Pricing and Risk Analysis in Hyperbolic Local Volatility Model with Quasi-Monte Carlo

Julien Hok
Investec Bank, London, UK
Sergei Kucherenko
BRODA Ltd, UK, e-mail: s.kucherenko@broda.co.uk

Abstract

Local volatility models usually capture the surface of implied volatilities more accurately than other approaches, such as stochastic volatility models. We present the results of applying Monte Carlo (MC) and quasi-Monte Carlo (QMC) methods for derivative pricing and risk analysis based on the hyperbolic local volatility model. In high-dimensional integration, QMC shows a superior performance over MC if the effective dimension of an integrand is not too large. In application to derivative pricing and computation of Greeks, effective dimensions depend on path discretization algorithms. The results presented for the Asian option show the superior performance of QMC methods, especially for the Brownian bridge discretization scheme.

Keywords

Monte Carlo methods in finance, quasi-Monte Carlo, Sobol, sequences, Brownian bridge, skew/smile models, hyperbolic local volatility model

1. Introduction

Monte Carlo (MC) methods are widely used in valuation of complex financial instruments. Although the convergence rate of MC methods is $O\left(\frac{1}{\sqrt{N}}\right)$, where N is the number of sampled points, does not depend on the number of variables n, it is rather slow. Switching from random numbers to quasi-random numbers such as low-discrepancy sequences (LDS) can significantly improve the convergence under some conditions. Methods based on LDS are known as quasi-Monte Carlo (QMC). Asymptotically, they can provide a rate of convergence $O(\frac{1}{N})$. The quality of Sobol' sequences depends heavily on the so-called direction numbers and in practice, very few Sobol' sequence generators show good efficiency in valuation of complex financial instruments (see e.g. Sobol' *et al.*, 2011; BRODA, 2020).

The solution to financial problems such as pricing and hedging can be formulated as a mathematical expectation of some functionals, which in practice is reduced to evaluation of the Wiener path integrals. The nominal dimensions of such integrals are the product of the number of time steps at which the asset prices are observed and the number of risk factors (the underlying assets). This can reach many thousands of dimensions. QMC methods can lose their superior performance in high-dimensional settings unless the problem's effective dimension is low. The concept of effective dimension was introduced by Caflisch *et al.* (1997). It was shown in many papers that QMC is superior to MC if the effective dimension of an integrand is not too large. Effective dimensions and the QMC convergence of path-dependent integrals depend on path discretization algorithms. There are two such algorithms that are widely used in finance: (1) the incremental (also known as standard) discretization algorithm and (2) the Brownian bridge discretization algorithm. Both algorithms have the same variance, hence their MC convergence rates are the same. However, the corresponding QMC algorithms have different efficiencies, with the Brownian bridge algorithm having a much higher convergence rate for the majority of payoffs. This combination of Sobol, points with the Brownian bridge construction was proposed by Moskowitz and Caflisch (1996) and has been found to be highly effective in finance applications (see e.g. Acworth *et al.*, 1998; Åkesson & Lehoczky, 2000; Bianchetti *et al.*, 2015; Caflisch *et al.*, 1997; Kucherenko & Shah, 2007).

Pricing and hedging of financial instruments has been primarily based on *Gaussian* models, where the underlying asset dynamics (interest rates, equity prices or exchange rates) are assumed to follow a *Hull–White* or *Black–Scholes* model. However, empirical asset return distributions tend to exhibit fat tails (kurtosis) and skewness (asymmetric distribution). The *skew* or *smile* in *implied volatility* surfaces (defined in Section 3) observed across various asset classes are market reality (see e.g. Gatheral, 2006; Overhaus *et al.*, 2007; Wilmott, 2007). We need more convenient models for the asset S able to produce more closely the implied volatility surfaces. Local volatility models, either parametric or non-parametric (see e.g. Cox, 1975; Derman & Kani, 1998; Dupire, 1994; Jäckel, 2008; Rubinstein, 1994), usually capture the surface of implied volatilities more accurately than other approaches, such as stochastic volatility models (see e.g. Ren *et al.*, 2007; Romo, 2012 for discussions). Here we extend previous analysis to a more realistic local volatility-type diffusion, namely the *hyperbolic local volatility* introduced by Jäckel (2008) and widely used in the quantitative finance industry (see e.g. Bompis & Hok, 2014; Hok & Tan, 2019; Hok *et al*., 2018).

The objective of this paper is to compare the application of MC and QMC methods for pricing financial derivatives and computation of Greeks using the hyperbolic local volatility model. The rest of this paper is organized as follows. Section 2 briefly reviews MC and QMC methods. Section 3 introduces the time-homogeneous hyperbolic local volatility model. Time discretization schemes are presented in Section 4. MC simulation of option pricing and computation of Greeks is considered in Section 5. Section 6 presents the results of prices and sensitivities (Greeks) computation. Finally, conclusions and directions of future work are given in Section 6.

2. MC and QMC algorithms

After transformation the option pricing problem can be reduced to the computation of the multidimensional integral

$$I[f] = \int_{H^n} f(x)dx. \tag{1}$$

Here, the function $f(x)$ is integrable in the n-dimensional unit hypercube H^n. The justification is given in Section 5. The MC quadrature formula is based on the probabilistic interpretation of an integral. For a random variable that is uniformly distributed in H^n:

$$I[f] = \mathbb{E}[f(x)], \tag{2}$$

where $\mathbb{E}[f(x)]$ is the mathematical expectation. The standard MC estimator of an expectation is

$$I_N[f] = \frac{1}{N}\sum_{i=1}^{N} f(x_i), \tag{3}$$

where $\{x_i\}$ is a sequence of random points in H^n of length N. The approximation $I_N[f]$ converges to $I[f]$ with probability 1. An integration error, ϵ, according to the Central Limit Theorem, has the expectation

$$\mathbb{E}(\epsilon^2) = \frac{\sigma^2(f)}{N}, \tag{4}$$

where $\sigma^2(f)$ is the function variance. Then, the expression for the root mean square error of the MC method is

$$\epsilon_N = (\mathbb{E}(\epsilon^2))^{\frac{1}{2}} = \frac{\sigma(f)}{\sqrt{N}}. \tag{5}$$

The MC convergence rate does not depend on the number of variables n, but it is rather slow. It is known that random number sampling is prone to clustering. As new points are added randomly, they do not necessarily fill the gaps between already sampled points. However, QMC methods are based on LDS (also known as *quasi-random numbers*). LDS are specifically designed to place sample points as uniformly as possible. Successive LDS points "know" about the position of previously sampled points and "fill" the gaps between them. The QMC algorithm for the evaluation of the integral (2) has a form similar to (3), where instead of a sequence of random points $\{x_i\}$, LDS points $\{q_i\}$ uniformly distributed in a unit hypercube H^n are used: $q_i = (q_i^1, \dots, q_i^n)$.

