Stochastic Simulations

Autumn Semester 2025

Prof. Fabio Nobile *Assistant: Eliott Van Dieren*

Project 4

submission deadline: 11 January 2026

QMC integration of non-smooth functions: application to pricing exotic options

1 Introduction and background

In high-dimensional integration, the goal is to approximate a multiple integral over an m-dimensional ($m \gg 1$) domain, e.g., the unit cube $[0,1]^m$:

$$I_m(\phi) = \int_{[0,1]^m} f(x)\mathrm{d}x = \int_0^1 \cdots \int_0^1 f(x_1, \ldots, x_m)\mathrm{d}x_1 \ldots \mathrm{d}x_m, \tag{1}$$

or, more generally, over $\mathbb{R}^m$ with respect to some distribution $\rho := \rho_1(x_1)\rho_2(x_2)\ldots\rho_m(x_m)$, i.e.,

$$I_m(\phi) = \int_{\mathbb{R}^m} f(x)\rho(x)\mathrm{d}x = \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} f(x_1, \ldots, x_m)\rho_1(x_1)\ldots\rho_m(x_m)\mathrm{d}x_1 \ldots \mathrm{d}x_m. \tag{2}$$

By now, quasi-Monte Carlo (QMC) has become a popular way to approximate the high-dimensional integration problem (1). Recall (c.f. lecture notes) that the approximation error of the quasi-Monte Carlo method can be controlled by the product of the star-discrepancy of the point set $P_N = \{x_1, \ldots, x_N\}$ and a term measuring the *smoothness* of the integrand. More precisely, the Koksma-Hlawka inequality states that the error can be bounded as

$$\left| \int_{[0,1]^m} f(x)\mathrm{d}x - \frac{1}{N}\sum_{i=1}^N f(x_i) \right| \le \underbrace{\left(\sum_{\mathbf{k}\subset\{1,\ldots,s\}} \int_{[0,1]^{|\mathbf{k}|}} \left| \frac{\partial^{|\mathbf{k}|}}{\partial x_{\mathbf{k}}} f(x_{\mathbf{k}}, \mathbf{1}) \right| \mathrm{d}x_{\mathbf{k}} \right)}_{:=V_{\mathrm{HK}}} \mathcal{D}^*(P_N), \tag{3}$$

where $x_{\mathbf{k}} = (x_{k_1}, x_{k_2}, ..., x_{k_p})$, if $|\mathbf{k}| = p \le s$, $\mathcal{D}^*(P_N)$ is the so-called star-discrepancy of the point set P_N and V_{HK} is the total variation norm in the sense of Hardy and Krause. Notice that for (3) to be a meaningful error estimator, f needs to have integrable mixed first order derivatives.

In this mini-project we implement a method proposed in [1] for numerically computing integrals of the form (2) using QMC in the presence of "kinks" (i.e. discontinuities in the gradients) or "jumps" (i.e. discontinuities in the function) on the integrand $f(x)$, where those kinks and jumps are not aligned to the coordinate axis, hence the integrand $f(x)$ fails to have bounded V_{HK}. Such functions commonly arise in financial mathematics, more concretely in option pricing, where the payoff function is often of the form $f(x) = \max(0, \Theta(x) - k)$, with

$\Theta : \mathbb{R}^m \to \mathbb{R}$ a smooth function for, say a call option, or $f(x) = x\mathbb{I}_{\{\phi(x)>a\}}$, for, e.g., binary digital options. In the first case, f may have a kink across the surface $\Theta(x) = k$, whereas in the second case f may have a jump across the surface $\Phi(x) = a$. More generally, let us consider integrands $f : \mathbb{R}^m \mapsto \mathbb{R}$ of the form

$$f(x) = \theta(x)\mathbb{I}_{\{\phi(x)\}}, \tag{4}$$

for smooth functions θ and ϕ, and assume that, $\exists j \in \{1, 2, \ldots, m\}$ such that $\frac{\partial \phi}{\partial x_j}(x) > 0 \ \forall x \in \mathbb{R}^m$ and that $\phi(x) \to \pm\infty$ as $x_j \to \pm\infty$.

