

Comparison of Efficient Path Generation and Variance Reduction Strategies in Asian Option Pricing: Based on Sobol Sequences and Control Variates

Haochen Cai

*Department of Mathematics, City University of Hong Kong, Hong Kong, China
Haohecai4-c@my.cityu.edu.hk*

Abstract. Monte Carlo (MC) simulation is the predominant numerical method for pricing path-dependent derivatives, such as arithmetic average Asian options, which lack closed-form solutions. However, its slow convergence rate, $O(\frac{1}{\sqrt{N}})$, requires large computational budgets. This paper investigates and quantitatively compares two primary strategies for enhancing MC efficiency: Variance Reduction (VR) techniques, specifically Antithetic Variates (AV) and Control Variates (CV), and Efficient Path Generation using Quasi-Monte Carlo (QMC) methods with Sobol sequences. It benchmarks the performance of Standard MC, AV, CV (using the geometric Asian option as the control variable), and Sobol QMC using a comprehensive set of performance metrics: Variance Reduction Ratios (VRR), Root Mean Square Error (RMSE) convergence rates, and computational time efficiency. The results demonstrate that, while CV offers the highest VRR, the Sobol QMC method consistently achieves the fastest convergence rate towards the true value across varying levels of path complexity, providing superior efficiency, especially when high precision is required. This study offers practical guidance for financial engineers in selecting the optimal simulation acceleration technique based on their specific accuracy and computational budget constraints.

Keywords: Arithmetic Asian Options, Quasi-Monte Carlo, Variance Reduction, Computational Efficiency, Curse of Dimensionality

1. Introduction

The precise valuation of path-dependent derivatives remains a bedrock of modern quantitative finance. Since Boyle [1] first introduced the Monte Carlo (MC) framework to option pricing, Stochastic simulation has become an indispensable tool for pricing complex financial derivatives lacking closed-form analytical solutions. Among these, arithmetic Asian options whose payoffs depend on the arithmetic average of the underlying asset price pose a significant computational challenge, as the sum of log-normal variables does not follow a log-normal distribution, precluding a simple Black-Scholes closed-form solution. While standard MC methods provide unprecedented flexibility, their practical application is strictly bottlenecked by a slow statistical convergence rate of

$O(N^{-1/2})$, where N is the number of simulated paths [2]. In real-time risk management and trading environments, this slow rate requires massive computational budgets, driving decades of research into numerical acceleration.

To bypass this statistical bottleneck, research has historically diverged into two main optimization paradigms: Variance Reduction (VR) techniques and Quasi-Monte Carlo (QMC) methods. In the domain of VR, the classic Control Variates (CV) technique pioneered by Kemna and Vorst [3] represents a major milestone, utilizing the analytically tractable geometric average Asian option as an internal "control" to effectively eliminate simulation noise. Parallel to VR, QMC approaches enhance the convergence rate itself toward $O(N^{-1})$ by replacing random sampling with deterministic low-discrepancy sequences [4, 5]. However, the effectiveness of QMC is often highly sensitive to the monitoring frequency of the option path, frequently battling the "curse of dimensionality" as the integration space grows larger [6, 7].

In recent years, contemporary literature has heavily favored the complex hybridization of these standalone methods. For instance, Xu, Zhang and Wang [8] developed an integrated pricing framework for arithmetic Asian options by combining variance reduction with QMC methods under a time-varying GARCH volatility structure. Concurrently, Case [9] conducted a rigorous comparative analysis of standalone Sobol' and Faure sequences against standard MC, demonstrating that the efficiency gains of low-discrepancy sampling are strictly bounded by path dimensions. Furthermore, alternative advanced paradigms such as Multilevel Monte Carlo (MLMC) [10] and specialized dimension-reduction path constructions [11] continue to expand the choices available to quantitative practitioners.

Despite these extensive theoretical developments and the insights provided by modern studies like Xu et al. [8] and Case [9], there is a noticeable gap in literature regarding a simplified, head-to-head benchmarking of these strategies in their "pure" standalone forms within modern vectorized programming environments. This paper addresses this pragmatic gap by conducting a rigorous empirical comparison of Standard MC, Antithetic Variates (AV), CV, and Sobol-based QMC for pricing European arithmetic Asian call options. By introducing a standardized "Efficiency Index" that explicitly penalizes both statistical variance and computational execution runtime, this study provides data-driven guidance on the raw performance trade-offs between payoff-based (CV) and sampling-based (QMC) optimizations across diverse market regimes.