Many practical studies have proven that the Sobol LDS (Sobol, 1967), is in many aspects superior to other LDS (see e.g. Bianchetti *et al*., 2015; Glasserman, 2013; Jäckel, 2002; Kucherenko & Shah, 2007; Wilmott, 2007). For this reason it was used in this work. Sobol, LDS were constructed by following the three main requirements (Sobol, 1967):

1. Best uniformity of distribution as $N \to +\infty$.
2. Good distribution for fairly small initial sets.
3. A very fast computational algorithm.

Points generated by the Sobol, LDS produce a very uniform filling of the space, even for a rather small number of points N, which is a very important case in practice.

In this work we used the Sobol sequence generator `SobolSeq` provided by BRODA Ltd. (BRODA, 2020).[1] Sobol' sequences produced by `SobolSeq` satisfy additional uniformity properties: Property A for all dimensions and Property A' for adjacent dimensions see Sobol' *et al*., 2011 for details). It has been found that BRODA's `SobolSeq` generator outperforms all other known LDS generators, both in terms of speed and accuracy (Renzitti, 2018; Sobol' *et al*., 2011).

For the best known LDS sequences, the estimate for the rate of convergence $I_N \to I$ is known to be $\frac{O(ln^n N)}{N}$, while for Sobol' LDS $\frac{O(ln^{(n-1)} N)}{N}$ if $N = 2^p$, where p is an integer. This rate of convergence is much faster than that for the MC method (5), although it depends on the dimensionality n. Consequently, the smaller n, the better this estimate. In practice, at $n > 1$, the rate of convergence $\frac{O(ln^n N)}{N}$, is not observed. It appears to be approximately $N^{-\alpha}$, $0 < \alpha \leq 1$, depending on the effective dimension. For financial problems, typically $0.5 < \alpha \leq 1$. Hence, the QMC method usually outperforms MC in terms of convergence. α can be dramatically increased by using effective dimension reduction techniques such as the *Brownian bridge*.

3. Time-homogeneous hyperbolic local volatility model

Since the advent of the Black–Scholes option pricing formula, the study of implied volatility has become a central preoccupation for both academics and practitioners. It is well known that actual option prices rarely conform to the predictions of explicit formulas because the idealized assumptions required for them to hold don't apply in the real world. Consequently, implied volatility (the volatility input to the Black–Scholes formula that generates the market European call or put price) in general depends on the strike K and the maturity of the option T. The collection of all such implied volatilities is known as the volatility surface. The effect that implied volatility $\sigma_{im}(T, K)$ is a decreasing function of strike is called *skew*. Figure 1 provides an illustration for the equity index STOXX50E. The graph shows a strong *skew* for all maturities and this shape is usually observed in the equity derivatives market. This means that the underlying asset price process cannot be explained using the Black–Scholes model, for which the implied volatility does not depend on the strike. Rather, we need to find a convenient model for the underlying asset to evaluate contingent claims. Local volatility models, either parametric or non-parametric (see e.g. Cox, 1975; Derman & Kani, 1998; Dupire, 1994; Jäckel, 2008; Rubinstein, 1994) usually capture the surface of implied volatilities more precisely than other approaches, such as stochastic volatility models (see Ren *et al*., 2007; Romo, 2012 for details).

For our analysis, we consider the time-homogeneous hyperbolic local volatility model (HLV), which is widely used in quantitative finance to capture the market skew. It corresponds to a parametric local volatility-type model in which the dynamic of the underlying risk-neutral measure $\mathbb{Q}$ is

$$dS(t) = rS(t)dt + \tilde{\sigma}(S(t))dW(t),\ S_0 = 1, \tag{6}$$

where r is the risk-free interest rate and

Figure 1: Implied volatilities for different strikes and maturities for STOXX50E on 9/12/2005.

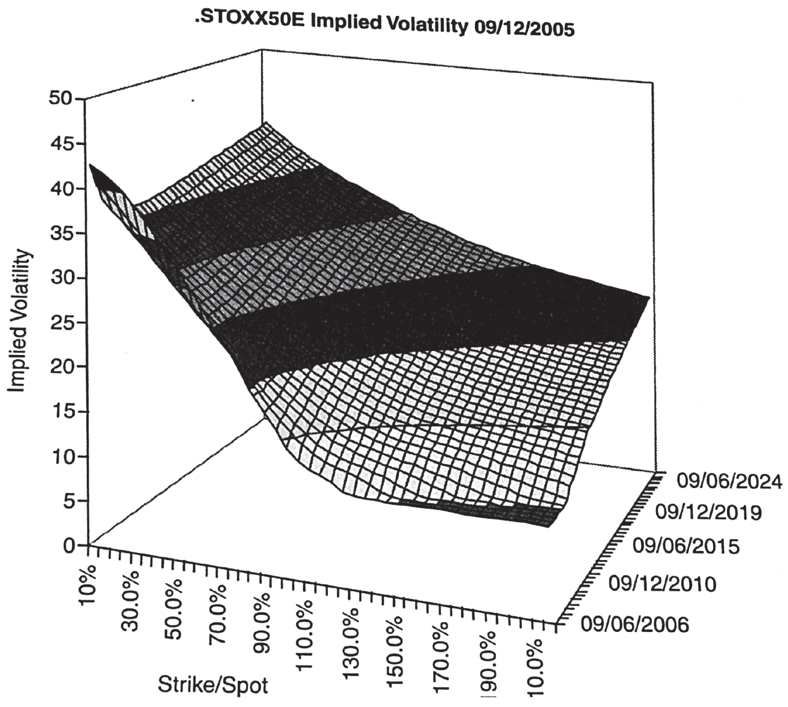

$$\tilde{\sigma}(S) = \nu\left\{ \frac{(1-\beta+\beta^2)}{\beta} S + \frac{(\beta-1)}{\beta}\left(\sqrt{S^2+\beta^2(1-S)^2} - \beta\right) \right\}. \tag{7}$$

