformity advantage of QMC decreases with increasing dimensionality while the performance of control variates is mainly related to the correlation between the control variable and the target variable.The two effects may therefore reinforce each other in low-dimensional settings but offset each other in high-dimensional settings. Therefore, this study takes the copper option data from the Shanghai Futures Exchange and examines whether the theoretical advantages of QMC and CV hold in the Chinese commodity option market, how their interaction changes with the number of time steps, and whether that interaction differs between European and Asian options.

This study makes three contributions.First, standard MC, AV, CV, QMC and QMC+CV are compared in a single system to enhance comparability among methods. Secondly, it turns the research problem into "how to carry out the research and which ways to use together". Third, provide real Shanghai copper option data for practical reference in method selection. Test whether AV, CV, QMC and QMC+CV are significantly different in efficiency, and determine whether the combination of QMC and CV has an additive, amplifying or offsetting effect.

2. Literature review

2.1. Monte Carlo simulation in option pricing

Since Boyle introduced Monte Carlo simulation to option pricing (as cited in Liu et al.), it has become a core computational tool in finance because it can accommodate path-dependent and high-dimensional derivatives [2].

From the perspective of volatility inputs, Siswanah et al. used Nasdaq data to compare MC pricing accuracy under historical volatility and Newton-Raphson implied volatility [3]. They found that MC estimates based on implied volatility had substantially lower mean absolute percentage errors and tracked market prices more closely. This finding suggests that MC pricing accuracy depends not only on the number of simulated paths but also on the quality of model inputs, particularly the volatility estimate. Kim et al. compared the Black-Scholes model, standard MC, and several machine-learning algorithms for option pricing and found that machine-learning methods produced lower pricing errors in some settings [4]. Nevertheless, MC remains valuable because of its modeling flexibility, which reinforces the need to improve its computational efficiency.

2.2. Antithetic variates

Antithetic variates (AV) is one of the earliest variance reduction techniques in Monte Carlo simulations. The method is relatively simple to implement, computationally inexpensive, and reasonably stable when the payoff is monotonic with respect to the price of the underlying asset. Its variance reduction benefit can decline substantially, however, when the payoff is non-monotonic or exhibits weak symmetry, which limits its applicability.

2.3. Control variates

Control variates (CV) are among the most theoretically important variance reduction techniques. For Asian options, Kemna and Vorst proposed using the analytical price of a geometric-average Asian option as a control variate for the MC estimate of an arithmetic-average Asian option [5]. This approach has since become standard. The effectiveness of CV depends on the correlation between

DOI: 10.54254/2754-1169/2026.39315

the control variate and the target estimator: a stronger correlation generally produces a larger variance reduction, whereas a weak correlation can reduce or even reverse the benefit of the adjustment. More recent studies have developed data-driven control variates, including neural-network approaches that identify effective controls automatically [6]. These methods provide an alternative to the analytical geometric-average control variate used in this study.

2.4. Quasi-Monte Carlo method

Quasi-Monte Carlo (QMC) replaces the pseudorandom numbers used in standard MC with low-discrepancy sequences, such as Sobol and Halton sequences. By distributing sample points more uniformly over the integration domain, QMC can reduce numerical error. Joy, Boyle, and Tan were among the first to apply QMC systematically in numerical finance and demonstrated its efficiency advantage over standard MC in low-dimensional pricing problems [7]. Niederreiter provides a systematic treatment of QMC numerical integration and low-discrepancy point sets and sequences, offering a theoretical foundation for the sampling approach used in this study [8].

2.5. Combined variance reduction methods and research gap

Each variance reduction method has a distinct range of effectiveness, which has motivated research on combined approaches. Xu, Zhang, and Wang systematically compared standard MC, AV, CV, QMC, and QMC+CV for pricing arithmetic-average Asian options, using the variance reduction ratio (VRR) as the primary performance measure [9]. They found that QMC alone produced only modest efficiency gains, whereas combining QMC with CV substantially increased the VRR, suggesting a positive interaction between the two methods.