2. Methodology

2.1. Underlying asset and target derivative

This study assumes the underlying asset price S_t follows a Geometric Brownian Motion (GBM) under the risk-neutral measure:

$$\frac{dS_t}{S_t} = rdt + \sigma dW_t \quad (1)$$

where r is the risk-free rate, σ is the volatility, and dW_t is the increment of a standard Wiener process. The simulation discretizes the time interval $[0, T]$ into M steps of size $\Delta t = T/M$. The price at the next step is simulated as:

$$S_{t_{i+1}} = S_{t_i} \exp \left(\left(r - \frac{1}{2} \sigma^2 \right) \Delta t + \sigma \sqrt{\Delta t} Z_i \right) \quad (2)$$

where $Z_i \sim N(0, 1)$ are independent standard normal random variables. The target derivative is European arithmetic average Asian call option, with a payoff C_{Asian} at maturity T defined by

$$C_{\text{Asian}} = e^{-rT} E \left[\max \left(\frac{1}{M} \sum_{i=1}^M S_{t_i} - K, 0 \right) \right] \quad (3)$$

where K is the strike price.

2.2. Simulation methods implemented

2.2.1. Standard Monte Carlo (Standard MC)

This serves as the benchmark. The prices are estimated by simulating N independent paths using standard pseudo-random number generators (PRNGs) and averaging the discounted payoffs.

2.2.2. Antithetic Variates (AV)

This VR technique exploits the symmetry of the normal distribution. For every path generated using the vector of standard normal samples $Z = (Z_1, \dots, Z_M)$, a second path is generated using the antithetic vector $-Z = (-Z_1, \dots, -Z_M)$. The estimator is the average of the payoffs from both paths, which significantly reduces the variance provided the payoffs are negatively correlated.

$$C_{\text{AV}} = \frac{1}{2N} \sum_{k=1}^N [\text{Payoff}(S^{(k)}(Z)) + \text{Payoff}(S^{(k)}(-Z))] \quad (4)$$

2.2.3. Control Variates (CV)

The CV method uses an auxiliary variable Y that is strongly correlated with the target payoff X and whose expected value $E[Y]$ is known analytically. The optimal controlled estimator X_{CV} is:

$$X_{\text{CV}} = X - \beta^* (Y - E[Y]) \quad (5)$$

where the optimal coefficient β^* that minimizes the variance is given by:

$$\beta^* = \frac{\text{Cov}(X, Y)}{\text{Var}(Y)} \quad (6)$$

For the arithmetic Asian option (X), the control variable is defined by geometric average Asian option price Y , as it is highly correlated with X and possesses a known Black-Scholes-like closed-form solution for its expected value $E[Y]$.

2.2.4. Quasi-Monte Carlo (QMC) with sobol sequences

In QMC, the pseudo-random numbers $U \sim U(0,1)$ are replaced by points from a low-discrepancy sequence (LDS). This study adopts Sobol sequence, which is known for its effectiveness in financial dimensions. The uniform Sobol sequences are transformed into standard normal variables Z via the inverse cumulative distribution function (Inverse CDF):

$$Z_i = \Phi^{-1}(U_i) \quad (7)$$

The QMC estimator then computes the average payoff using the deterministic paths derived from these carefully sampled variables.

2.3. Performance metrics

To rigorously compare efficiency, the following metrics will be used:

2.3.1. Variance Reduction Ratio (VRR)

The ratio of the variance of the benchmark estimator to the variance of the accelerated estimator, demonstrating the raw power of variance reduction.

$$VRR = \frac{\text{Var}(C_{\text{Standard}})}{\text{Var}(C_{\text{Method}})} \quad (8)$$

2.3.2. Root Mean Square Error (RMSE) and convergence rate

The RMSE measures both bias and variance, representing the distance from a reference price C^* . This is used to plot the convergence curve against the number of paths N .

$$\text{RMSE}(N) = \sqrt{\frac{1}{J} \sum_{j=1}^J (C_j(N) - C^*)^2} \quad (9)$$

where C^* is a highly precise reference price (e.g., derived from a very large N standard MC run).