Here, $\nu > 0$ is the level of volatility, $\beta \in (0, 1]$ is the skew parameter, and W is the standard Brownian motion. This model was introduced in Jáckel (2009). It behaves similarly to the constant elasticity of variance (CEV) model, and has been used for numerical experiments in Bompis and Hok (2014), Hok and Tan (2019), and Hok *et al.* (2018). The advantage of this model is that zero is not an attainable boundary, and that allows us to avoid some numerical instabilities present in the CEV model when the underlying asset price is close to zero (see e.g. Andersen & Andreasen, 2000). It corresponds to the Black–Scholes model for $\beta = 1$ and exhibits a skew for the implied volatility surface when $\beta \neq 1$. Figure 2 illustrates the impact of the parameter β on the skew of the volatility surface. We observe that the skew increases significantly with decreasing value of β. For example, with $\nu = 0.3$, $\beta = 0.2$, the difference in volatility between strikes at 50 percent and at 100 percent is about 15 percent.

4. Time discretization schemes

4.1. Euler discretization of the stochastic differential equation

Consider the problem of pricing an option on a single asset whose value at time t is denoted by $S(t)$. We assume the asset follows a HLV process defined by the stochastic differential equation (SDE) (6). To guarantee a positive price in the simulation, let's suppose $Y(t) = \ln(S(t))$ and by Ito's formula, we have

$$dY(t) = \left[r - \frac{1}{2}\sigma^2(Y(t))\right]dt + \sigma(Y(t))dW_t,\ Y(0) = \log(S(0)), \tag{8}$$

with $\sigma(Y) = \frac{\tilde{\sigma}(e^Y)}{e^Y}$.

As the solution is not known in closed form, we proceed by using the discretization Euler–Maruyama scheme (Glasserman, 2013; Kloeden & Platen, 2013; Maruyama, 1955). In a discrete case of n equally distributed time steps, it has the following form:

$$\begin{aligned} Y^n(t_{i+1}) = Y^n(t_i) + \left[r - \frac{1}{2}\sigma^2(Y^n(t_i))\right](t_{i+1} - t_i) \\ + \sigma(Y^n(t_i))\sqrt{t_{i+1} - t_i}(W(t_{i+1}) - W(t_i)), \end{aligned} \tag{9}$$

with $Y^n(0) = \log(S(0))$, $\Delta t = \frac{T}{n}$, $t_i = i\Delta t,\ i = 0, \ldots, n$.

In addition to the statistical noise discussed in Section 2, there is also a discretization error associated with a chosen discretization scheme. Theorem 10.2.2 in Kloeden and Platen (2013) provides conditions for the Euler–Maruyama scheme to have a strong error convergence of order $\frac{1}{2}$. Under stronger conditions as in Kloeden and Platen (2013, theorem 14.5.2), theorem 14.5.2, the scheme reaches a weak error convergence of order 1.

4.2. Discretization of the Wiener process

We consider two algorithms for the discretization of the Brownian motion W in Eq. (8). The first one is known as the incremental (standard) discretization algorithm. Its construction follows directly from the definition of $W(t)$. The second one is the alternative discretization algorithm, which is based on the use of conditional distributions.

The standard (incremental) discretization algorithm is defined by the relation

$$W(t_i) = W(t_{i-1}) + \sqrt{\Delta t} Z_i,\quad 1 \le i \le n, \tag{10}$$

where (Z_i) are independent standard normal variates. In the standard discretization algorithm, the evolution of an asset value is generated by normal variates with equal weights.

Figure 2: Impact of the value β on the hyperbolic local volatility for fixed volatility level $\nu = 0.3$.

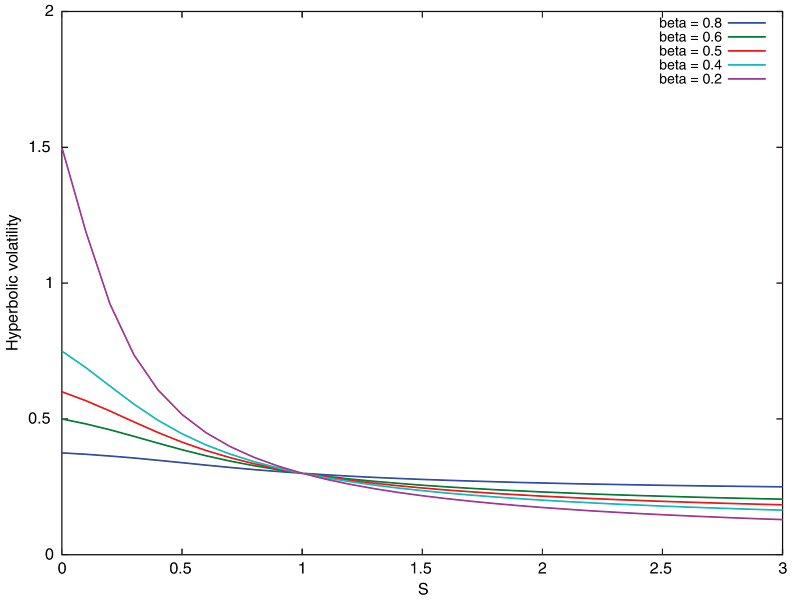

In the Brownian bridge discretization, the value of $W(t_i)$ is generated from values of $W(t_l)$, $W(t_m)$, $l \leq i \leq m$, at earlier and later time steps. Unlike the standard discretization, which generates $W(t_{i+1})$ sequentially along the time horizon, the Brownian bridge discretization first generates the variable at the terminal point

$$W(T) = \sqrt{T}Z_1$$

and then fills other points using already found values of $W(t_i)$. The generalized Brownian bridge formula is given by

$$W(t_i) = (1-\gamma)W(t_l) + \gamma W(t_m) + \sqrt{\gamma(1-\gamma)(m-l)\Delta t}Z_i, \tag{11}$$

where $\gamma = \frac{i-l}{m-l}$ (Morokoff, 1998). It can be seen from Eq. (11) that the variance of the stochastic part of the Brownian bridge formula is $\gamma(1-\gamma)(m-l)\Delta t$. It decreases at successive levels of refinement and the first few points contain most of the variance. This variance is less than that in Eq. (10) as $\gamma(1-\gamma)(m-l) < 1$. Both algorithms have the same variance, hence their MC convergence rates are the same. However, QMC algorithms have different efficiencies, with the Brownian bridge algorithm having a much higher convergence rate (see e.g. Bianchetti *et al.*, 2015; Caflisch *et al.*, 1997; Kucherenko & Shah, 2007; Sobol & Kucherenko, 2005).