Several questions remain about whether this synergistic effect is generalizable. Commodity options may exhibit price dynamics such as mean reversion and stochastic volatility, which can alter the choice of control variate and thus the effectiveness of the QMC+CV combination. Moreover, prior studies have relied mainly on simulation experiments, with limited empirical testing based on data from actual commodity option markets.

The literature reveals three main gaps. First, few studies examine the mechanisms underlying interactions between combined variance reduction methods or determine whether the resulting efficiency gains vary systematically with option type (European or Asian) and simulation dimension (number of time steps). Second, research on commodity options remains limited, as most studies focus on equity or equity-index options. The empirical performance of these methods in commodity option markets has therefore not been tested systematically. Third, differences in evaluation metrics, experimental designs, and benchmark prices make it difficult to compare the practical limits of each method across studies. To address these gaps, this study uses Shanghai copper options to build a unified MC simulation framework and compare standard MC, AV, CV, QMC, and QMC+CV for European and Asian option pricing.

3. Methods

3.1. Model setup

This study assumes that Shanghai copper futures prices follow a geometric Brownian motion (GBM) under the risk-neutral measure:

$$dS_t = rS_t dt + \sigma S_t dW_t \tag{1}$$

DOI: 10.54254/2754-1169/2026.39315

where r denotes the risk-free interest rate, σ denotes volatility, and W_t is a standard Brownian motion. Dividing the time horizon T into n equal intervals of length Δt yields the following discrete-time path simulation equation:

$$S_{t+\Delta t} = S_t \exp[(r - \frac{\sigma^2}{2})\Delta t + \sigma\sqrt{\Delta t}Z], Z \sim N(0,1) \tag{2}$$

The model parameters are estimated from historical returns on Shanghai copper futures, and the risk-free rate is proxied by the three-month SHIBOR over the corresponding period. The Black-Scholes formula provides the analytical benchmark for European options [10]. Because arithmetic-average Asian options have no closed-form solution, a high-precision numerical estimate is used as the benchmark price.

3.2. Standard MC (basis)

Standard MC simulates N price paths and estimates the option price as the discounted average of the simulated payoffs:

$$\widehat{C}_{MC} = e^{-rT}\frac{1}{N}\sum_{i=1}^{N}\text{Payoff}_i \tag{3}$$

Its convergence rate is O(N^−1/2), and its standard error is σ̂_payoff/√N. Standard MC serves as the benchmark against which all variance reduction methods are evaluated in this study.

3.3. Antithetic variaties

For each simulated path, the antithetic variates method generates a standard normal draw Z^(i) together with its antithetic counterpart −Z^(i). The negative association between the two path payoffs reduces the variance of the estimator:

$$\widehat{C}_{AV} = e^{-rT}\frac{1}{N}\sum_{i=1}^{N}\frac{Payoff(Z^{(i)}) + Payoff(?Z^{(i)})}{2} \tag{4}$$

AV reduces variance without increasing the total number of simulated paths and therefore has a low computational cost. It is particularly suitable for options whose payoffs are monotonic in the underlying asset price, but its effectiveness is more limited for non-monotonic payoffs.

3.4. Control variaties

The control variates method introduces an auxiliary variable with a known expected value that is strongly correlated with the target estimator. The MC estimate is adjusted using this variable to reduce variance. For Asian option pricing, the analytical price of a geometric-average Asian option, C_Geo^exact, is used as the control variate:

$$\widehat{C}_{CV} = \widehat{C}_{Arith}^{MC} - \beta^*(\widehat{C}_{Geo}^{MC} - C_{Geo}^{exact}) \tag{5}$$

The optimal control coefficient is:

DOI: 10.54254/2754-1169/2026.39315

$$\beta^* = \frac{\widehat{Cov}(\widehat{C}^{MC}_{Arith}, \widehat{C}^{MC}_{Geo})}{\widehat{Var}(\widehat{C}^{MC}_{Geo})} \tag{6}$$

The optimal coefficient β* is estimated from the sample covariance. Because geometric-average and arithmetic-average Asian option prices are highly correlated, the CV method can, in principle, achieve substantial variance reduction in Asian option pricing.