2.3.3. Computational time efficiency

The CPU time required for each method to achieve a predefined level of accuracy ϵ . This metric is crucial for comparing the overhead introduced by QMC and the calculations in CV (optimal β^* estimation) against the resultant gain in accuracy.

3. Numerical experiments

To ensure a rigorous and reproducible evaluation of the proposed simulation techniques, this study establishes a standardized experimental framework. This section outlines the baseline market parameters, defines the explicit operational protocols for the five core numerical experiments, and details the computational environment.

3.1. Market parameters and asset specifications

The parameters are standardized across all baseline experiments as detailed in Table 1.

Table 1. Baseline parametric framework

Parameter	Symbol	Baseline Value
Initial Stock Price	S_0	100

Table 1. (continued)

Strike Price	K	100 / 80
Risk-Free Rate	r	5%
Volatility	σ	20%
Time to Maturity	T	1.0 Year
Monitoring Frequency	M	50

3.2. Core numerical experiments

Five sequential experiments are designed to comprehensively evaluate the performance of Standard Monte Carlo (SMC), Antithetic Variates (AV), Control Variates (CV), and Sobol-based Quasi-Monte Carlo (QMC).

3.2.1. Experiment 1: reference price determination and benchmarking

Purpose: To establish an asymptotic, high-precision reference price (C^*) to compute absolute errors and convergence matrices in subsequent stages, while validating the financial logic of the simulation engine.

Concrete Operations:

(1) Reference Pricing: Since arithmetic Asian options lack an analytical closed-form solution, a large-scale simulation using Standard MC with $N = 10^7$ paths and $M = 50$ is executed to obtain the benchmark C^* . The resulting statistical standard error (SE) is guaranteed to be below 3×10^{-3} , formalized as:

$$SE = \frac{\widehat{s_N}}{\sqrt{N}} \quad (10)$$

where $\widehat{s_N}$ is the sample standard deviation of the discounted terminal payoffs.

(2) Convergence Verification: The pricing engine's behavior is cross-checked across expanding sample tiers ($N \in \{10^4, 10^5, 10^6, 10^7\}$) to ensure the steady elimination of sampling noise.

3.2.2. Experiment 2: empirical RMSE convergence analysis

Purpose: To empirically test and verify the theoretical asymptotic error-decay rates—specifically checking the classical $O(N^{-1/2})$ rate against the enhanced Quasi-Monte Carlo $O(N^{-1})$ limit.

Concrete Operations:

(1) Path Scaling: The number of simulated paths scales across a uniform sequence of pure powers-of-two: $N \in \{2^6, 2^8, 2^{10}, 2^{12}, 2^{14}, 2^{16}, 2^{18}, 2^{20}\}$ (ranging from 64 to 1048576 paths). This

discrete structure ensures optimal, unbiased radix configurations for the Sobol low-discrepancy sequence.

(2) Statistical Replication: For each fixed path allocation N , the simulation is independently repeated across $R = 50$ replications to construct robust statistical bounds.

(3) Convergence Estimation: The results are mapped onto a log-log coordinate system. The empirical convergence exponent α is derived via ordinary least squares (OLS) regression: $\ln(\text{RMSE}_N) = -\alpha \ln(N) + \beta$, directly revealing the error-decay velocity of each architecture.

3.2.3. Experiment 3: Variance Reduction Ratio (VRR) and moneyness sensitivity framework

Purpose: To isolate and quantify the noise-canceling capabilities of the variance reduction strategies under different option regimes, specifically evaluating how non-linear payoff boundaries affect algorithmic performance.

Concrete Operations:

(1) Moneyness Isolation: The experiment bypasses extreme tail-risk scenarios and isolates the two most heavily traded market profiles: an At-The-Money (ATM) regime ($K = 100$) and an In-The-Money (ITM) regime ($K = 80$).

(2) Variance Auditing: The total path budget is fixed at $N = 65,536$ with an observation density of $M = 50$. The sample variance (Var) of the discounted option payoffs is computed for SMC, AV, CV, and QMC.