Remark 4.1 Standard normal variates are computed as $Z_i \overset{L}{=} \Phi^{-1}(U_i)$, where Φ is the cumulative function of the normal distribution and U_i is a random variables with uniform distribution in [0, 1]. So in practice, one simulates independent uniform random variables and uses this transformation to obtain independent standard Gaussian variables.

5. Monte Carlo simulation of option pricing and computation of Greeks

5.1. Option pricing

We consider a geometric average Asian call option whose payoff function is given by

$$P_A = \max(\overline{S} - K, 0), \tag{12}$$

where $\overline{S}$ is a geometric average at n equally spaced time points:

$$\overline{S} = \left(\prod_{i=1}^{n} S_i\right)^{\frac{1}{n}}, \tag{13}$$

where S_i is the asset price at time $t_i = i\frac{T}{n}$, $1 \leq i \leq n$.

In a risk-neutral environment, the value of a geometric average Asian call option with maturity T and strike K is the discounted value of its payoff:

$$AC(T, K) = e^{-rT}\mathbb{E}^{\mathbb{Q}}[P_A]. \tag{14}$$

In the HLV model, there is no analytical formula for Eq. (14). We are going to estimate the price by the MC method. There are two steps. Firstly, we approximate the asset price $S(t_i)$ with $S^n(t_i) = e^{Y^n(t_i)}$ by discretizing the SDE (8) as described in Section 4.1. Secondly, the MC method approximates the expectation of the Asian payoff (12) with a simple arithmetic average of payoffs taken over a finite number N of simulated price paths:

$$AC_N(T, K) = e^{-rT}\left[\frac{1}{N}\sum_{i=1}^{N} \max(\overline{S}^{(i)} - K, 0)\right], \tag{15}$$

where $\overline{S}^{(i)}$ is an approximation of $\overline{S}$ using the simulated price paths i. So $e^{-rT}\max(\overline{S}^{(i)} - K, 0)$ can be written as $f(U_{i1}, U_{i2}, \ldots, U_{in})$ following Remark 4.1 where all (U_{ij}) are independent uniform variates, which together with (14) justifies formula (1).

5.2. Sensitivity factors

Sensitivity factors or *Greeks* are derivatives of the price $AC(T, K)$ w.r.t. specific parameters like spot price or volatility. They are very important quantities which need to be computed for hedging and risk management purposes. In the present work, we consider in particular the following Greeks:

$$\Delta = \frac{\partial AC(T, K)}{\partial S(0)}, \tag{16}$$

$$\Gamma = \frac{\partial^2 AC(T, K)}{\partial S(0)^2}, \tag{17}$$

$$\vartheta_\nu = \frac{\partial AC(T, K)}{\partial \nu}, \tag{18}$$

$$\vartheta_\beta = \frac{\partial AC(T, K)}{\partial \beta}, \tag{19}$$

called Delta, Gamma, ν-Vega, and β-Vega, respectively. *Delta* represents the hedge of the financial instrument w.r.t. the risky underlying S. In the *dynamic hedging*, it corresponds to the number of assets held, which must be continuously changed to maintain a *delta-neutral* position. Gamma is the second derivative of the price with respect to the underlying. Since gamma is the sensitivity of the delta to the underlying, it is a measure of how much or how often a position must be rehedged in order to maintain a delta-neutral position. Vega represents the sensitivity of the option price to volatility. As discussed in Section 3, ν-Vega measures the price sensitivity to the level of volatility and β-Vega to the volatility skewness. As there are no analytical formulas, the Greeks above are estimated by MC simulation and finite differences, using the central difference formulas

$$\Delta \approx \frac{AC_N(T, K, S(0) + \epsilon_s) - AC_N(T, K, S(0) - \epsilon_s)}{2\epsilon_s}, \tag{20}$$

$$\Gamma \approx \frac{AC_N(T, K, S(0) + \epsilon_s) + AC_N(T, K, S(0) - \epsilon_s) - 2AC_N(T, K, S(0))}{\epsilon_s^2}, \tag{21}$$

$$\vartheta_\nu \approx \frac{AC_N(T, K, \nu + \epsilon_v) - AC_N(T, K, \nu - \epsilon_v)}{2\epsilon_v}, \tag{22}$$

$$\vartheta_\beta \approx \frac{AC_N(T, K, \beta + \epsilon_v) - AC_N(T, K, \beta - \epsilon_v)}{2\epsilon_v}. \tag{23}$$

In the MC simulations for Greeks, we use path recycling of both pseudo-random sequences and LDS to minimize the variance of the Greeks, as suggested in Glasserman (2013) and Jäckel (2002). Notice that the analysis of the root mean square error (RMSE) for Greeks is, in general, more complex than that for prices, since the variance of the MC simulation mixes with the bias due to the approximation of derivatives with finite differences. For the sensitivity factors estimation to be meaningful and not entirely hidden by MC noise, the shifts ϵ_s and ϵ_v are chosen to be large enough and represent about 1 percent of the current spot and the volatility parameters β respectively, (see Glasserman, 2013; Jäckel, 2002 for detailed discussions).

6. Numerical results

In this section we present the numerical results from simulations of prices and sensitivity factors for Asian call options. The following parameters were used for simulation: $S_0 = 100$, $r = 3\%$, $T = 1$, $\nu = 30\%$, $\beta = 0.5$, number of discrete time steps $n = 256$. This corresponds to a high-dimension case. For thoroughness, we consider *in-the-money*, *at-the-money*, and *out-of-the-money* options with strike 80, 100, and 120, respectively.