3.5. Quasi-Monte Carlo

QMC replaces pseudorandom numbers with low-discrepancy sequences such as Sobol sequences, thereby reducing integration error through a more uniform distribution of sample points. Its theoretical convergence rate can approach O((log N)^d/N), which is faster than that of standard MC in low-dimensional problems. As dimensionality increases, however, the uniformity advantage of low-discrepancy sequences is weakened, reducing the marginal efficiency gain from QMC. This study uses Sobol sequences and applies the Box-Muller transform to obtain standard normal variates for path simulation.

The combined method applies the CV adjustment to paths generated by QMC:

$$\widehat{C}_{QMC+CV} = \widehat{C}^{QMC}_{Arith} - \beta^*(\widehat{C}^{QMC}_{Geo} - C^{exact}_{Geo}) \tag{7}$$

QMC reduces integration error by improving the uniformity of the sampling points, whereas CV reduces estimator variance by introducing a correlated auxiliary variable. Because the two techniques operate through different mechanisms, they may generate additional efficiency gains when combined. The central question of this study is whether the gain from QMC+CV equals the sum of the gains from the individual methods or instead reflects a positive or negative interaction.

To quantify the interaction between QMC and CV, this study defines the following excess-gain index,

$$\Delta_{interaction} = VRR_{QMC+CV} - VRR_{QMC} - VRR_{CV} + 1 \tag{8}$$

3.6. Evaluation metrics

The variance reduction ratio is defined as:

$$VRR = \frac{Var(\widehat{C}_{MC})}{Var(\widehat{C}_{method})} \tag{9}$$

A positive value of $\Delta_{interaction}$ indicates a synergistic effect, a negative value indicates an offsetting effect, and a value close to zero indicates that the two methods contribute approximately independently. The analysis tracks $\Delta_{interaction}$ across option types (European and Asian) and numbers of time steps, n.

DOI: 10.54254/2754-1169/2026.39315

Two additional metrics are used for cross-method comparisons:

$$VRR = \frac{Var(\widehat{C}_{MC})}{Var(\widehat{C}_{method})} \tag{10}$$

$$RMSE = \sqrt{\frac{1}{M}\sum_{j=1}^{M}(\widehat{C}_j - C^*)^2} \tag{11}$$

where C* is the benchmark price and M is the number of independent experiments. In addition, the time-adjusted efficiency ratio $TER = \frac{VRR}{\frac{t_{method}}{t_{MC}}}$ is used to evaluate the practical efficiency of each method while accounting for differences in computation time.

3.7. Data and experimental design

This article uses historical data of Shanghai Futures Exchange copper futures (2019-2024) and options (2024) to estimate historical volatility $\widehat{\sigma} = std(ln\frac{S_t}{S_{t-1}}) \times \sqrt{252}$ Log returns are used to estimate historical volatility, and the corresponding three-month SHIBOR is used as the risk-free rate. The experiments cover three moneyness categories: at-the-money (ATM), in-the-money (ITM), and out-of-the-money (OTM) options. The number of simulated paths is set to N $\in$ {1,000, 5,000, 10,000, 50,000}, and the number of time steps is set to n $\in$ {12, 52, 252} to examine how each method's efficiency changes with problem dimensionality. All experiments use the same random seed and computational resources to ensure comparability across methods.

4. Results

This section reports the result of numerical pricing of Shanghai copper options and the underlying futures contracts. The first analysis takes into account exchange-traded European-style copper options, and compares the convergence of MC, AV and QMC to the analytical Black-Scholes price. The second constructs arithmetic-average Asian options using the same market parameters and evaluates MC, AV, CV, QMC, and QMC+CV against a high-precision QMC+CV benchmark.

4.1. Volatility estimation

The volatility estimates include values of approximately 0.5639 and 0.5333. For some option contracts, volatility can be estimated directly from the historical returns of the corresponding futures contract; for others, the main futures contract is used as a proxy because the corresponding historical series is too short. Using multiple volatility sources therefore improves sample coverage.