This design is structured to test the Non-Linear Truncation Hypothesis, which dictates that the non-linear kink of the $\max(\cdot, 0)$ operator at $K = 100$ neutralizes the negative covariance mechanism of AV, whereas an ITM position ($K = 80$) restores linearity and triggers a performance awakening.

3.2.4. Experiment 4: high-dimensional stress test (curse of dimensionality)

Purpose: To evaluate the structural robustness of the deterministic Scrambled Sobol sequence and the Control Variate framework as the path dimensionality grows and dilutes space-filling properties.

Concrete Operations:

(1) Dimensional Scaling: The option's monitoring frequency is adjusted across an expanding integration spectrum: $M \in \{10, 20, 50, 100, 250, 350, 500\}$, scaling the path architecture from monthly to daily monitoring intervals over a 1-year horizon.

(2) Control Boundaries: The sample size is held constant at an intermediate institutional level ($N = 16,384$) to isolate dimensional degradation from sample scale advantages.

(3) Error Mapping: For each dimension M , a unique high-density reference target C_M^* is established. The absolute pricing error ($\epsilon_M = |\widehat{C_M} - C_M^*|$) is mapped across all dimensions to monitor the performance decay of the low-discrepancy sequence.

3.2.5. Experiment 5: market microstructure replication and logarithmic joint efficiency assessment

This experiment evaluates the joint trade-off between statistical precision and computational execution time across diverse asset volatilities, path dimensions, and contract horizons. By incorporating explicit computational costs, this design provides a robust quantitative foundation for algorithmic selection in practical financial engineering.

To eliminate sample-scale biases, the total path allocation is fixed at a standard production capacity of $N = 65536$. The pricing engine is independently executed across three regimes:

(1) Group A (Low-Noise Stable Market: Blue-Chips): Calibrated with $\sigma = 12\%$, $r = 4.5\%$, $T = 0.5$ years, and $M = 24$ (representing semi-monthly monitoring).

(2) Group B (High-Noise High-Dimension Market: Tech Growth and High-Frequency Derivatives): Calibrated with $\sigma = 65\%$, $r = 5.0\%$, $T = 1.0$ year, and $M = 250$ (representing daily monitoring).

(3) Group C (Mid-Dimension Extended-Horizon Market: Commodities and Energy Options): Calibrated with $\sigma = 35\%$, $r = 2.0\%$, $T = 3.0$ years, and $M = 36$ (representing monthly monitoring).

To quantify the precision gain achieved per unit of computational effort, we utilize a work-normalized variance metric. The raw computational efficiency ($\mathcal{E}$) is defined as the inverse product of sample variance and processing time:

$$\mathcal{E} = \frac{1}{\text{Variance} \times \text{Computation Time}} \quad (11)$$

where *Variance* represents the sample variance of the discounted terminal payoffs, and *Computation Time* is the mean CPU wall-clock execution time computed across 50 independent trials in an isolated background thread to eliminate operating system noise.

To clearly capture and display the vast disparities in performance magnitudes across regimes, the final reportable metric is defined as the Logarithmic Joint Efficiency Index ($\mathcal{E}_{\log}$) by directly taking the base-10 logarithm of raw efficiency:

$$\mathcal{E}_{\log} = \log_{10}(\mathcal{E}) \quad (12)$$

Under this framework, optimal engineering choices can be determined by directly performing cross-sectional comparisons of the absolute $\mathcal{E}_{\log}$ scores within each respective market microstructure.

3.3. Implementation details

The simulation framework is implemented using a vectorized architecture in Python 3.10. Vectorized path generation and matrix operations are managed via NumPy, while the low-discrepancy grid optimization relies on the `scipy.stats.qmc.Sobol` module, configured with Owen's scrambling algorithm (`scramble=True`) and boundary truncation ($\pm 10^{-7}$) for numerical stability.

To guarantee the reliability and reproducibility of the time-error and joint efficiency metrics, all experiments are executed under an isolated background thread condition on a 2.3 GHz Quad-Core Intel Core i7 processor with 32GB RAM.