There are no analytical solutions for geometric average Asian call option prices and sensitivity factors in the HLV model. Numerical simulations using MC and QMC methods were performed to compare convergence of the methods with the standard and the Brownian bridge discretization schemes. The Mersenne Twister generator, which is considered to be one of the most efficient uniform random number generators, was used for MC simulation (Matsumoto & Nishimura, 1998). The Sobol, sequence generator SobolSeq with additional uniformity properties described in Section 2 was used for QMC simulations (BRODA, 2020; Sobol’ *et al.*, 2011). To eliminate randomness which may be caused by the seed point for MC or a specific set of Sobol’ points, we performed $L = 10$ independent runs. For the MC method all runs were statistically independent. For QMC integration for each run a different part of the Sobol, sequence was used. We denote by $Q_N^{(l)}$ a quantity (price or sensitivity factors) on the lth run for N path replications. Each estimated quantity $\overline{Q}_N$ was averaged over $L = 10$ independent runs, i.e.

$$\overline{Q}_N = \frac{1}{L}\sum_{l=1}^{L} Q_N^{(l)}. \tag{24}$$

The reference values were estimated by the Sobol, QMC method with $m = 262,144$ simulation paths and given by

$$\overline{Q}_{ref} = \frac{1}{L}\sum_{l=1}^{L} Q_m^{(l)}. \tag{25}$$

Figures 3 to 7 show the results of simulating an Asian call price and sensitivity factors versus the number of paths obtained using MC with the standard and QMC method with the Brownian bridge discretizations. We note that convergence of MC does not depend on the type of discretization scheme, hence we did not show the results of MC with the Brownian bridge discretization. All results of MC simulations show that simulated solutions slowly converge to the reference solutions, while the convergence curves are highly oscillating. In contrast, QMC with Brownian

Figure 3: Asian call price with strike (a) 80; (b) 100; (c) 120.

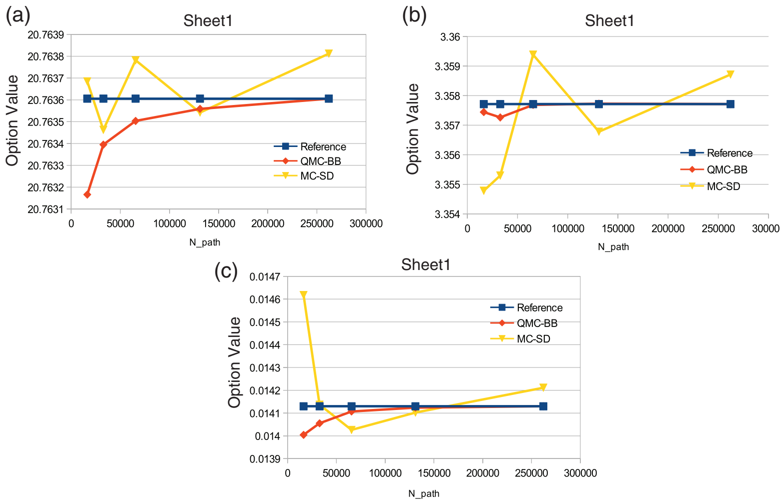

Figure 4: Asian call delta value with strike (a) 80; (b) 100; (c) 120.

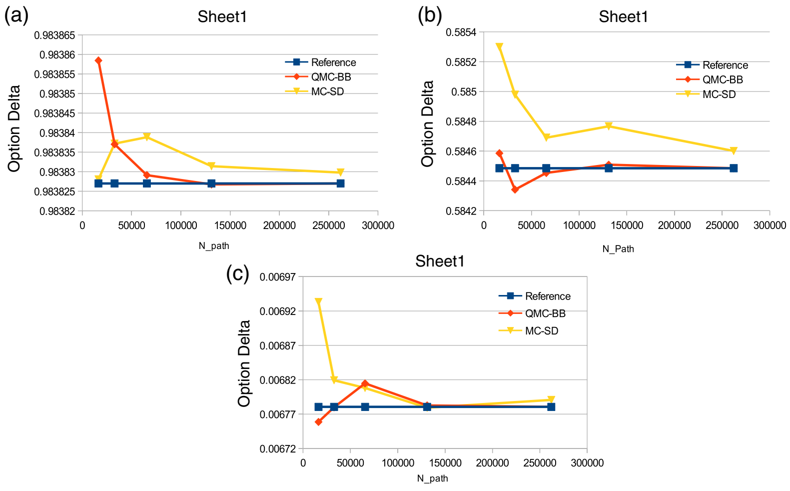

Figure 5: Asian call gamma value with strike (a) 80; (b) 100; (c) 120.

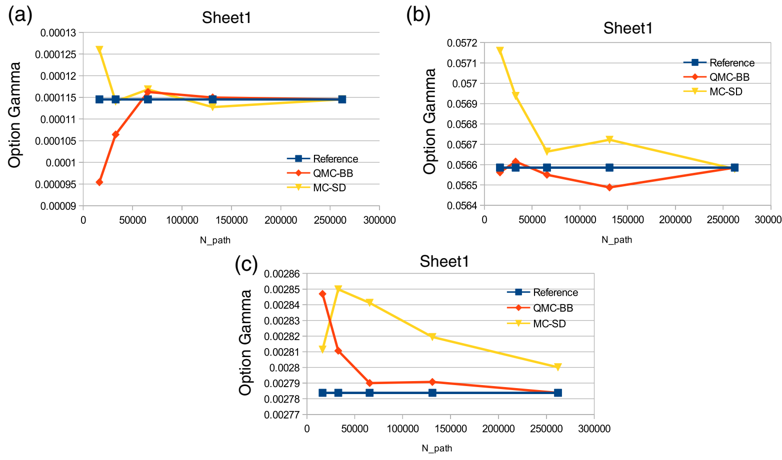

Figure 6: Asian call vega level value with strike (a) 80; (b) 100; (c) 120.