The method proposed in [1] is relatively simple: notice that we can write (2) as

$$I_m(\phi) = \int_{\mathbb{R}^m} f(x)\rho(x)\mathrm{d}x = \int_{\mathbb{R}^{m-1}} \underbrace{\left(\int_{\mathbb{R}} f(x_j, x_{-j})\rho_j(x_j)\mathrm{d}x_j\right)}_{:= p(x_{-j})} \rho_{-j}(x_{-j})\mathrm{d}x_{-j}. \tag{5}$$

where we have written $x = (x_j, x_{-j})$, with $x_{-j} = (x_1, x_2, \ldots, x_{j-1}, x_{j+1}, \ldots, x_m)$, and similarly for ρ_j and ρ_{-j}. Instead of evaluating $\int_{\mathbb{R}^m} f(x)\rho(x)\mathrm{d}x$ directly by a QMC formula, we could alternatively, first evaluate the inner integral for a set of given points x_{-j}, and then evaluate the outer integral in $m-1$ dimensions via, say MC or QMC. What is interesting about this approach is that the resulting inner function $p(x_{-j})$ is smooth, i.e., it doesn't have a kink or a jump anymore. To see this, notice that for any fixed $x_{-j} \in \mathbb{R}^{m-1}$ by the assumptions on the derivatives of ϕ, the function $x_j \mapsto f(x_j, x_{-j})$ has a jump at the (unique) point where $\phi(x_j, x_{-j}) = 0$. It follows from the implicit function theorem that for each x_{-j}, there exists a unique value $x_j^* := \psi(x_{-j})$ for which $\phi(x_j, x_{-j}) < 0$ if $x_j < x_j^*$ and $\phi(x_j, x_{-j}) > 0$ if $x_j > x_j^*$. Thus, given the chosen form of $f(x)$, we can write

$$p(x_{-j}) := \int_{\mathbb{R}} f(x_j, x_{-j})\rho_j(x_j)\mathrm{d}x_j = \int_{\psi(x_{-j})}^{+\infty} \theta(x_j, x_{-j})\rho_j(x_j)\mathrm{d}x_j, \tag{6}$$

with both θ and ψ smooth, and

$$I_m(\phi) = \int_{\mathbb{R}^{m-1}} p(x_{-j})\rho_{-j}(x_{-j})\mathrm{d}x_{-j}. \tag{7}$$

Notice that once $\psi(x_{-j})$ has been found, the one dimensional integral (6) can be evaluated using, e.g., a Gauss quadrature and (7) can be recast into an $(m-1)$-dimensional integral over the hypercube $[0, 1]^{m-1}$ with, hopefully, a smooth integrand.

1.1 Application to option pricing

Consider the problem of pricing an option with maturity $T > 0$ based on the stock price S, which is given as the solution to the stochastic differential equation

$$dS = rS\,dt + \sigma S\,dw_t\ , \quad S(0) = S_0\ .$$

Here S_0 represents the initial value of the underlying asset, K is the so-called strike price, r is the interest rate, σ is the volatility and w_t denotes a standard one-dimensional Wiener process. One can show that $S_t = S_0 \exp\big((r - \sigma^2/2)t + \sigma w_t\big)$. It follows that S_t has a log-normal distribution for any $t > 0$. For $m \in \mathbb{N}$, let $t_i = i\Delta t$ with $\Delta t = T/m$ denote the discrete observation times of the stock price S (e.g. daily at market closure). The problem of simulating asset prices can be reduced to the problem of simulating discretized Brownian motion paths taking values $w = (w_{t_1}, ..., w_{t_m})$. Notice that w is normally distributed with covariance matrix having entries $C_{i,j} = \min\{t_i, t_j\}$, $i, j = 1, \ldots, m$. In particular, we will be interested in pricing Asian options with the following payoffs:

$$\Psi_1(w_{t_1}, \ldots, w_{t_m}) := \left(\frac{1}{m} \sum_{i=1}^{m} S_{t_i}(w_{t_i}) - K \right)_+ = \left(\frac{1}{m} \sum_{i=1}^{m} S_{t_i}(w_{t_i}) - K \right) \mathbb{I}_{\left\{ \frac{1}{m} \sum_{i=1}^{m} S_{t_i}(w_{t_i}) - K \right\}},$$
$$\Psi_2(w_{t_1}, \ldots, w_{t_m}) := \mathbb{I}_{\left\{ \frac{1}{m} \sum_{i=1}^{m} S_{t_i}(w_{t_i}) - K \right\}}.$$

Here, Ψ_1 corresponds to the payoff of an Asian call option and Ψ_2 corresponds to the payoff of a binary digital Asian option. Notice that the value $V_i,\ i = 1, 2$, of such options is given by

$$V_i := e^{-rT} \mathbb{E}[\Psi_i(w_{t_1}, \ldots, w_{t_m})] = \frac{e^{-rT}}{(2\pi)^{m/2} \sqrt{\det(C)}} \int_{\mathbb{R}^m} \Psi_i(w_{t_1}, \ldots, w_{t_m}) e^{-\frac{1}{2} w^T C^{-1} w} \mathrm{d}w, \quad (8)$$

Thus, the problem of pricing an option with m discrete time steps, can be treated as a problem of integration in m dimensions, however, with a discontinuous or only Lipschitz continuous function.

2 Goals of the project

Consider the task of estimating V_1, V_2, with parameters $K = 100, S_0 = 100, r = 0.1, \sigma = 0.1, T = 1$, and different discretization parameters, namely $m = 32, 64, 128, 256, 512$. To recast the integral (8) into a hypercube, consider rewriting it first in terms of independent standard normal random variables and then transforming them into uniform random variables in $[0, 1]$. For the former transformation, consider taking a Cholesky factorization of C, or a reparametrization into Gaussian increments $\xi_i = \frac{w_{t_i} - w_{t_{i-1}}}{\sqrt{t_i - t_{i-1}}}$, or using a Lévy-Ciesielski construction (see [2, §3.1]).

1. First estimate the value of the integral, as well as the error of the estimation, using a crude Monte Carlo and a randomized QMC, without the pre-integration trick, i.e., by generating (Q)MC sample $(x_1, \ldots, x_N)$ over the m-dimensional unit cube $[0, 1]^m$. Try increasing values of the sample sizes $N = 2^7, \ldots, 2^{13}$ and plot the estimated error versus N. Comment on the observed convergence rate.

2. Implement now the pre-integration trick. First, generate randomized (Q)MC points over $[0, 1]^{m-1}$, and decide on which direction x_j to perform the integration. Discuss your choice (or, alternatively, compare numerically different choices) depending also on the adopted re-parametrization of the Brownian path. Use the pre-integration approach to estimate V_i using both QMC and MC. Estimate your (Q)Monte Carlo error using sample sizes $N = 2^7, \ldots, 2^{13}$. Plot the estimated error versus N. Compare your results to those obtained in point 1 in terms of error and computational cost.

3. Consider now $K = 120$ and the goal of computing V_1 and V_2 with given relative accuracy 10^{-2} in the mean squared sense. What sample size should be used for a crude MC estimator? Propose a variance reduction technique to improve the performance of the crude MC estimator. **Note:** You can change the value of K to suit your computational budget.

4. Can the variance reduction technique proposed in point 3 be used also for the QMC estimators (with and without the pre-integration trick)? If yes, how? Repeat point 3 for the QMC estimators. What do you observe?

References

[1] Andreas Griewank, Frances Y Kuo, Hernan Leövey, and Ian H Sloan. High dimensional integration of kinks and jumps—smoothing by preintegration. *Journal of Computational and Applied Mathematics*, 344:259–274, 2018.

[2] René L. Schilling and Lothar Partzsch. *Brownian Motion.* De Gruyter, Berlin, Boston, 2012.