4. Results and conclusions

Based on a rigorous empirical valuation of arithmetic average Asian options, the experimental findings are structured chronologically by experiment to evaluate numerical accuracy and computational efficiency:

4.1. Experiment 1: baseline pricing and benchmarking framework

Table 2. Numerical convergence and baseline validation for arithmetic Asian option valuation

Number of Paths (N)	Estimated Option Price	Standard Error (SE)
10000	5.9855	0.081462
100000	5.8533	0.025618
1000000	5.8612	0.008109
10000000 (Benchmark C*)	5.8524	0.002562

Table 2 presents the pricing results and standard errors of Standard MC with different sample sizes. The reference price $C^* = 5.8524$ when $N = 10^7$, with standard error below 3×10^{-3} . The error decreases in line with the theoretical $O(N^{-1/2})$ rule.

4.2. Experiment 2: empirical RMSE convergence trajectory

In Figure 1, Standard MC and AV follow the theoretical $O(N^{-1/2})$ convergence slope. CV maintains the lowest RMSE across all sample sizes. Sobol QMC shows faster error decay at small N, achieving the smallest RMSE when the number of paths is large.




Figure 1. Empirical error convergence trajectory under Standard ATM regime in Experiment 2

4.3. Experiment 3: structural sensitivity and VRR tracking

The bar chart in Figure 2 maps the Variance Reduction Ratio (VRR) across distinct moneyness blocks. Under ATM condition ($K = 100$), the VRR of CV reaches 1273.5. Under ITM condition ($K = 80$), VRR rises to 1563.5. By contrast, AV and QMC show no obvious variance reduction effect ($VRR \approx 1.0$). The non-linear payoff truncation weakens the performance of AV, while a higher linearity of ITM options further enhances CV's effect.




Figure 2. Structural sensitivity and AV performance awakening test in Experiment 3

4.4. Experiment 4: path dimensionality stress

Figure 3 shows that on a log scale axis tracking absolute price error, the green CV line demonstrates high structural immunity, remaining bounded within a low-error interval ($10^{-3} \leq \text{Error} \leq 10^{-2}$) across all dimensional points. QMC achieves optimal accuracy when $M = 100$, while its error rises sharply when $M > 100$ due to the curse of dimensionality. Standard MC and AV show violent error fluctuations in high-dimensional scenarios, and the anti-correlation effect of AV completely fails when $M = 500$.




Figure 3. Absolute error under high-dimensional paths in Experiment 4

4.5. Experiment 5: market logarithmic joint efficiency index

In all three market regimes shown by Figure 4, the efficiency ranking remains consistent: $CV > QMC > AV > \text{Standard MC}$. In low-volatility markets, CV has an obvious advantage. In high-volatility and high-dimensional markets, the overall efficiency of all methods declines, but CV still ranks first. QMC performs well in medium-dimensional long-horizon markets.




Figure 4. Efficiency performance $\log_{10} E$ tracing under empirical market conditions in Experiment 5

5. Discussion

5.1. Harmonization of empirical findings with theoretical expectations

All experimental outcomes align with core quantitative finance theories. Experiment 2 verifies the Central Limit Theorem: Standard MC (SMC) and Antithetic Variates (AV) converge at the theoretical $O(N^{-1/2})$ rate, while Sobol Quasi-Monte Carlo (QMC) exhibits faster initial convergence as predicted by the Koksma-Hlawka inequality. Experiment 3 confirms that Control Variates (CV) performance hinges on payoff linearity. Shifting options from ATM ($K = 100$) to ITM ($K = 80$) reduces truncation distortion in the payoff function, lifts the correlation between arithmetic and geometric Asian payoffs, and raises CV's Variance Reduction Ratio (VRR) from 1273.5 to 1563.5. Experiment 4 exposes QMC's high-dimensional limitations. Sobol QMC accuracy peaks at $M = 100$ and deteriorates sharply as monitoring steps rise to 500. This U-shaped error trend stems from the curse of dimensionality: Sobol's low-discrepancy space-filling properties weaken in high-dimensional hypercubes, creating undersampled gaps that damage integration precision. Experiment 5's efficiency metrics reflect real-world computational trade-offs. Though CV and QMC cut estimation variance, they add extra matrix and logic computation overhead. In high-frequency high-volatility markets, this runtime cost offsets part of their accuracy gains.