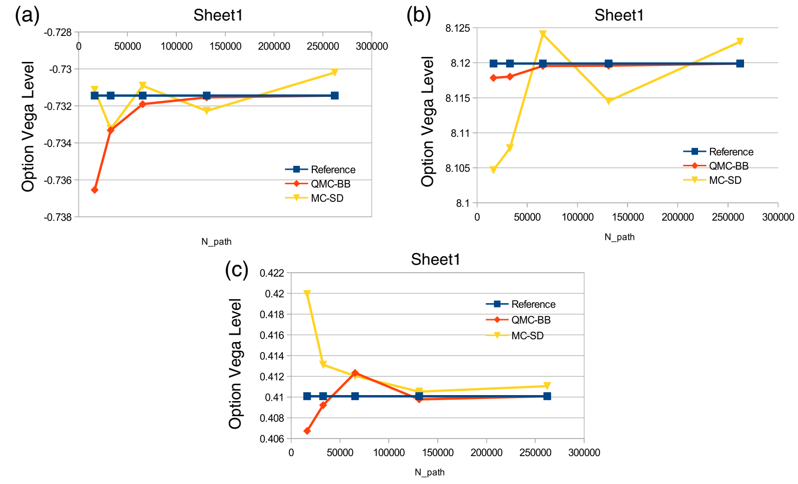

bridge converges much faster and in a more stable way. It is also typically one-sided with a few exceptions. We also notice slight variations in the speed of convergence, depending on whether the option is *in-the-money*, *at-the-money*, or *out-of-the-money*. More efficient convergence behavior changes depending on the type of financial product: price or Greeks, and it is different for different Greeks. For example, for QMC for the call price, the most efficient (fastest) convergence is observed for *at-the-money* or *out-of-the-money* calls, while for delta it is *in-the-money* and *at-the-money*.

We also computed the RMSE for each quantity (price and sensitivity factors), defined as

$$\epsilon = \left(\frac{1}{L}\sum_{l=1}^{L}(\overline{Q}_{ref} - Q_N^{(l)})^2\right)^{\frac{1}{2}}. \qquad (26)$$

As in the previous computations, RMSE was averaged over $L = 10$ independent runs. Figures 8 to 12 show the RMSE versus the number of paths for MC and QMC methods on a log–log scale. We fitted the regression lines $N^{-\alpha}$ to extract convergence rates α. Depending on the strike, the convergence rate for QMC+BB varies

Figure 7: Asian call vega skew value with strike (a) 80; (b) 100; (c) 120.

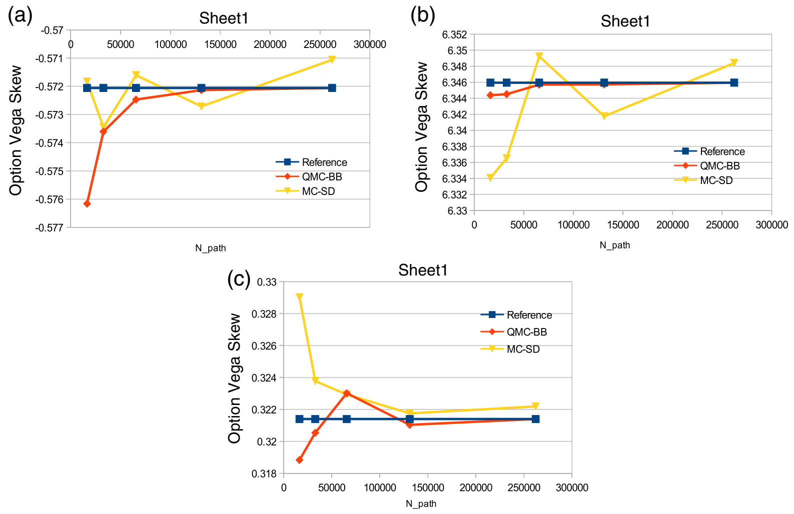

Figure 8: Log–log plot of the RMSE for Asian call price with strike (a) 80; (b) 100; (c) 120.

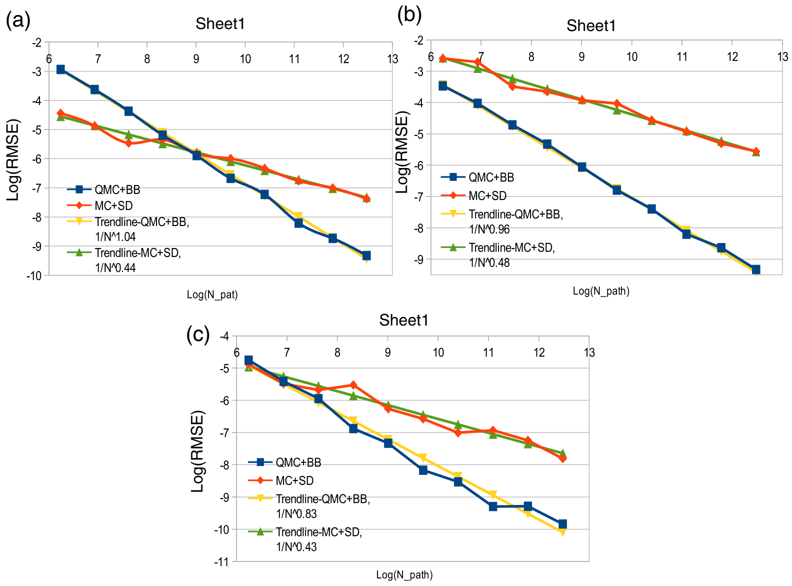

Figure 9: Log–log plot of the RMSE for Asian call delta with strike (a) 80; (b) 100; (c) 120.

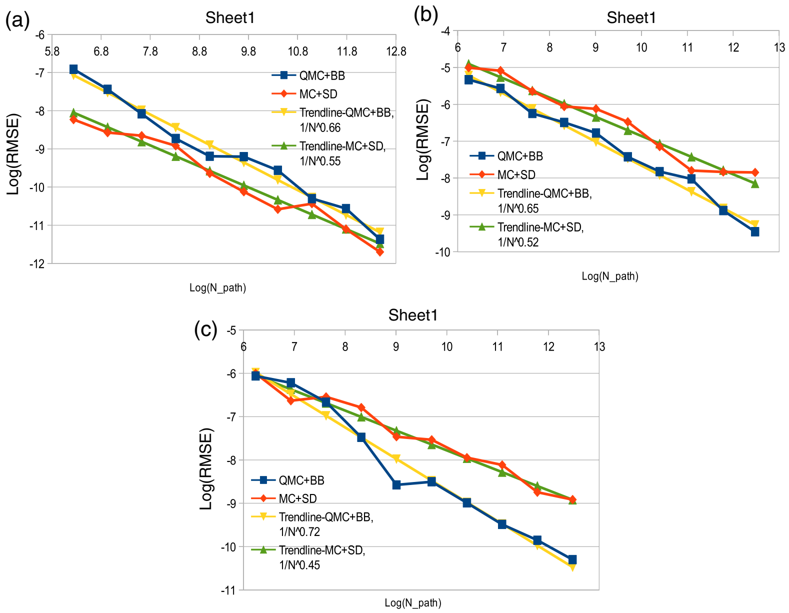

Figure 10: Log–log plot of the RMSE for Asian call gamma with strike (a) 80; (b) 100; (c) 120.