4.2. European option pricing result

The European option results show that pricing errors relative to the Black-Scholes benchmark decline as the number of simulated paths increases. At N = 1,000, the RMSE values for MC, AV, and

DOI: 10.54254/2754-1169/2026.39315

QMC are 197.04, 175.66, and 21.04, respectively. At N = 50,000, the corresponding values fall to 35.62, 21.09, and 0.75. QMC therefore converges much faster than the other two methods and produces estimates close to the analytical benchmark even with relatively few simulated paths. By contrast, AV provides only a modest reduction in RMSE but delivers a stable variance reduction, with an average VRR of approximately 4.46.

The experiment also reveals substantial differences between Black-Scholes prices and observed market prices in some cases. For several contracts, for example, the Black-Scholes price is substantially higher than the market price. This pattern suggests that simulation error is not the main source of the empirical pricing discrepancy; the use of historical rather than implied volatility is a more plausible explanation. The European option analysis therefore evaluates how efficiently each simulation method converges to the Black-Scholes benchmark rather than whether the Black-Scholes model explains observed market prices.

Table 1. RMSE and variance reduction ratio (VRR) of MC, AV, and QMC relative to the black-scholes benchmark (European option)

N	MC RMSE to BS	AV RMSE to BS	QMC RMSE to BS	AV VRR	QMC VRR
1,000	197.04	175.66	21.04	4.46	1.02
5,000	107.57	103.02	6.21	4.44	0.99
10,000	63.18	66.63	3.37	4.48	1.00
50,000	35.62	21.09	0.75	4.46	1.00

4.3. Asian option pricing results

In the Asian option experiments, control variates (CV) perform best. A high-precision QMC+CV estimate is used as the benchmark. CV reduces variance substantially, producing a mean VRR above 2,000—far higher than the reductions achieved by AV or QMC alone. AV also provides a stable improvement, with a mean VRR ranging from 6.5 to 7.2. In contrast, QMC alone has a VRR close to 1, indicating limited effectiveness for high-dimensional, path-dependent options when used without CV.

In terms of RMSE, CV and QMC+CV perform well in most cases. At N = 1,000 and n = 12, the RMSE for MC is 165.29, compared with 147.96 for CV and 147.93 for QMC+CV. RMSE values increase slightly as the number of time steps rises to n = 252, indicating that pricing becomes more difficult when the average is computed more frequently along the simulated path.

Table 2. RMSE and variance reduction ratio (VRR) of MC, AV, CV, QMC, and QMC+CV for arithmetic-average asian option pricing

Method	RMSE at N=1000,n=12	RMSE at N=50000,n=252	Mean VRR
MC	165.29	187.14	1.00
AV	164.10	185.91	6.5–7.2
CV	147.96	182.58	2064–2231
QMC	150.51	188.10	1.0–1.1
QMC+CV	147.93	183.02	1972–2243