5.2. Analysis of anomalous volatility and discrepancies

Despite these theoretical validations, experiment 4 reveals notable numerical volatility and non-monotonic error spikes that require careful analysis:

SMC and AV Oscillations: The chaotic zigzag patterns of the SMC and AV curves across the time-step dimension highlight the vulnerabilities of pure pseudo-random sampling within a fixed path budget. At $M = 100$, SMC may benefit from a fortuitous path distribution ("lucky sampling") that artificially drops its absolute error near 10^{-2} . Conversely, at $M = 500$, AV's absolute error spikes sharply to nearly 10^{-1} . This spike occurs because standard sequential path generation maps independent random variables chronologically, causing the local variance from early increments to propagate unchecked. Without dimension reduction, high-dimensional spaces dilute the anti-correlation required by AV, causing the paired paths to decouple and act as independent noisy simulations.

6. Conclusion

This study comprehensively compares four Monte Carlo simulation approaches for pricing arithmetic Asian options: standard Monte Carlo, antithetic variates, control variates and Sobol sequence-based quasi-Monte Carlo. The results indicate that control variates deliver the strongest variance reduction capability, while Sobol quasi-Monte Carlo achieves the fastest convergence rate and superior overall efficiency when high precision is required. Nevertheless, quasi-Monte Carlo suffers from performance degradation under high-dimensional scenarios due to the curse of dimensionality, and antithetic variates show unstable accuracy with rising path dimensions. Across different market environments, control variates maintain the highest comprehensive computational efficiency, followed by Sobol quasi-Monte Carlo, while standard Monte Carlo and antithetic variates rank the lowest. Empirical results are consistent with classical quantitative finance theories. In high-dimensional tests, the space-filling advantage of Sobol sequences gradually diminishes, leading to growing pricing errors. Meanwhile, pseudo-random sampling used by standard Monte Carlo and antithetic variates produces volatile errors in high-dimensional settings, and the anti-correlation effect of antithetic variates fails when dimensions expand. In terms of combined efficiency considering both variance and runtime, all tested methods retain a stable performance ranking in low-volatility, high-volatility and long-horizon market regimes. For financial practitioners, control variates are the preferred choice for general scenarios with limited computational budget, while Sobol quasi-Monte Carlo is more suitable for tasks demanding high accuracy and low dimensionality.

This study uses the geometric Brownian Motion model with constant volatility, ignoring real-market characteristics including stochastic volatility, fat-tailed returns and price jumps. Additionally, the fixed simulation paths cannot dynamically allocate computational resources, which may cause slow convergence in volatile markets. The high efficiency of control variates relies on the analytical solution of geometric Asian options under GBM, leading to poor portability for other exotic options or advanced asset models. For future research, deep learning can be adopted to construct universal control variates for complex financial models. Adaptive sample allocation algorithms can also be developed to dynamically adjust simulation scale and further improve computational efficiency in practical pricing tasks.

References

- [1] Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics*, 4(3), 323-338.
- [2] Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer Science & Business Media.
- [3] Kemna, A. G., & Vorst, A. C. (1990). A pricing method for options based on average asset values. *Journal of Banking & Finance*, 14(1), 113-129.
- [4] Niederreiter, H. (1992). *Random Number Generation and Quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics.
- [5] Joy, C., Boyle, P. P., & Tan, K. S. (1996). Quasi-Monte Carlo methods in numerical finance. *Management Science*, 42(6), 926-938.
- [6] Kuo, F. Y., & Sloan, I. H. (2005). Lifting the curse of dimensionality associated with quasi-Monte Carlo integration. *Notices of the AMS*, 52(11), 1320-1328.
- [7] Dick, J., Kuo, F. Y., & Sloan, I. H. (2013). High-dimensional integration: The quasi-Monte Carlo way. *Acta Numerica*, 22, 133-288.
- [8] 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), 594.
- [9] Case, G. (2025). Comparative Study of Monte Carlo and Quasi-Monte Carlo Techniques for Enhanced Derivative Pricing. *arXiv preprint arXiv: 2502.17731*.
- [10] Giles, M. B. (2008). Multilevel Monte Carlo path simulation. *Operations Research*, 56(3), 607-617.

- [11] Imai, J., & Tan, K. S. (2006). A general dimension-reduction technique for derivative pricing. *Journal of Computational Finance*, 9(2), 29.