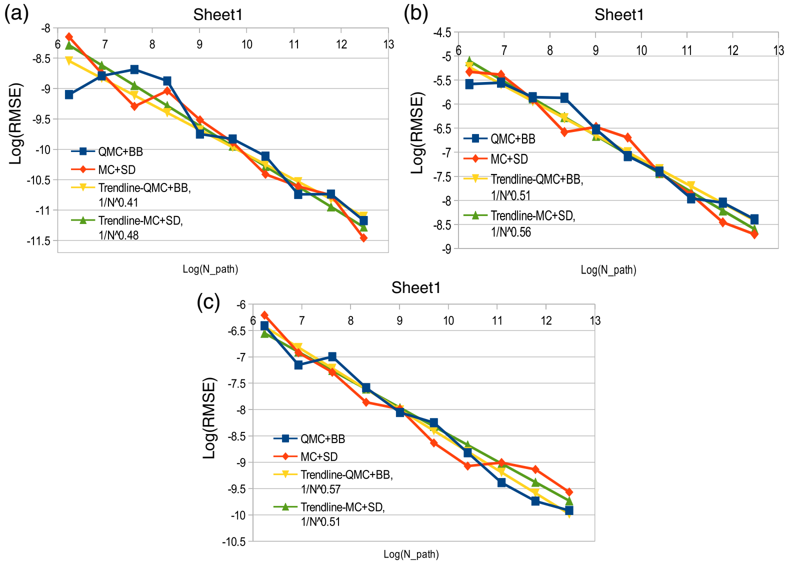

Figure 11: Log–log plot of the RMSE for Asian call vega level with strike (a) 80; (b) 100; (c) 120.

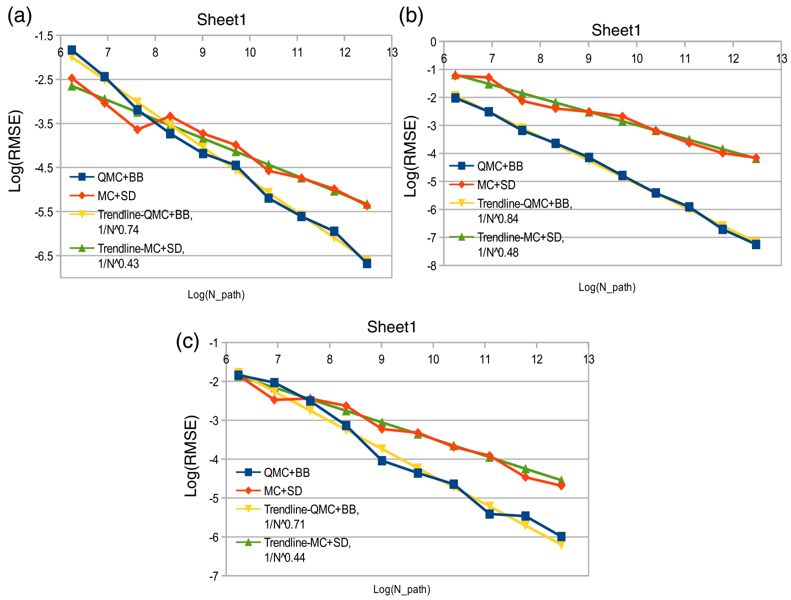

Figure 12: Log–log plot of the RMSE for Asian call vega skew with strike (a) 80; (b) 100; (c) 120.

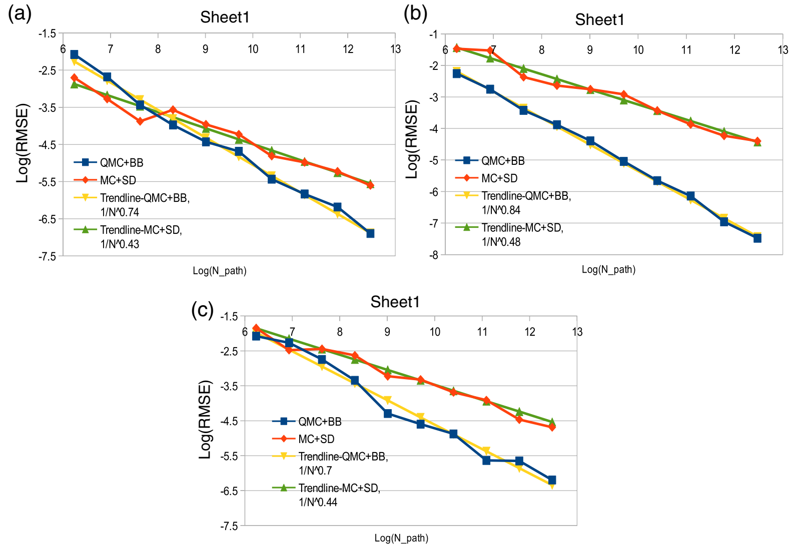

between 0.83–1.0 for price, 0.65–0.72 for Delta, 0.5 for Gamma, and 0.71–0.84 for Vega. For MC+SD this rate is very close to the theoretically predicted limit for MC of 0.5 (5).

7. Conclusions

We present and discuss the results of application of MC and QMC methods for derivative pricing and risk analysis based on the hyperbolic local volatility model. Local volatility models usually capture the surface of implied volatilities more accurately than other approaches. The results presented for the Asian option show the superior performance of the QMC methods, especially for the Brownian bridge discretization scheme. Effective dimensions fully explain the superior efficiency of QMC due to the specifics of Sobol' sequences. The initial coordinates of Sobol' LDS are much better distributed than the later high-dimensional coordinates. The Brownian bridge discretization scheme changes the order in which time steps are sampled. It uses well-distributed coordinates from each n-dimensional LDS vector to determine most of the structure of a path, and reserves other coordinates to fill in finer details. This results in a reduction of the effective dimensions and significantly improved accuracy of the QMC method.