DOI: 10.54254/2754-1169/2026.39315

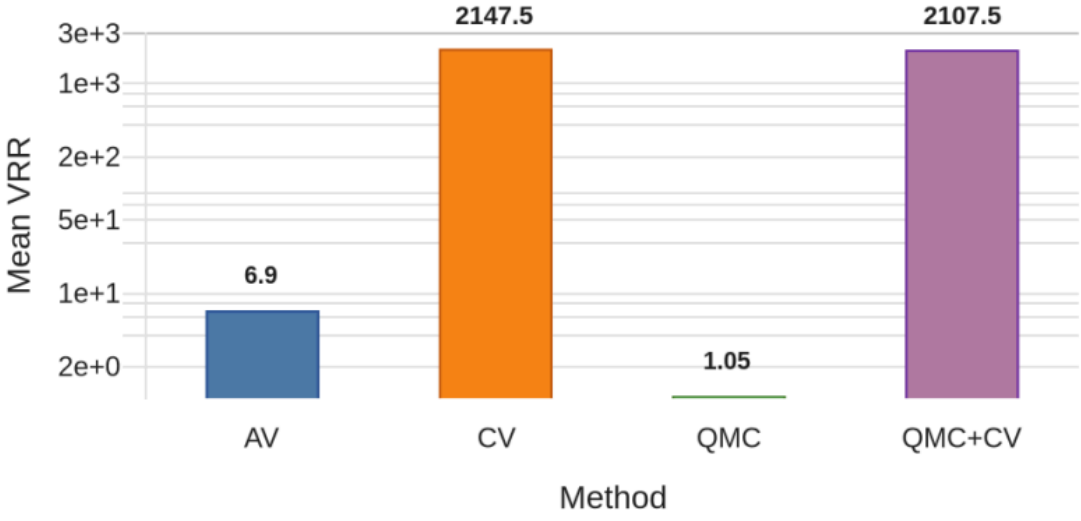

Figure 1. Asian options: variance reduction ratio by method

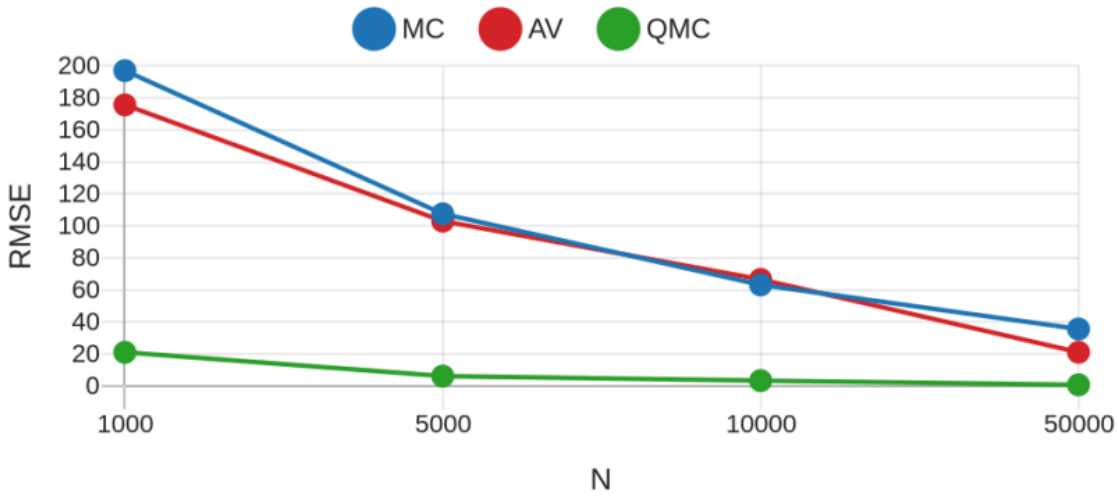

Figure 2. European option pricing: RMSE relative to Black-Scholes by simulation size

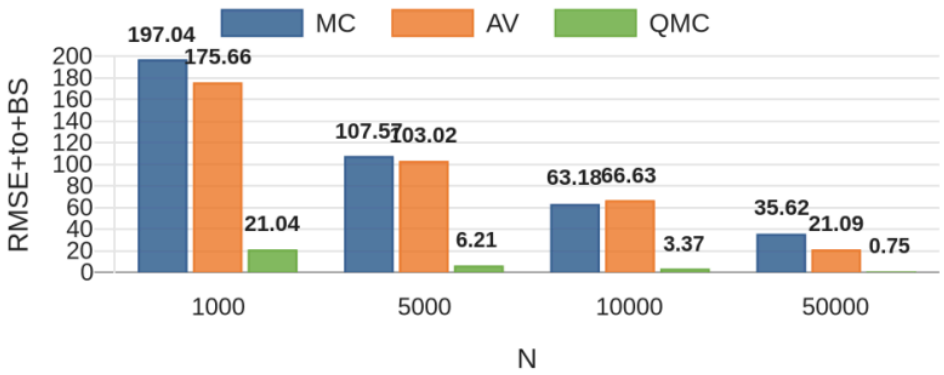

Figure 3. European options: RMSE relative to Black-Scholes by simulation size

DOI: 10.54254/2754-1169/2026.39315

4.4. Interaction effect between QMC and CV

The interaction between QMC and CV is not uniform across settings. With fewer time steps, such as n = 12 or n = 52, QMC+CV often exhibits a positive interaction, indicating an additional gain from combining the methods. At n = 252, however, the interaction weakens and may become negative. One possible explanation is that CV already accounts for most of the variance in Asian option prices, leaving little scope for QMC to provide an additional benefit.