Julien Hok holds a PhD in financial mathematics from Ecole Polytechnique France. He started as a quantitative analyst in equity at Santander in London for six years and worked at Citi Group for two years at London in interest rates. After that he joined CA-CIB as quantitative analyst in the hybrid desk at London for four years. Recently he has joined INVESTEC Bank as a quantitative analyst for the Equity derivatives desk at London.

Sergei Kucherenko received his MSc degree and PhD in applied mathematical physics from Moscow Engineering Physics Institute in Russia. He held a number of research and faculty positions in various universities in Russia, the United States, the UK and Italy. He also worked in investment banking. Currently he holds a position of Senior Research Fellow at Imperial College London. He is also a founder of BRODA Ltd. which provides consultancy services to investment banks and financial companies in the area of MC and Quasi MC simulation and other advanced numerical techniques used in quantitative finance.

Endnotes

1. The most recent generator released by BRODA is `SobolSeq65536` with maximum dimensionality 65,536 and Properties A and A′.

References

[1] Acworth, P. A., Broadie, M., and Glasserman, P. 1998. A comparison of some Monte Carlo and quasi Monte Carlo techniques for option pricing. In H. Niederreiter, P. Hellekalek, G. Larcher, and P. Zinterhof (eds), *Monte Carlo and Quasi-Monte Carlo Methods 1996* New York:Springer; 1–18.
[2] Åkesson, F. and Lehoczky, J. P. 2000. Path generation for quasi Monte Carlo simulation of mortgage-backed securities. *Management Science* 46(9), 1171–1187.
[3] Andersen, L. and Andreasen, J. 2000. Volatility skews and extensions of the Libor market model. *Applied Mathematical Finance* 7(1), 1–32.
[4] Bianchetti, M., Kucherenko, S., and Scoleri, S. 2015. Pricing and risk management with high-dimensional quasi Monte Carlo and global sensitivity analysis. *Wilmott* 2015(78), 46–70.
[5] Bompis, R. and Hok, J. 2014. Forward implied volatility expansion in time-dependent local volatility models. *ESAIM: Proceedings and Surveys* 45, 88–97.
[6] BRODA 2020. High-dimensional Sobol' sequence generators. *http://www.broda.co.uk/*.
[7] Caflisch, R. E., Morokoff, W. J., and Owen, A. B. 1997. Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension. *The Journal of Computational Finance* (1), 27–46.
[8] Cox, J. 1975. Notes on option pricing I: Constant elasticity of diffusions. Working paper.
[9] Derman, E. and Kani, I. 1998. Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility. *International Journal of Theoretical and Applied Finance*, 1(01), 61–110.
[10] Dupire, B. 1994. Pricing with a smile. *Risk*, 7(1), 18–20.

[11] Gatheral, J. 2006. *The Volatility Surface: A Practitioner's Guide*. Chichester: Wiley Finance.
[12] Glasserman, P. 2013. *Monte Carlo Methods in Financial Engineering*, vol. 53. New York: Springer Science & Business Media.
[13] Hok, J. and Tan, S.-H. 2019. Calibration of local volatility model with stochastic interest rates by efficient numerical PDE methods. *Decisions in Economics and Finance* 42(2), 609–637.
[14] Hok, J., Ngare, P., and Papapantoleon, A. 2018. Expansion formulas for European quanto options in a local volatility FX-Libor model. *International Journal of Theoretical and Applied Finance* 21(02).
[15] Jäckel, P. 2002. *Monte Carlo Methods in Finance*, vol. 71. Chichester: Wiley.
[16] Jäckel, P. 2008. Hyperbolic local volatility. *http://www.jaeckel.org/HyperbolicLocalVolatility.pdf*.
[17] Jäckel, P. 2009. Quanto skew. *http://www.jaeckel.org/QuantoSkew.pdf*.
[18] Kloeden, P. E. and Platen, E. 2013. *Numerical Solution of Stochastic Differential Equations*, vol. 23. New York: Springer Science & Business Media.
[19] Kucherenko, S. and Shah, N. 2007. The importance of being global. Application of global sensitivity analysis in Monte Carlo option pricing. *Wilmott Magazine* 4, 2–10.
[20] Maruyama, G. 1955. Continuous Markov processes and stochastic equations. *Rendiconti del Circolo Matematico di Palermo* 4(1), 48–90.
[21] Matsumoto, M. and Nishimura, T. 1998. Mersenne Twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 8(1), 3–30.
[22] Morokoff, W. J. 1998. Generating quasi-random paths for stochastic processes. *Siam Review* 40(4), 765–788.
[23] Moskowitz, B. and Caflisch, R. E. 1996. Smoothness and dimension reduction in quasi-Monte Carlo methods. *Mathematical and Computer Modelling* 23(8–9), 37–54.
[24] Overhaus, M., Bermúdez, A., Buehler, H., Ferraris, A., Jordinson, C., and Lamnouar, A. 2007. *Equity Hybrid Derivatives*, vol. 374. Wiley Online Library.
[25] Ren, Y., Madan, D., and Qian, M. 2007. Calibrating and pricing with embedded local volatility models. *Risk Magazine*, 20(9), 138.
[26] Renzitti, S. 2018. Accelerating CVA and CVA sensitivities using quasi Monte Carlo methods. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3193219.
[27] Romo, J. M. 2012. The quanto adjustment and the smile. *Journal of Futures Markets* 9(32), 877–908.
[28] Rubinstein, M. 1994. Implied binomial trees. *The journal of finance* 49(3), 771–818.
[29] Sobol, I. M. 1967. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics* 7(4), 86–112.
[30] Sobol, I. M. and Kucherenko, S. 2005. On global sensitivity analysis of quasi-Monte Carlo algorithms. *Monte Carlo Methods and Applications* 11(1), 83–92.
[31] Sobol', I. M., Asotsky, D., Kreinin, A., and Kucherenko, S. 2011. Construction and comparison of high-dimensional Sobol' generators. *Wilmott* 2011(56):64–79.
[32] Wilmott, P. 2007. *Paul Wilmott Introduces Quantitative Finance*. Hoboken, NJ: Wiley.