Overall, the European option experiments show that QMC has a substantial advantage in low-dimensional pricing problems. The Asian option experiments identify CV as the most effective variance reduction method, while QMC+CV delivers comparatively stable performance. The results also show that no single numerical method is uniformly efficient as option structures become more complex. For path-dependent options, combining methods may be beneficial, but the gain depends on dimensionality and the effectiveness of the component techniques.

5. Conclusion

Using Shanghai copper options, this study systematically compares standard Monte Carlo, antithetic variates, control variates, quasi-Monte Carlo, and QMC+CV for pricing European and Asian options. It addresses a gap in the literature by examining not only the performance of each method in isolation but also whether the combined effect of QMC and CV is additive, amplifying, or offsetting.

The results of the European options show that, in a low-dimensional setting, QMC is much faster than standard MC and AV. This finding is consistent with prior research and suggests that the empirical characteristics of the Shanghai copper option market do not eliminate QMC's theoretical advantage in low-dimensional problems. For Asian options, CV is more effective than AV or QMC alone, supporting the classical control variate approach proposed by Kemna and Vorst [5].

The interaction between QMC and CV, however, varies with the number of time steps. This combination can have a positive effect when few time steps are specified. As the number of time steps increases, the interaction weakens and may become negative. A combination of methods thus does not necessarily provide better performance, but it is dependent on the dimension and payoff structure of the priced option.

This study has several limitations. First,The volatility is calculated based on historical data and not implied volatilities, which could account for some of the difference between Black-Scholes and market prices.Secondly, the analysis only covers copper options, and results may not be directly generalizable to other commodities. Third, the model has no stochastic volatility and only incorporates geometric Brownian motion, with no consideration for price jumps.

References

[1] Brignone, R., Gonzato, L., & Sgarra, C. (2024). Commodity Asian option pricing and simulation in a 4-factor model with jump clusters. *Annals of Operations Research, 336*(1), 275-306.

[2] Liu, H., Liu, Y., Tang, Q., & Zhu, J. (2026). Pricing catastrophe risk during transitions of physical and economic environments. *European Journal of Operational Research, 331*(2), 615-628.

[3] Siswanah, E., Saleh, M., Qudsiyah, E., & Hakim, M. M. (2024). Performance of two numerical methods in option price calculation. *Engineering Letters, 32*(3), 541-559.

[4] Kim, S., Kim, J., & Song, J. (2024). Option pricing and profitability: A comprehensive examination of machine learning, Black-Scholes, and Monte Carlo method. *Communications for Statistical Applications and Methods, 31*(5), 585-599.

DOI: 10.54254/2754-1169/2026.39315

[5] Kemna, A. G. Z., & Vorst, A. C. F. (1990). A pricing method for options based on average asset values. *Journal of Banking & Finance, 14*(1), 113-129.
[6] El Filali Ech-Chafiq, Z., Lelong, J., & Reghai, A. (2021). Automatic control variates for option pricing using neural networks. *Monte Carlo Methods and Applications, 27*(2), 91-104.
[7] Joy, C., Boyle, P. P., & Tan, K. S. (1996). Quasi-Monte Carlo methods in numerical finance. *Management Science, 42*(6), 926-938.
[8] Niederreiter, H. (1992). *Random number generation and quasi-Monte Carlo methods*. Society for Industrial and Applied Mathematics.
[9] Xu, L., Zhang, H., & Wang, F. L. (2023). Pricing of arithmetic average Asian option by combining variance reduction and quasi-Monte Carlo method. *Mathematics, 11*(3), Article 594. https://doi.org/10.3390/math11030594
[10] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy, 81*(3), 637-654.